

Scalable Global Data Access with HTTP/WebDAV

Alejandro Álvarez Ayllón
on behalf of the LCGM development team

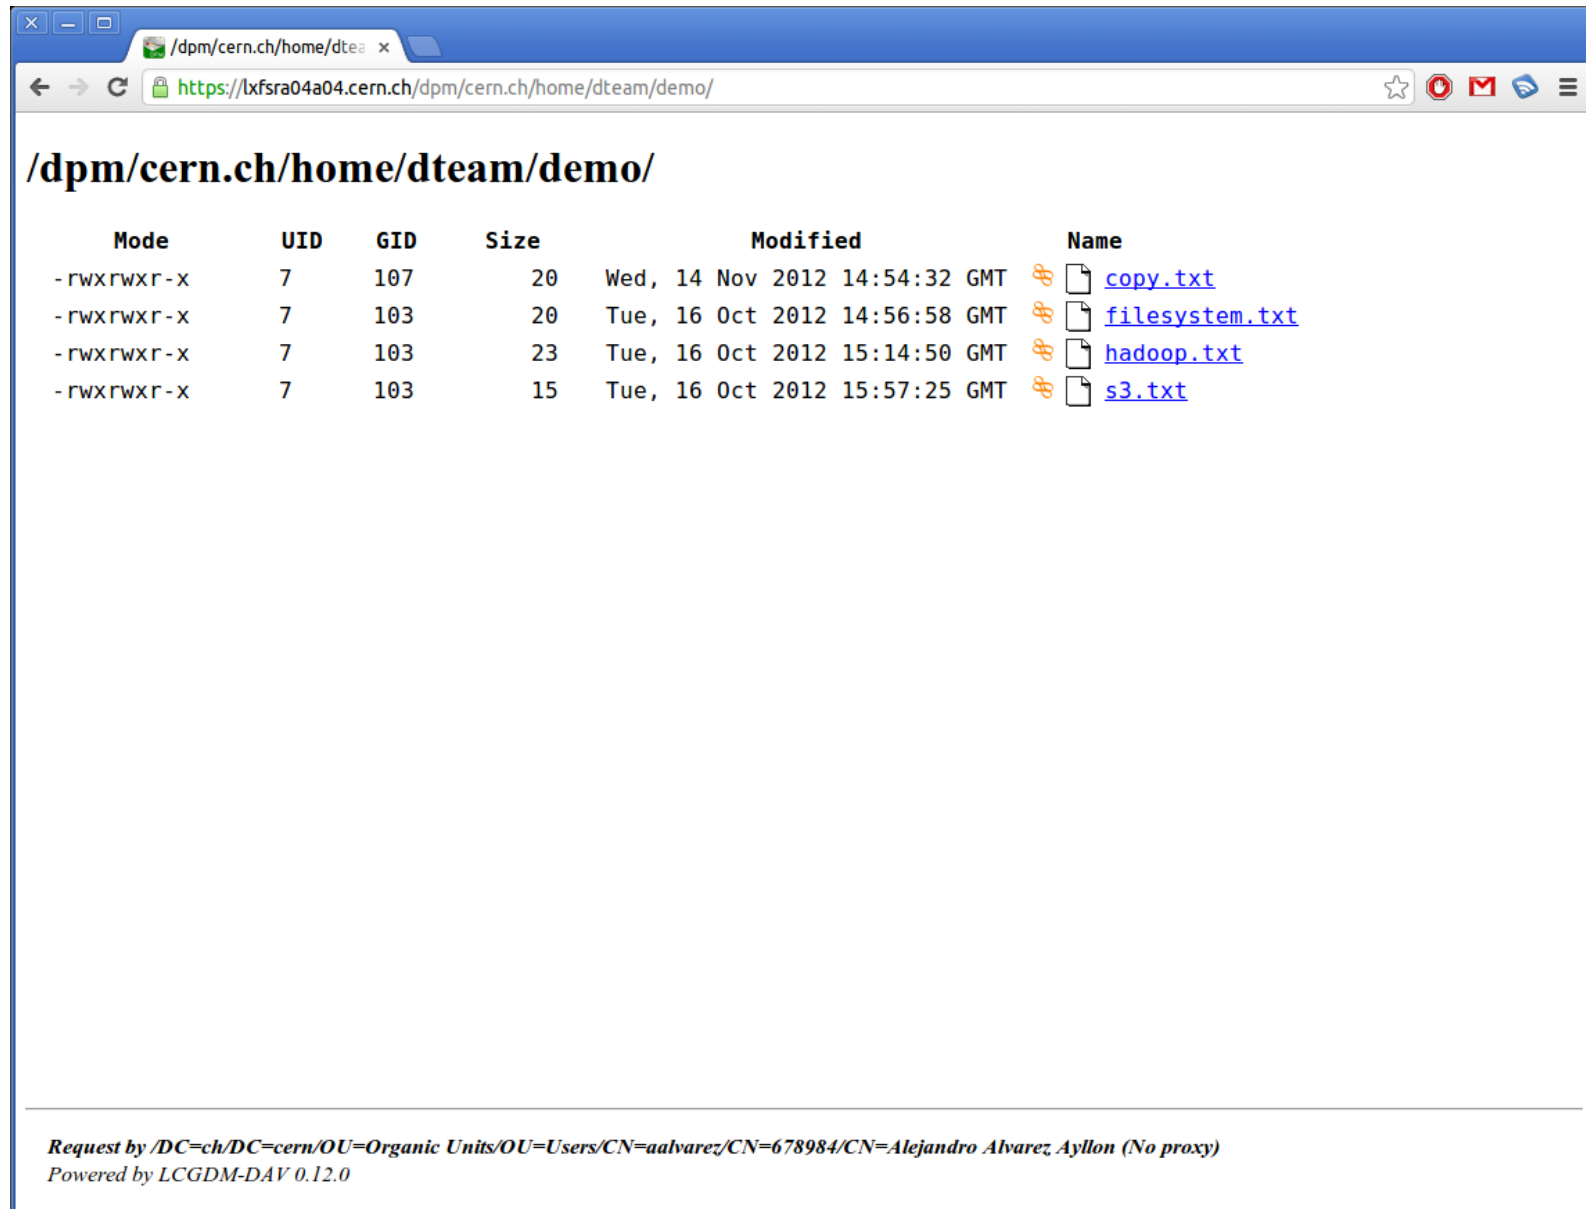
Overview

- HTTP/DAV in front of DPM
 - Use case: Installing in an existing DPM
 - Browsing using Nautilus
 - Random I/O
 - Third party copies
- HTTP/DAV in front of LFC
 - Use case: Accessing CERN's read-only LFC
 - Redirection
 - Fall-back with Metalink

HTTP/DAV in front of DPM

- We have already shown how installing lcgdm-dav-server is quite easy
 - And also, how to boost performance using the new plugins
- And we have seen how to navigate using a web browser
 - Which is already good enough!

HTTP/DAV in front of DPM




The screenshot shows a web browser window with the address bar displaying `https://lxfsra04a04.cern.ch/dpm/cern.ch/home/dteam/demo/`. The page content is a directory listing for `/dpm/cern.ch/home/dteam/demo/`. The listing is a table with the following columns: Mode, UID, GID, Size, Modified, and Name. There are four entries, each representing a text file.

Mode	UID	GID	Size	Modified	Name
-rwxrwxr-x	7	107	20	Wed, 14 Nov 2012 14:54:32 GMT	copy.txt
-rwxrwxr-x	7	103	20	Tue, 16 Oct 2012 14:56:58 GMT	filesystem.txt
-rwxrwxr-x	7	103	23	Tue, 16 Oct 2012 15:14:50 GMT	hadoop.txt
-rwxrwxr-x	7	103	15	Tue, 16 Oct 2012 15:57:25 GMT	s3.txt

Request by /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=aalvarez/CN=678984/CN=Alejandro Alvarez Ayllon (No proxy)
Powered by LCGDM-DAV 0.12.0

HTTP/DAV in front of DPM

- But we can mount it using Nautilus



HTTP/DAV in front of DPM

- Random I/O
 - A file can be read partially using Range headers
 - Using cURL, we can test it manually
`curl -L http://host/path/file -r start-end`
- Third party copies
 - Copy from one HTTPS to another
 - Source must understand logic
 - Destination doesn't
 - But you must be authorized, obviously!

HTTP/DAV in front of LFC

- We said we could browse a DPM using regular web browsers, and DAV clients (i.e. Nautilus)
- But we can do that too with the LFC!
 - Even better, if our client supports redirection, we can browse an LFC and access files from there

HTTP/DAV in front of LFC

- Fallback with Metalink
 - What happens if one of the replicas is broken?
 - Metalink gives us the full list, and a client (i.e. aria2) can try each one until it reaches a good one
 - Even better, if the file is big, it can use at the same time several replicas to increase throughput