

DAVIX: a client library for WebDAV and HTTP-derived protocols oriented to high performance data and metadata access

Tuesday 12 March 2013 11:40 (20 minutes)

HTTP-based protocols for data and metadata access are becoming more and more important in the "big data" application domain. HTTP is now the "de facto" standard protocol underlying several big cloud storage systems (Amazon S3, SWIFT, WebDAV) and is now commonly accepted in the grid storage world (lcgdm-dav WebDAV, dCache Jetty DAV, WebDAV for Storm).

On top of that, HEP computing (and in general, HPC) needs several features that can technically be achieved with HTTP-based protocols, and that can be at the same time not very easy to design and implement.

Examples of such features include a broad choice of authentication mechanisms (GSI proxies, x509 proxies, passwords, S3 Cloud tokens), support for them in WebDAV, support of the LCGDM extensions for data management, support of session re-use, transparent support of the HTTP redirections in all the primitives, support for vectored IO, metalinks, failover, thread safety and many others.

To the best of our knowledge, finding a third-party client (open source or not) that supports all this is very difficult, as publicly available clients tend to support various permutations of the features, not them all. There are also other examples of client libraries that are good quality implementations of just a low level access layer.

We think that users should be presented a coherent interface that "just works" and that shields them from having to deal with parsing complex XML responses or with security libraries or with too many options for the supposedly simple task of interacting with a data repository.

For these reasons when we needed a complete client to implement parts of the design of our "federated HTTP ecosystem" we decided to use libneon to write a complete and performant client that exposes a simple POSIX-like (open/read/write/close) API that supports all the aforesaid features. We named this client DAVIX.

Now DAVIX is used in the Dynamic Federations project, in GFAL2 and in FTS3. Given its generality, we think that many other applications (e.g ROOT I/O) would get a big benefit using it, exploiting the many data-access related features that it has and its very broad degree of compatibility.

Author: FURANO, Fabrizio (CERN)

Presenter: FURANO, Fabrizio (CERN)