



WebDAV in front of LFC and SE's

Alejandro Álvarez Ayllón

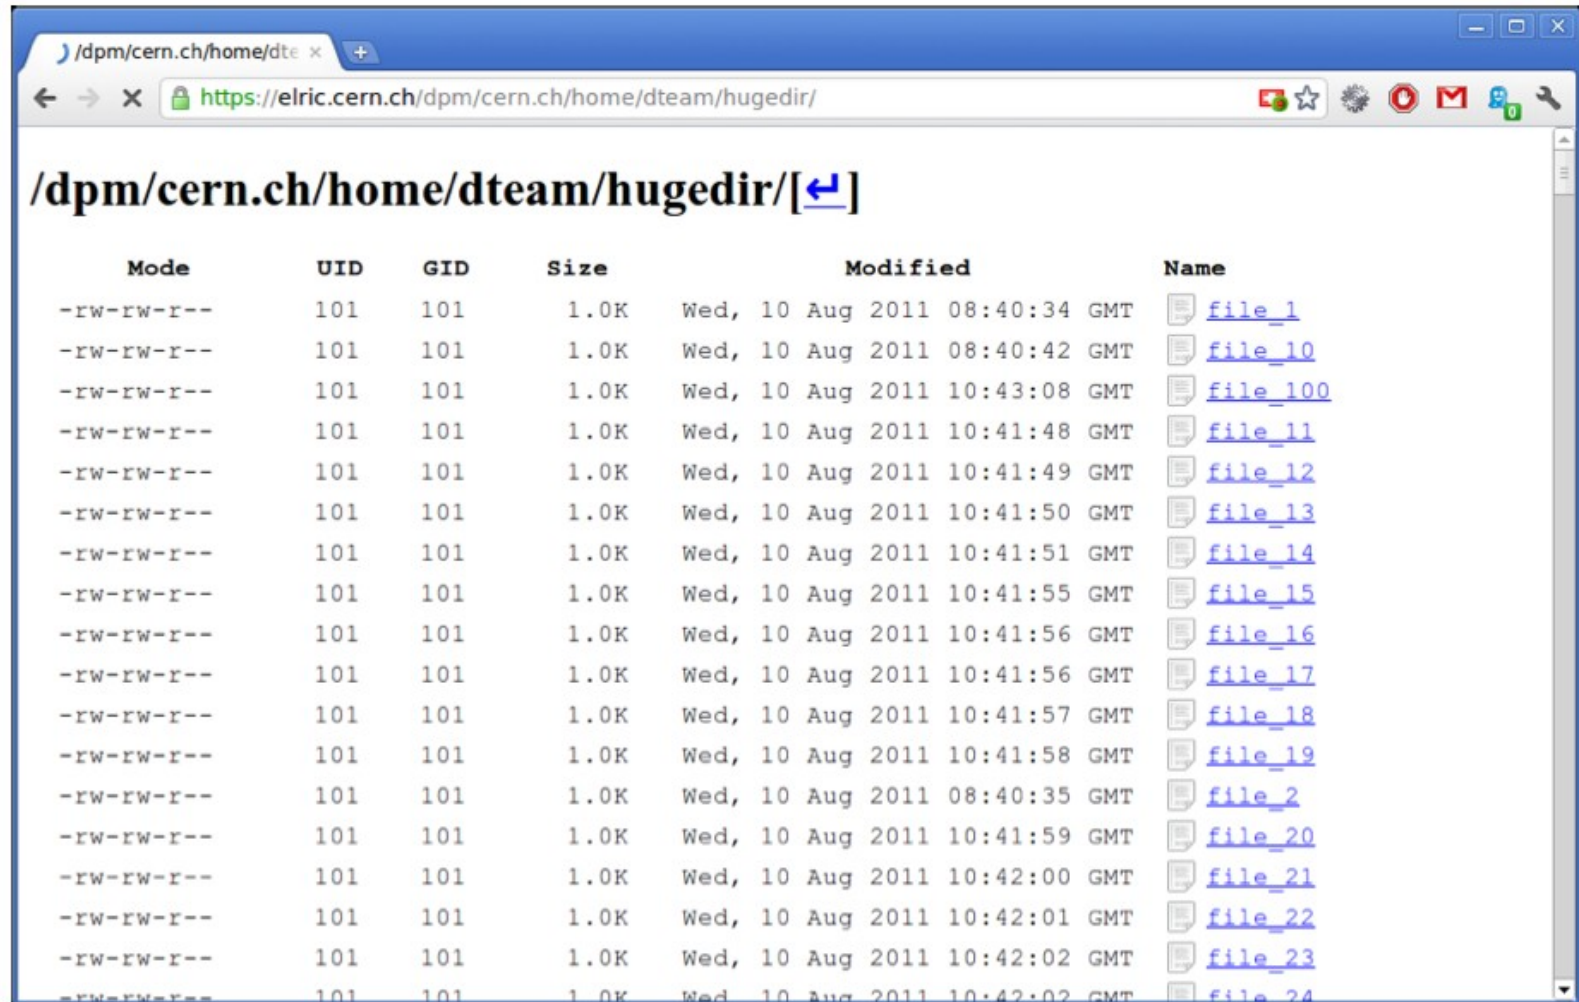
Overview

- DAV and DPM
- DAV and LFC
- Properties and replicas
- DAV, LFC, DPM and others
- Third party copies
- Issues and future work

DAV and DPM

- WebDAV access for DPM servers was the first milestone we had
 - Testing under high load is ongoing
 - This has been already made public (as beta)
- It uses redirection to drive the user from the Head node to the disk that has the replica
 - Uses internal `dpm_get` to obtain the best option if several replicas are available

DAV and DPM



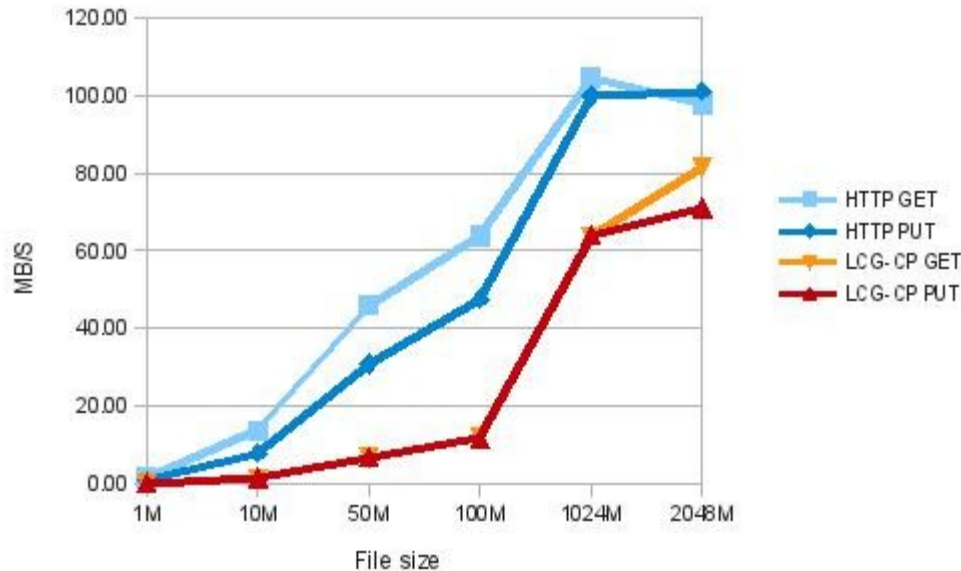
Browser address bar: <https://elric.cern.ch/dpm/cern.ch/home/dteam/hugedir/>

Page title: /dpm/cern.ch/home/dteam/hugedir/[↵]

Mode	UID	GID	Size	Modified	Name
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 08:40:34 GMT	file_1
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 08:40:42 GMT	file_10
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:43:08 GMT	file_100
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:41:48 GMT	file_11
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:41:49 GMT	file_12
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:41:50 GMT	file_13
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:41:51 GMT	file_14
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:41:55 GMT	file_15
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:41:56 GMT	file_16
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:41:56 GMT	file_17
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:41:57 GMT	file_18
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:41:58 GMT	file_19
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 08:40:35 GMT	file_2
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:41:59 GMT	file_20
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:42:00 GMT	file_21
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:42:01 GMT	file_22
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:42:02 GMT	file_23
-rw-rw-r--	101	101	1.0K	Wed, 10 Aug 2011 10:42:02 GMT	file_24

DAV and DPM

HTTP vs LCG-CP



LAN (1 Gbit/s link)

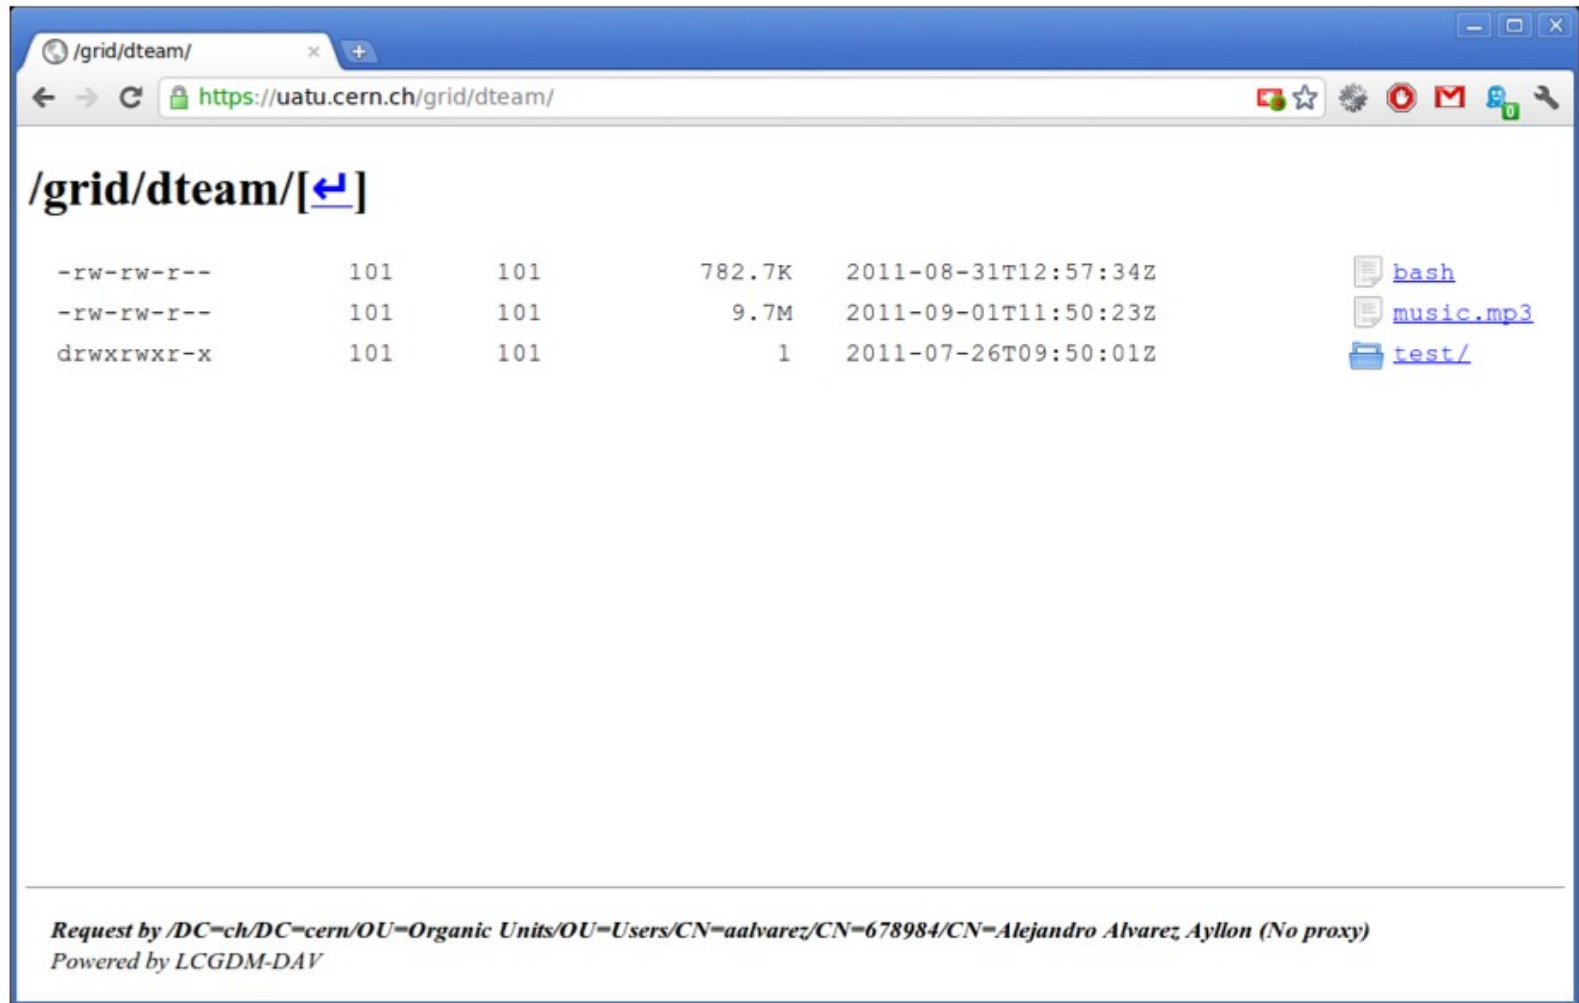
- No difference detected with different number of streams
- lcg-cp configured to use gridftp

- WAN (CERN → Taipei) **ONGOING**
 - For small files ($\leq 100\text{M}$) **HTTPS** is better for reading
 - For writing too, but for $\leq 50\text{M}$
 - For bigger files, GridFTP (srmcp) starts being better
 - 2.70M/s vs 1.84M/s for a 2GB file




DAV and LFC

- The next step would be to have LFC support too
 - The good news is: we already do!
- DPNS and LFC API's are pretty much the same
 - So the existing code almost worked in a LFC service
 - dpm_get is not available, so this was the only missing point
 - Easy to solve using Cns_getreplica when no DPM host is available
- HTTP redirection, just one step further
 - LFC → Head → Disk

DAV and LFC



The screenshot shows a web browser window with the address bar containing `https://uatu.cern.ch/grid/dteam/`. The main content area displays the path `/grid/dteam/` and a table of files. The table has columns for permissions, owner, group, size, and date. To the right of each row are icons and links for `bash`, `music.mp3`, and `test/`.

<code>-rw-rw-r--</code>	101	101	782.7K	2011-08-31T12:57:34Z	 bash
<code>-rw-rw-r--</code>	101	101	9.7M	2011-09-01T11:50:23Z	 music.mp3
<code>drwxrwxr-x</code>	101	101	1	2011-07-26T09:50:01Z	 test/

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

Properties and replicas

- WebDAV supports properties
 - So that already gives us how to handle replicas
 - Our own namespace, but a common could be agreed
 - One replica per line
 - Replicas added/removed with +/- prefixes

```
http://apache.org/dav/props/ executable = F
DAV: iscollection = 0
LCGDM: guid =
LCGDM: mode = 100664
LCGDM: sumtype =
LCGDM: sumvalue =
LCGDM: replicas = srm://elric.cern.ch/dpm/cern.ch/home/dteam/music.mp3
srm://yyrkoon.cern.ch/data/dteam/music.mp3
DAV: owner = /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=aalvarez/CN=678984/CN=Alejandro Alvarez Ayllon
DAV: group = dteam
dav:/grid/dteam/> set namespace LCGDM:
dav:/grid/dteam/> propset music.mp3 replicas -srm://yyrkoon.cern.ch/data/dteam/music.mp3
```


DAV, LFC, DPM and others

- LFC contains a list of replicas
 - Protocol not relevant, as the URI will be parsed
- The redirection will be done using HTTP and URL's
 - So, as far as the SE provides HTTP, it will work regardless of the implementation
 - **Warning!** The DAV server is not omniscient
 - The path part of the URL **must** match the path part of the SFN
 - HTTPS is assumed
 - The port has to be standard (443)
 - Can this be enforced?
 - As long as HTTP/S URI's are not used, at least

Third party copies

- Third party copies are supported using COPY method
 - Tested from DPM to dCache
- Delegation is needed
 - A common mechanism should be agreed between DAV implementations so one single client can be used

Issues and future work

- In a LFC, entries can be modified but not created. How?
 - PUT's to be supported in LFC's?
 - Directories can be created, though
- Selection algorithm needed when several replicas are available
 - Proximity, throughput, protocol,...
- Improve third party copy implementation
- More and heavier testing



Thank you!

EMI is partially funded by the European Commission under Grant Agreement RI-261611