



Contribution ID: 189

Type: poster presentation

Seamless access to HTTP/WebDAV distributed storage: the LHCb storage federation case study and prototype

In this contribution we describe the activities and the technical aspects that led to the construction of a public prototype for LHCb file access that is built on HTTP and WebDAV, supporting file access for distributed computing data management and data processing activities as well as seamless interactive access via web browsers.

The LHCb replica naming scheme provides characteristics that makes it an excellent candidate for evaluating the kind of interaction that an HTTP deployment can give on top of an existing, already running computing model.

Deploying HTTP access in the context of LHCb gives the possibility of accessing all 19 individual experiment storage areas, with 5 different disk-only and tape storage technologies deployed, using web tools like browsers, curl, wget and the Davix client.

These options allow a wide range of applications frameworks such as ROOT to access data, using techniques that may range from simple, single file transfers to vectored reads. Coordinating the deployment of WebDAV access in the sites has been one of the main efforts, and is still ongoing.

Given the very high level of coherency of the LHCb site namespaces, the quick pace at which sites joined the effort, and the desire to quickly setup an ambitious prototype, we decided to use the already existing storage federation testbed that is run at DESY, by configuring it with the HTTP/WebDAV endpoints of the LHCb storage elements as they joined the exercise.

We obtained a Dynamic Federations instance that is already covering a good part of the LHCb data namespace. The front-end shows an unified repository that is composed by the content of all the sites that it aggregates. The aggregation is performed on the fly, by distributing quick WebDAV queries towards the endpoints and scheduling, aggregating and caching the responses. We will further show how the HTTP protocol access will be incorporated into the LHCb/DIRAC distributed computing tool and can be beneficial for individual end users.

Primary authors: DEVRESSE, Adrien (CERN); FURANO, Fabrizio (CERN); ROISER, Stefan (CERN)

Presenter: ROISER, Stefan (CERN)

Track Classification: Track3: Data store and access