



---

## B2STAGE

### About

Document that describes the B2STAGE data staging EUDAT service.

**Modified:** 28 February 2018

### Synopsis

The B2STAGE service allows data staging, i.e. data transfer into and out of EUDAT data nodes. Data staged into EUDAT are assigned a unique Persistent Identified (PID). EUDAT exposes two protocols for staging data, as follows:

- GridFTP (via the EUDAT Data Storage Interface) is aimed at large data transfer and numerous files. It allows for third party transfers. The target group are power users who need to access data in B2SAFE and move them to compute sites.
- HTTP is for small and medium files. The HTTP API allows for access to B2SAFE metadata. The target group are community developers who want to integrate data and features from B2SAFE into their community-specific applications.

The current B2STAGE implementations are only limited to files managed by the EUDAT B2SAFE service.

### Introduction

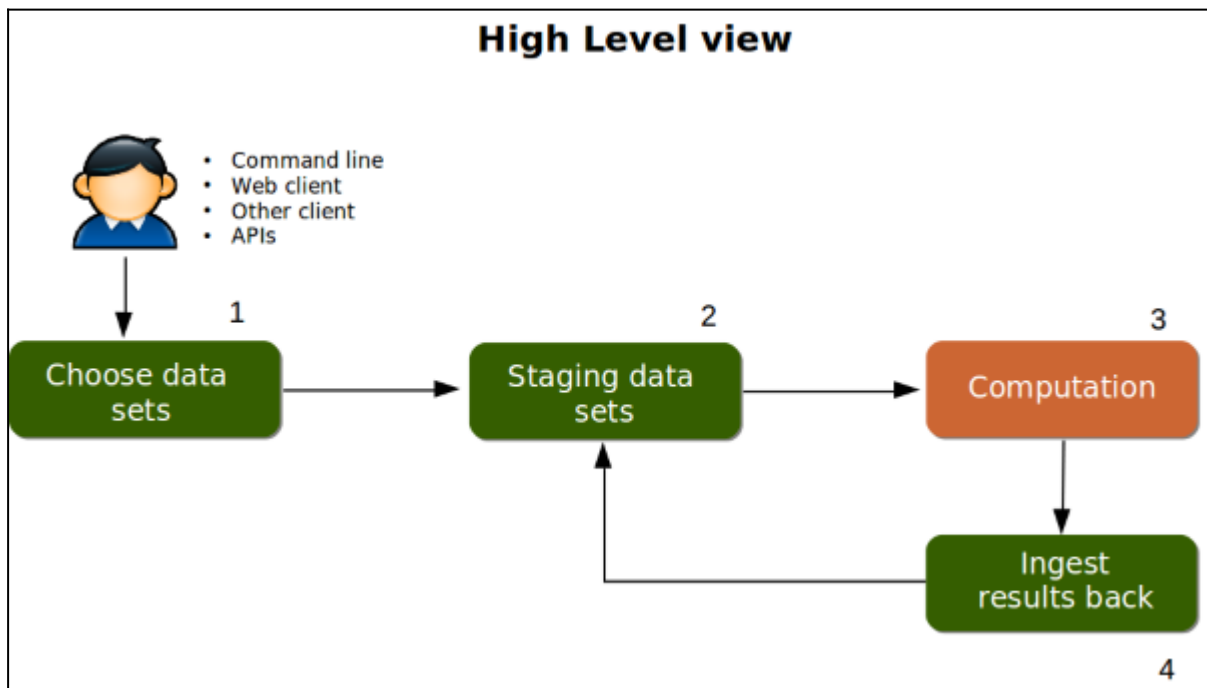
Modern Big Data computing involves a large amount of data. When many HPC centers are involved, data exchange to/from one center must be taken into account. Another scenario for data movement is when data needs to be moved between different (EUDAT and community) applications or to HPC centres for data processing. In general, data movement arises in order to satisfy at least one of the following needs:

1. data preservation and access optimizations;
2. dynamic replication to HPC workspace for processing;
3. dynamic access to data for integration into community-specific applications

In EUDAT, the first group is addressed by the **B2SAFE (Safe Replication)** service, the second by the **B2STAGE gridFTP (Data Staging)** service and the third one by the **B2STAGE HTTP API**.

In EUDAT, the key functionality of the B2STAGE gridFTP Service is transferring relevant data sets between HPC centers and EUDAT in order to store them, process them and, possibly, move the results back. Data could also be already stored in one or more EUDAT data centers, as result of the Safe Replication activity. Eventually, output data are identified through a Persistent Identifier (PID).

With the HTTP API EUDAT offers programmatic access to data in B2SAFE and thus allows for smooth integration of such data into other applications and data services.



**Figure 1. Typical user data workflow of EUDAT B2STAGE gridFTP.**

Figure 1 depicts the workflow of B2STAGE gridFTP. In general, it involves the following four steps:

1. The user chooses which data sets they want to move. It is possible to identify them using the PID (see below).
2. The user moves the data sets. It is possible to move data sets into EUDAT as well as from EUDAT. It is possible to move data sets to/from the user's desktop as well as to/from an HPC centre.
3. Possibly, the user uses the data sets to calculate new data sets, and
4. the new data sets could be ingested into EUDAT again.

The same workflow holds true for applications relying on data in EUDAT B2SAFE which process the data on behalf of the user. In this case one would employ the HTTP API.

Of course user communities and users in general move data in many different ways outside B2STAGE; B2STAGE addresses, in EUDAT, the problem of how to **efficiently move a large amount of data**. There are three main aspects to consider for this service:

1. performance;
2. reliability;
3. security.

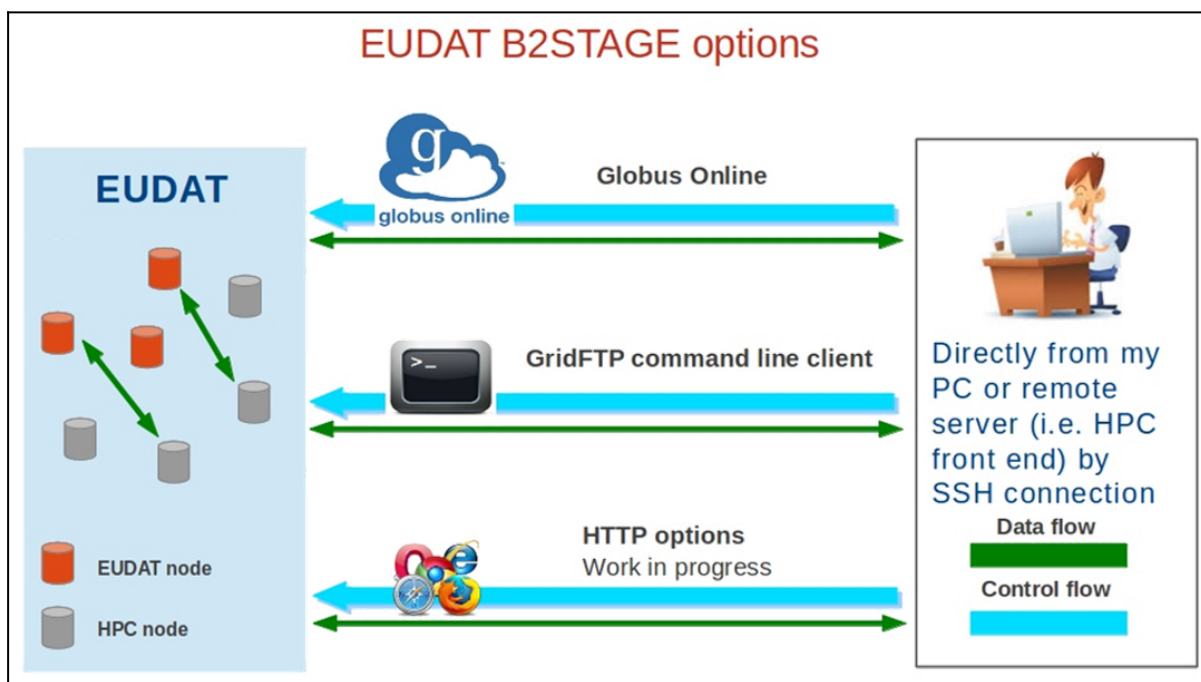
Several protocols target these aspects, including SCP/SFTP, HPSS and so on. However, these were designed for a LAN context in which latency is not as important as in the WAN context prevalent in EUDAT and tackled by B2STAGE. Furthermore, performance is not the only aspect to be considered: a tool able to transfer large amount of data must be easy to use and highly reliable; moreover, it could offer third-party transfer and the possibility to control/limit the transfer throughput to avoid saturating the network. To meet all these requests, **EUDAT supports** the following:

- a solution based on **GridFTP**, the *de facto* standard for high-performance, secure, reliable data-transfer in

the HPC community and which allows for third party transfers.

- a mainly RESTful **HTTP API** interface, with well-defined states, for lighter use and access to B2SAFE metadata.

Please see below for a comparison between the two.



**Figure 2. EUDAT B2STAGE options.**

With reference to Figure 2, it is important to observe that B2STAGE gridFTP allows two different flavours of data transfer: *third-party transfer*, i.e. a transfer between an EUDAT node and a non-EUDAT node (such as an HPC farm) instrumented from the user's PC; and *client-server transfer*, i.e. a direct transfer between an EUDAT node and the user's PC or, via SSH, the user's login node on an HPC farm. The HTTP method only allows client-server connections; this is a feature of the technology.

## B2STAGE HTTP API versus B2STAGE gridFTP

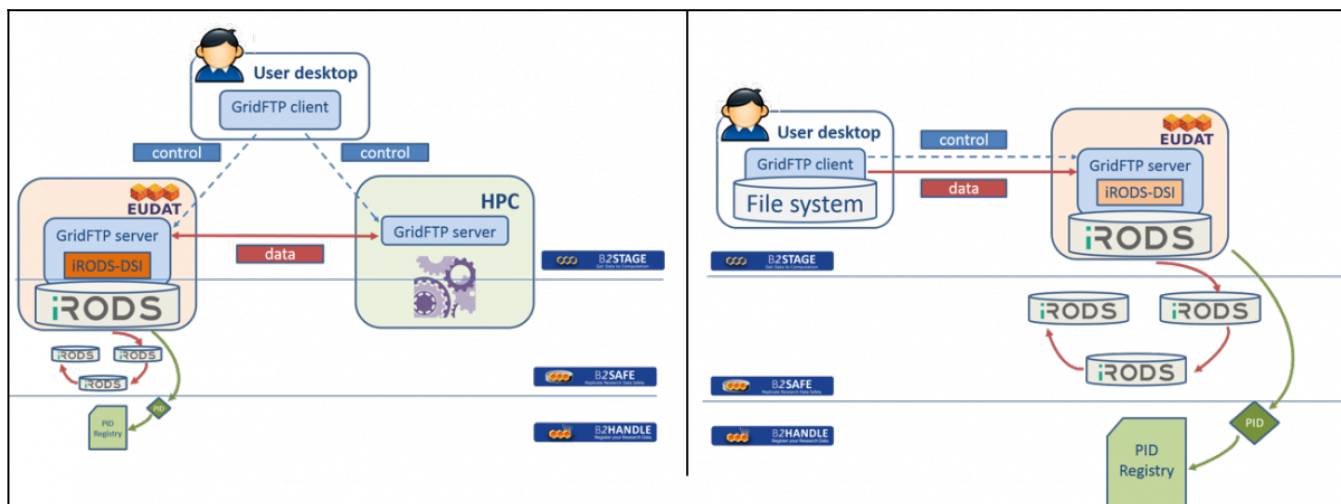
Users accessing data in B2SAFE by B2STAGE gridFTP need to have access to a *gridFTP client*, while for using the HTTP API no special clients are needed, a simple web-browser or curl command line tool suffice.

B2STAGE gridFTP is suitable for transfers of *large data files or many data files*, while the HTTP API can only support *small to medium-sized data files*. However, B2STAGE gridFTP does not support access to metadata, while the HTTP API users can read *iRODS metadata* created by the B2SAFE service.

## How B2STAGE works

The B2STAGE service is currently integrated only with the B2SAFE service, and thus couples efficient transfer into EUDAT with persistent identification and Safe Replication. The B2STAGE service integrates GridFTP and HTTP access with the iRODS technology. EUDAT also supports client-side tools to ease the user's B2STAGE GridFTP experience. This is depicted in Figure 3. When data arrive at an EUDAT node to be deposited, the B2SAFE service

ensures that a PID is generated by B2HANDLE for each artefact, and this is recorded in the EPIC PID Register. The iRODS Server also handles any replication required for these artefacts, according to the community policies that apply to the user who initiated the transfer. These apply to both third-party (shown on the left of Figure 3) and client-server transfer (shown on the right of Figure 3).



**Figure 3. Third party B2STAGE gridFTP. On the left, the user controls data flow between EUDAT and an HPC centre. On the right, the user controls data flow between their desktop and EUDAT.**

B2STAGE gridFTP plays the role of an interface in front of iRODS to perform high-speed and or large data transfer into EUDAT. The DSI component in Figure 3 was developed for the GridFTP-iRODS communication. DSI was developed in accordance to the [GridFTP Data Storage Interface](#) specifications which allow a GridFTP server to be a transfer interface to numerous data storage systems, including iRODS. Persistent identification is handled by the Handle system. The client and the server components of B2STAGE are discussed briefly below, with the emphasis on the client component.

## The B2STAGE servers

Both flavours of B2STAGE are implemented as server-side extensions (see Figure 3 for B2STAGE gridFTP), i.e. both need to be deployed by or in close collaboration with the system administrator of iRODS/B2SAFE.

The GridFTP data staging functionality of B2STAGE is realized by extending the iRODS system with a GridFTP interface, implemented by the EUDAT DSI Component. This permits the transfer of data through a reliable, high-performance protocol. Information for EUDAT systems administrators to deploy DSI is available from the [EUDAT User Documentation](#) site.

The HTTP API consists of several components. The main component is Rapydo which employs an ecosystem of docker containers to install and manage infrastructure components for the HTTP API. Rapydo controls a docker container which runs a Flask server and an nginx server. The Flask server provides the actual implementation of the HTTP API while nginx provides the HTTP server which exposes the API to the users. The Flask framework interacts with B2ACCESS for Authentication and authorisation and B2HANDLE for persistent identifier resolving. Furthermore, it interacts through a python API with iRODS and B2SAFE. Finally, Swagger, a definition of the HTTP API, is used as an HTTP API web frontend which offers a graphical way to explore the HTTP API interactively. Documentation for the server-side deployment of the HTTP functionality of B2STAGE is under development; previews are available the [on GitHub](#).



---

## Clients for B2STAGE GridFTP and Access

Any existing client supporting the GridFTP protocol can be employed for B2STAGE, including globus-url-copy, Globus On Line, UberFTP, gTransfer, etc. Among the available clients there is also the [EUDAT File Manager](#) which supports a range of transfer protocols (i.e. GridFTP, FTP, native iRODS, etc.) and provides an intuitive and user-friendly interface.

B2STAGE is tightly integrated with B2SAFE/iRODS, so users will first need an account for B2SAFE/iRODS. To make use of the B2STAGE gridFTP endpoints, users also need a personal certificate (X.509) to access the service. These certificates are issued by a certificate authority. The administrator of the B2SAFE/iRODS and B2STAGE instance can point you to the respective certificate authority.

## Clients for B2STAGE HTTP API and Access

The HTTP API knows two ways of authenticating users 1) via B2ACCESS and 2) via iRODS i.e. users known in the iRODS iCAT database. Both types of users can request an API token and make use of the HTTP API.

There are no special clients needed. Most commands can be issued with the standard curl command line tool or even via the web browser. For better understanding of the HTTP API, the installation comes with a Swagger interface, a graphical representation of the API which also allows to issue some commands.

## A practical case for B2STAGE GridFTP

Considering a use case in which a user decides to perform some computation on the data sets he has stored in EUDAT (see step 3 of Figure 1), in short the following steps are necessary:

1. get an account on the EUDAT node you want to access;
2. (optionally) get an account on the HPC farm in order to run some code;
3. (optionally) get a Globus Online account;
4. get an X.509 certificate;
5. associate your certificate to all the accounts listed before;
6. get a B2STAGE client tool (e.g. GridFTP command line client, Globus Online and UberFTP);
7. (optionally) install GlobusConnect on your PC.

Of course, in case you need to obtain remote accounts you should refer to the specific guidelines provided by the remote sites (this does not concerns EUDAT directly).

To move on to a practical use-case, consider the case of a user wanting to transfer data into the EUDAT node at the CINECA supercomputing center. Here are the steps to be performed:

1. Obtain x.509 certificate ([see here](#))
2. Create and activate Globus online account ([see here](#))
3. Use GridFTP on Fermi and PLX farm in CINECA ([see here](#))
4. Add end point in Globus Online ([see here](#))
  - (EUDAT node in CINECA is `data.repo.cineca.it:2811`)

At this point you are able to transfer files from/to the CINECA HPC farm (FERMI or PLX) and the EUDAT node in CINECA using Globus Online or equivalent tools such as globus-url-copy. If you also want to transfer data to/from your PC you need to install Globus Connect on it ([see here](#)).



---

## EGI Use of B2STAGE

EGI has adopted B2STAGE. You can browse their [documentation](#) for hints and highlights of their the use and adaptations of the service.

## Support

Our [B2STAGE presentations](#) discuss how to deploy the service as a data manager, and also how to use it as an end-user.

You can access B2STAGE hands-on training material from our [github](#); note in particular:

- data managers: the sessions dedicated to [installing a gridFTP server](#), [installing B2STAGE](#) and [installing a \(test\) environment for the HTTP API for the B2SAFE service](#) (Module 11).
- end-users: [using B2STAGE GridFTP](#) and [using B2STAGE HTTP API](#).

*Support* for B2STAGE is available via the EUDAT ticketing system through the [webform](#).

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

## Document Data

**Version:** 2.0

**Authors:**

Christine Staiger, [christine.staiger@surfsara.nl](mailto:christine.staiger@surfsara.nl)

Giovanni Morelli, [g.morelli@ Cineca.it](mailto:g.morelli@ Cineca.it)

Giacomo Mariani, [g.mariani@ Cineca.it](mailto:g.mariani@ Cineca.it)

**Editors:**

Kostas Kavoussanakis, [k.kavoussanakis@epcc.ed.ac.uk](mailto:k.kavoussanakis@epcc.ed.ac.uk)

Carl Johan Håkansson, [cjhak@kth.se](mailto:cjhak@kth.se)

Sri Harsha Vathsavayi, [sriharsha.vathsavayi@csc.fi](mailto:sriharsha.vathsavayi@csc.fi)

[Read more](#)