

Data Access – Standards & Status

Alejandro Álvarez Ayllón
on behalf of the LCGM development team



EMI is partially funded by the European Union Horizon 2020 research and innovation programme under the Marie Skłodé Curie Agreement RI-261611

- ▶ HTTP / DAV
- ▶ NFS4.1 / pNFS
- ▶ XROOT
- ▶ GridFTP
- ▶ SRM

- ▶ Frontend based on Apache2 + mod_dav
- ▶ In production since 1.8.3
- ▶ Can be used for GET/PUT, but also direct access
 - ▶ Some extensions for full GridFTP equivalence
 - ▶ *Multiple streams using Range/Content-Range*
 - ▶ *Third party copies using COPY+ Gridsite Delegation*
 - ▶ Random I/O
 - ▶ *Possible to do vector reads and other optimizations*
- ▶ With 1.8.4 is already DMLite based

- ▶ Metalink support
 - ▶ An **Internet standard** file format that allows to specify multiple download locations for one single resource
 - ▶ Some clients allow parallel download, hash verification, etc...
 - ▶ And all of them **fallback mechanisms**
 - ▶ Can be used with aria2, cURL, jDownloader,...
 - ▶ Integrated by Mandriva (urpmi+aria2), Fedora (yum) and openSUSE (ZYpp+aria2)

HTTP: Client support

	curl	browser
OS	Any	Any
GUI	NO	YES
CLI	YES	NO
X509	YES	YES
Proxies	YES	Only IE so far
Redirect	YES	YES
PUT	YES	NO

DAV: Client support




	TrailMix	Cadaver	Davlib	Shared Folder	DavFS 2	Nautilus	Dolphin
OS	Firefox < 4	*nix	Mac OS X	Windows	*nix	Gnome	KDE
GUI	YES	NO	YES	YES	N/A	YES	YES
CLI	NO	YES	NO	NO	N/A	NO	NO
X509	YES	YES	NO	YES	YES	NO	NO
Proxies	?	NO	NO	YES	NO	NO	NO
Redirect	YES	NO	YES	Not PUT*	NO	NO	YES


(*) Windows 7 doesn't redirect on GET either

HTTP: GET/PUT

PUT



GET



- ▶ Direct access to the data, with a standard NFS client
- ▶ Available with DPM 1.8.3 (read only)
 - ▶ Not yet based on DMLite
- ▶ Implemented as a plugin to the Ganesha server
- ▶ Only Kerberos authentication for now
 - ▶ Issue with X509 client support in Linux (Server ready)
 - ▶ We're investigating how to add this

- ▶ Supported by Unix (Solaris, AIX,...) and Unix-like (Linux, FreeBSD) operating systems
- ▶ Mac OS X, Microsoft Windows,...
- ▶ As mentioned, X509 support is not available yet

- ▶ Not really a standard, but widely used in HEP
- ▶ Initial implementation in 2006
 - ▶ No multi-vo support, limited authz, performance issues
- ▶ New version 3.1 (rewrite) available with 1.8.4
 - ▶ Multi VO support
 - ▶ Strong authn/authnz with X509, but ALICE token still available
 - ▶ Better performance
- ▶ Based on the standard XROOTD server
 - ▶ Plugins for XrdOss, XrdCmsClient and XrdAccAuthorize
- ▶ Support for per-VO xrootd federations
 - ▶ A number of DPMs are now in the Atlas FAX project

- ▶ Available via dedicated repository
 - ▶ Including xrootd packages (EPEL packaging)
- ▶ Expect xrootd 3.2.x in EMI third party first
 - ▶ Eventually, EPEL
- ▶ Working on a dmlite-based version
 - ▶ Even better performance
 - ▶ Multiple pool types, different namespace implementations, etc.

- ▶ Working on a new version based on DMLite
- ▶ Investigating redirection support
 - ▶ Seems possible!

- ▶ SRM imposes significant latency for data access
 - ▶ It has its use cases, but it is a killer for regular file access
 - ▶ For data access, only for protocols not supporting redirection
- ▶ Ideally, it will be kept for space management only
- ▶ Still, a initial proof-of-concept DMLite based implementation has been done
 - ▶ Only few calls are supported
 - ▶ *Ls, Get, Put*
 - ▶ More on this at 16.20
 - ▶ *“Data Access - The Missing Standards”*