



Web enabled data

management with DPM & LFC



Alejandro Álvarez Ayllón, Alexandre Beche, Fabrizio Furano, Martin Hellmich, Oliver Keeble, Ricardo Brito Da Rocha
Special thanks to Academia Sinica Grid Computing (Taiwan) for validation and testing.

WebDAV

allows clients to access a web server as if it was a filesystem, **extending HTTP** to provide a framework handling **metadata** through the web.

HTTP/S

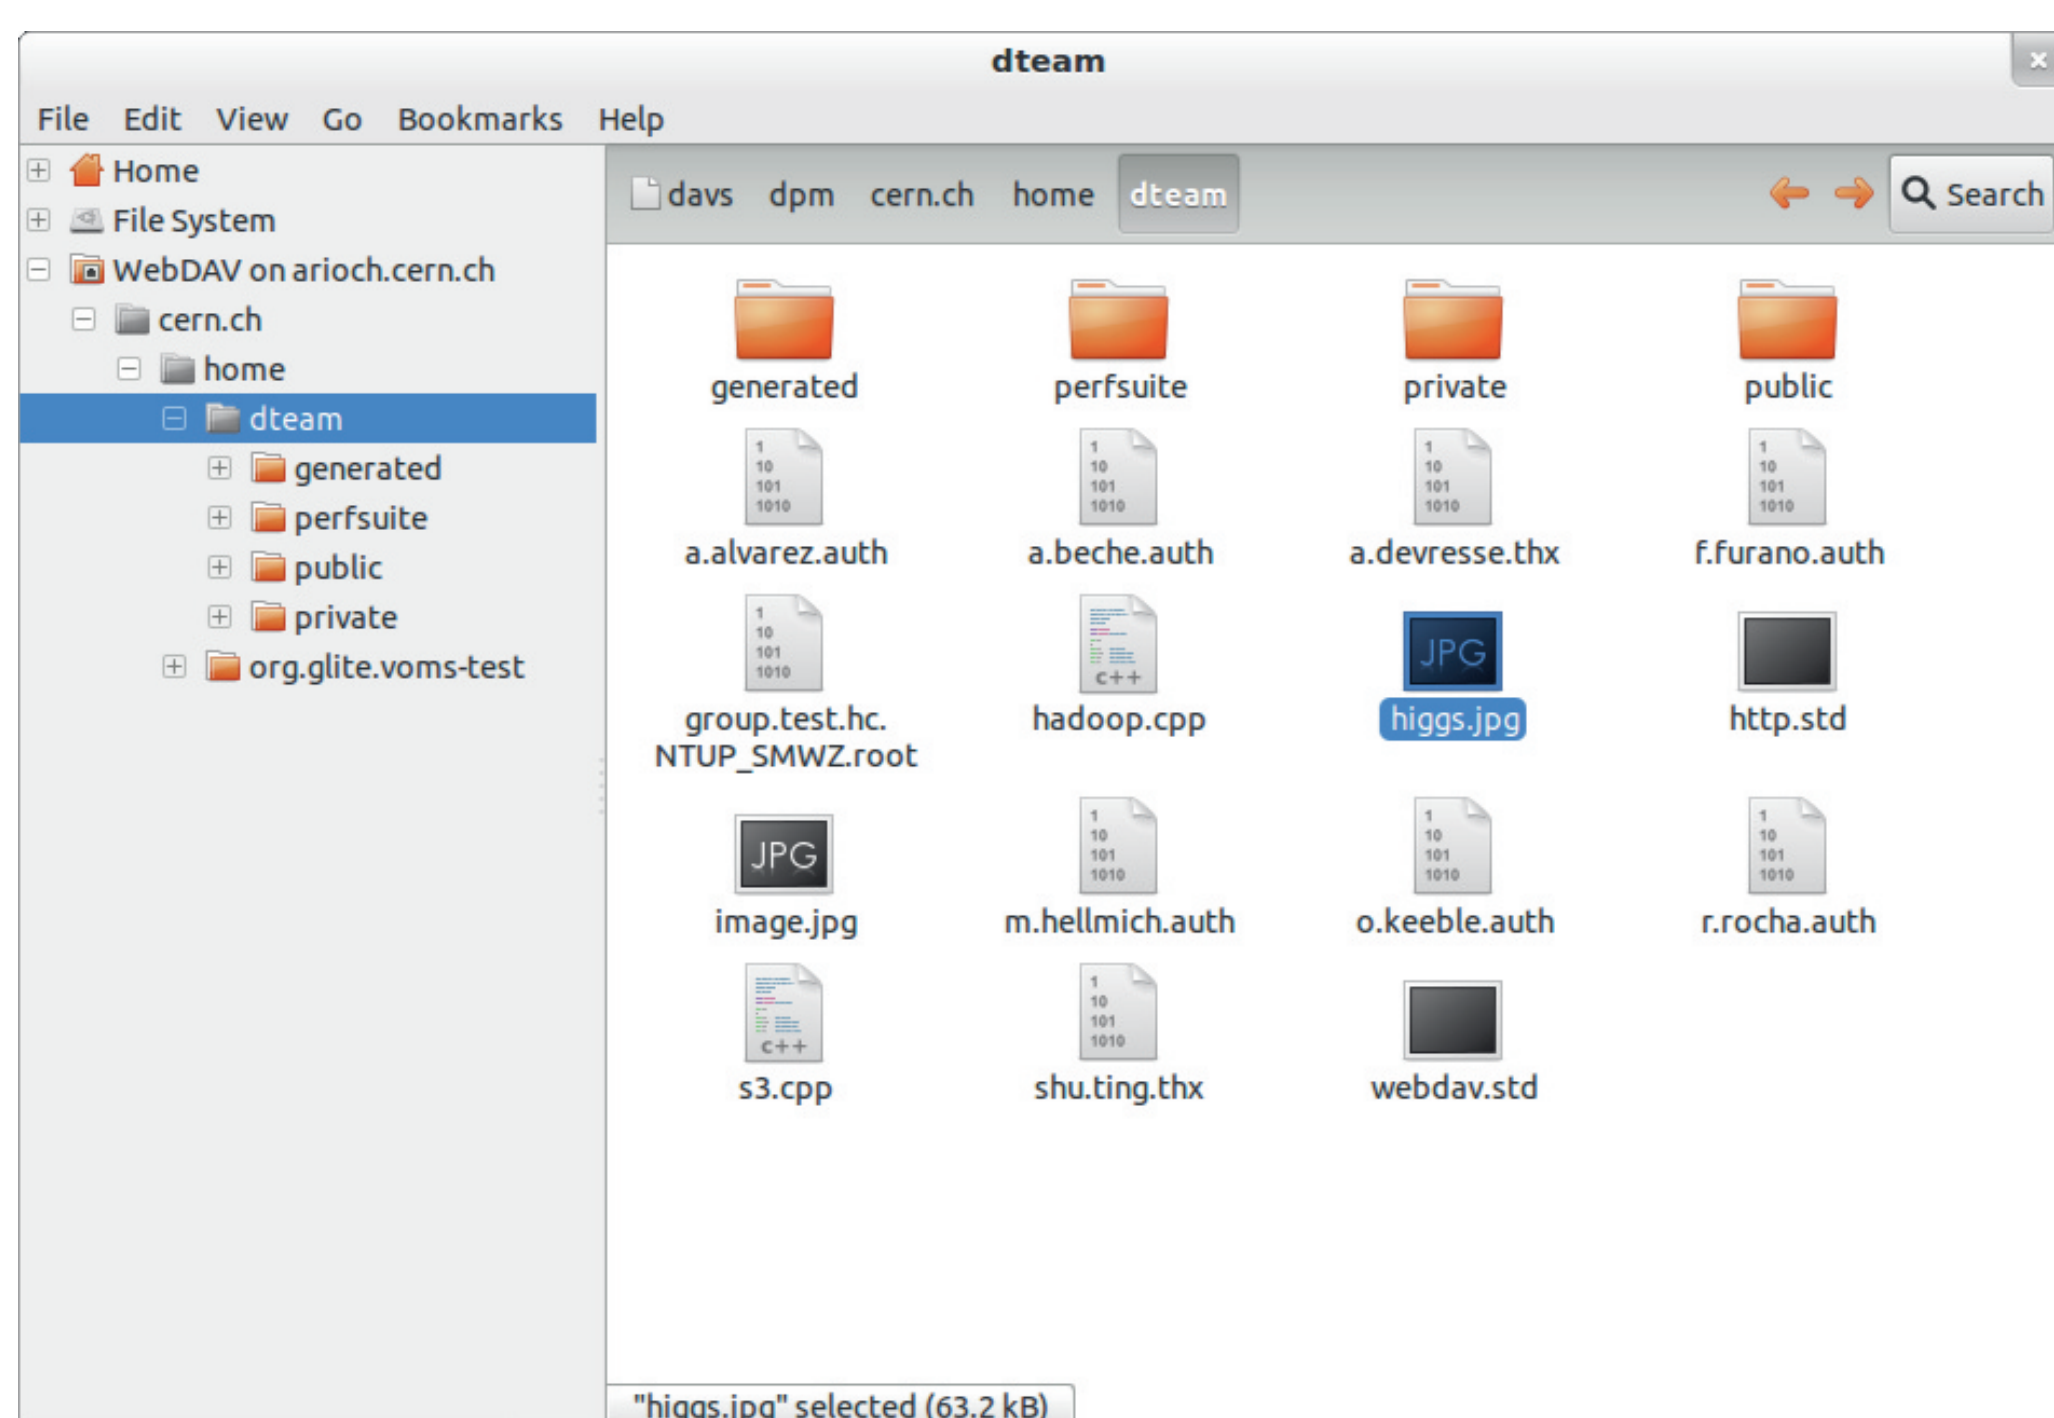
is the most well known and widely implemented web protocol. It can be used as an **access** and **transfer protocol**. Encryption and authentication can be used through SSL (HTTPS).

Objectives

Provide a **standards based** solution for data and metadata access on the grid, as an **alternative** to existing protocols - GridFTP for transfers and the custom protocol for DPM/LFC metadata queries - allowing the access from almost any system out-of-the-box. For transfer, we describe "Multistreams" and "3rd Party Copies", as they are key features of GridFTP.

Ease of use: WebDAV for metadata

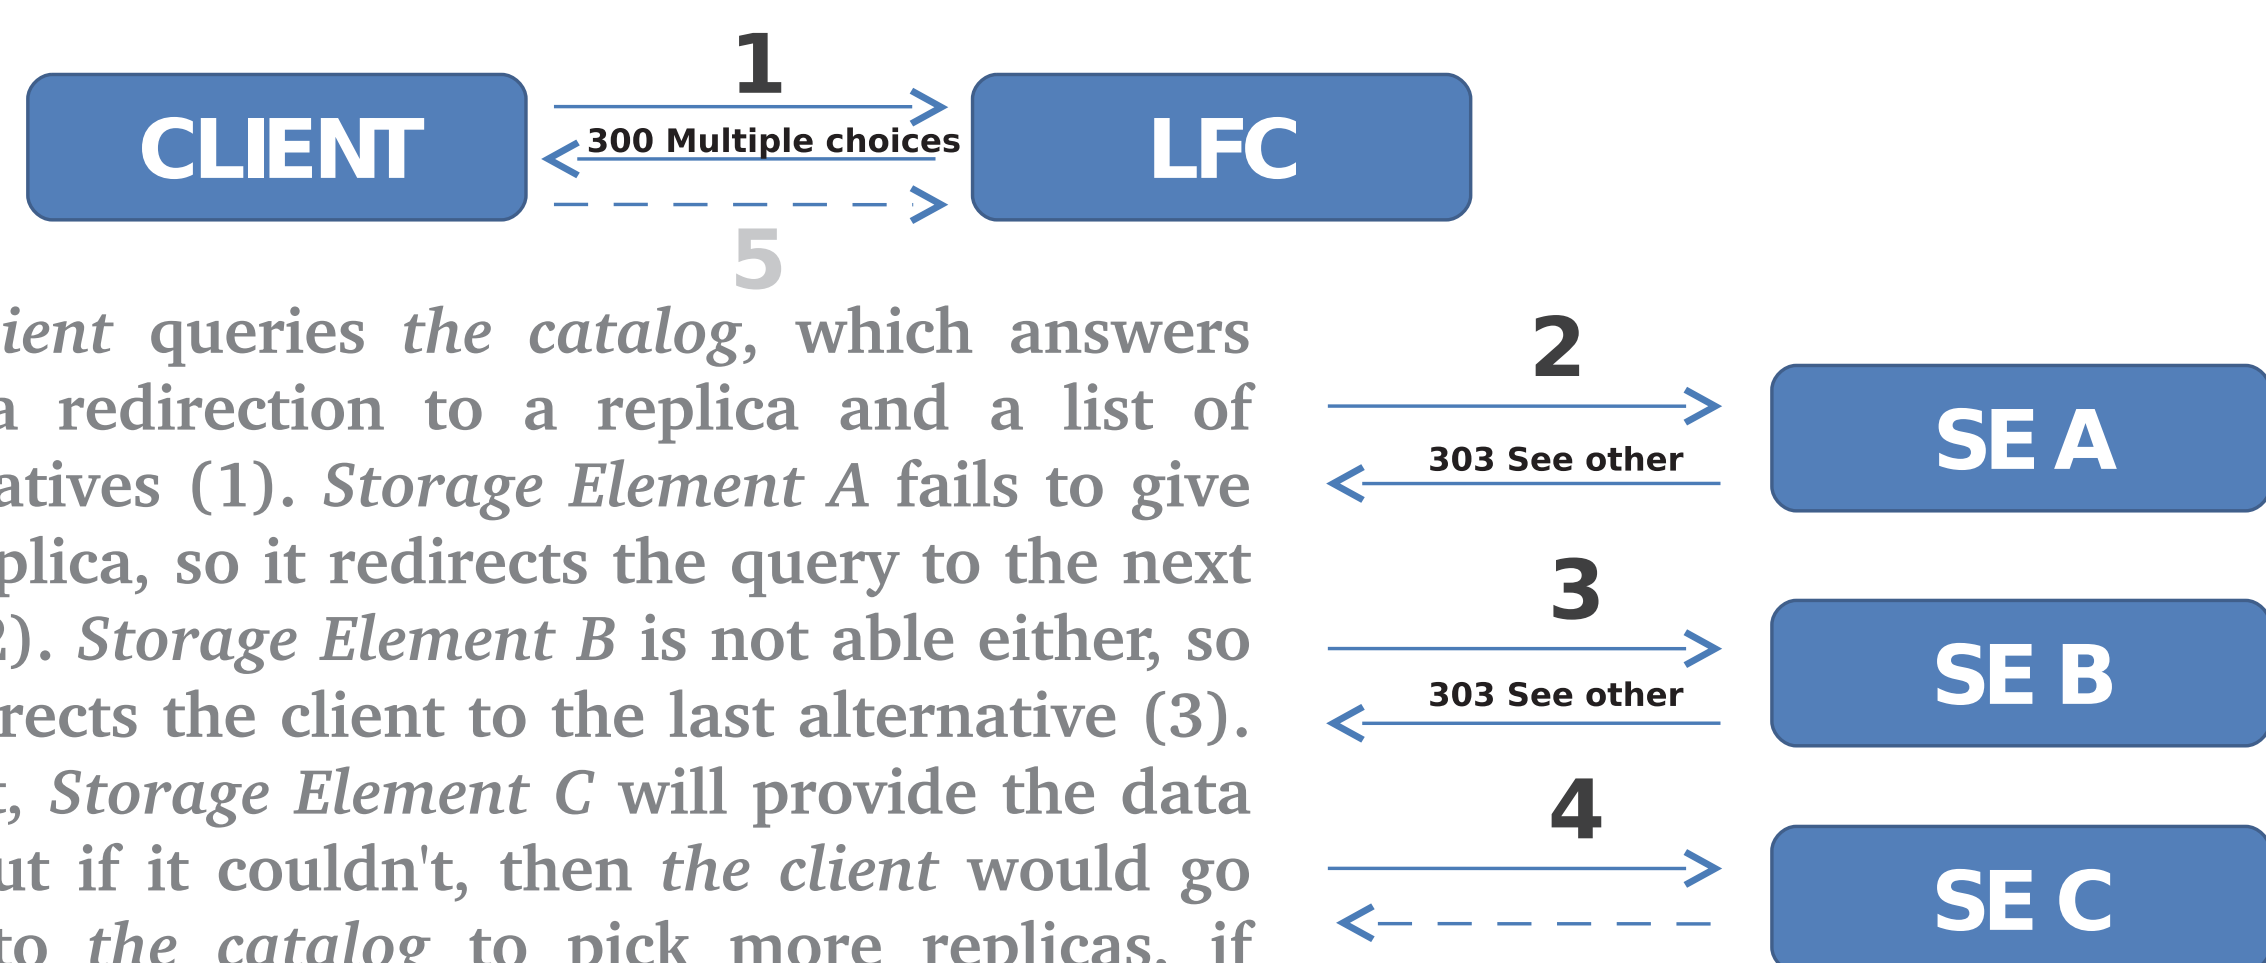
It can be used to browse a DPM or an LFC as a **regular filesystem**.



Nautilus connected to DPM through WebDAV


Transparency: Global Access Service

Provides transparency and fault tolerance on access from a catalog to a storage element. A redirection mechanism is defined supporting a server-driven **fall-back mechanism**.

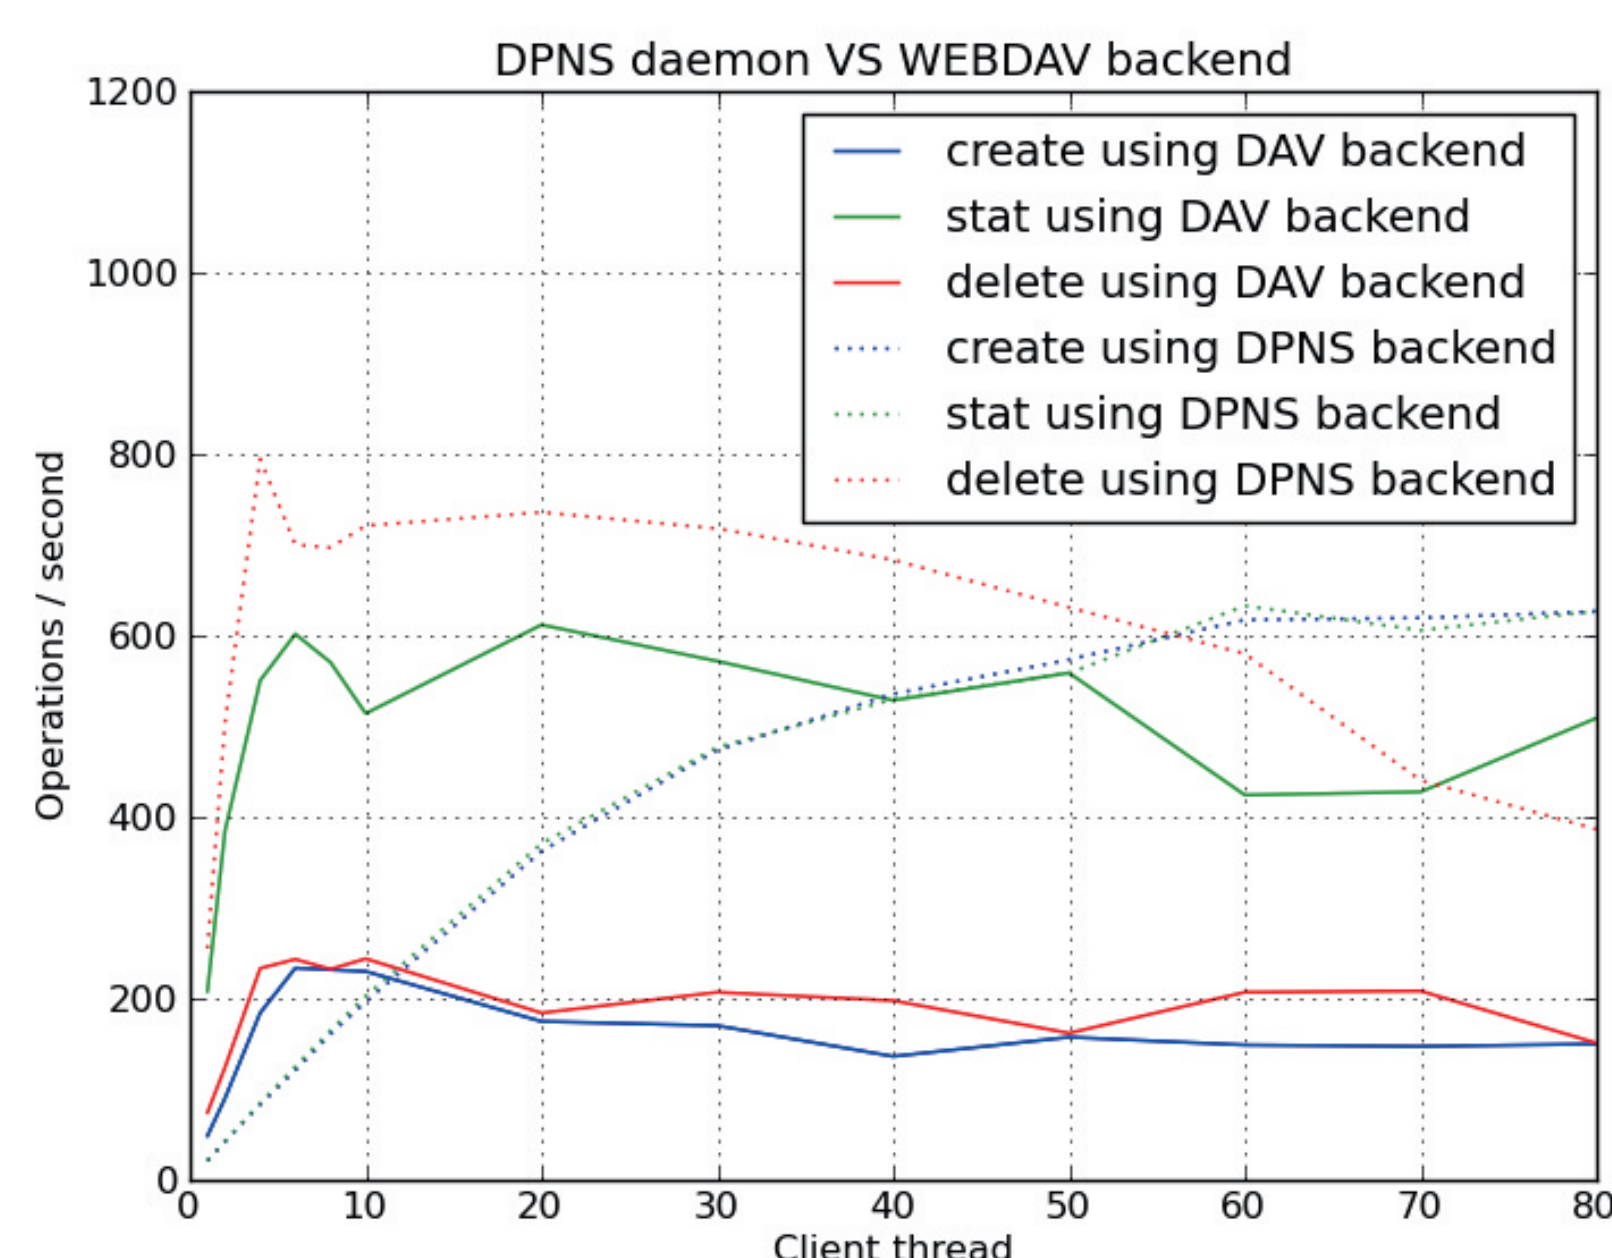


The client queries the catalog, which answers with a redirection to a replica and a list of alternatives (1). Storage Element A fails to give the replica, so it redirects the query to the next one (2). Storage Element B is not able either, so it redirects the client to the last alternative (3). At last, Storage Element C will provide the data (4), but if it couldn't, then the client would go back to the catalog to pick more replicas, if available (5).

Client support for WebDAV



Metadata performance (Preliminary data)




These tests were executed with Keep-Alive connections (sessions with DPNS)
WebDAV is on average comparable to DPNS.

High throughput: Multistreams

A file is split in chunks, and each chunk is transferred in parallel. This often **improves the general throughput**, as it mitigates the limitations of bandwidth per connection.

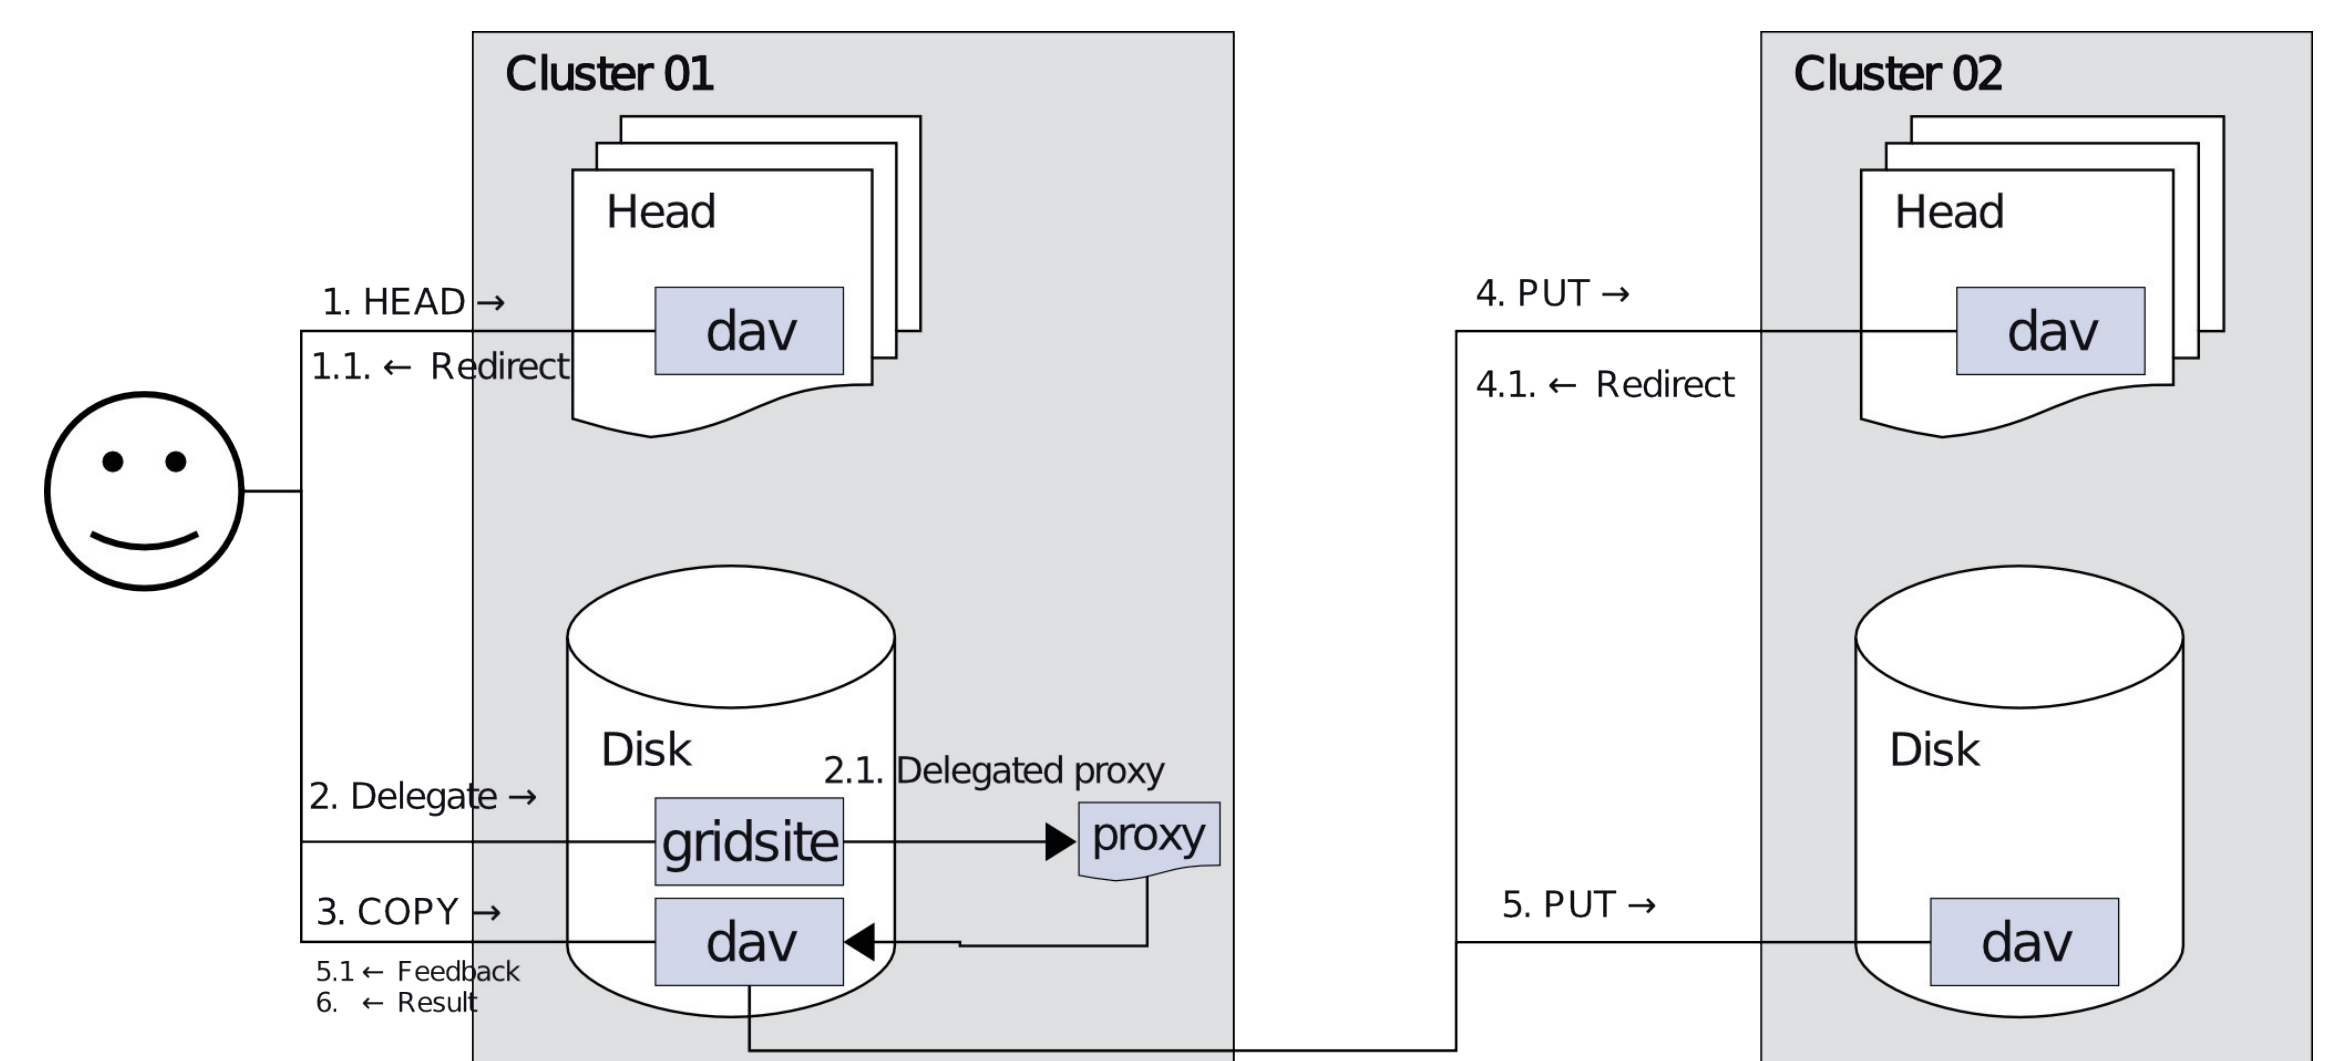
It can be done on HTTP thanks to Range and Content-Range headers.



The impact of multiple streams is higher for bigger files.

Replication: Third party copies


Copying data between servers is one typical use case. The most efficient way is to **delegate the copy** to the source or the destination server. Using WebDAV COPY method we implement a push model where the origin orchestrates the process.



Client support for HTTP



Transfer performance (8 streams on CERN → ASGC)



Two different GridFTP clients have been compared. lcg-cp and HTTP registers the file.
HTTP does as good as pure GridFTP, but including namespace operations.

Conclusions

Our HTTP/WebDAV implementation is production ready and a reliable alternative as an access and transfer protocol. It follows existing and familiar standards, making it easier to access data stored on the Grid. We are also experimenting with new backends to improve performance and functionality. It is currently available on EMI1, EMI2, EPEL, Fedora and GEP for legacy installations (gLite 3.2).

