



Grid Data Access:
Proxy Caches and User Views

EGI Technical Forum
19 September 2011

Jan Just Keijser

Cristian Cirstea

Jeff Templon

BiG Grid

the dutch e-science grid



Outline

- Who?
- Why?
- What?
- How?
- Results
- What's next?




By who and for whom?

By who

- Developed by Cristian Cirstea
 - Student at TU Eindhoven in the Netherlands
 - 9 month master thesis project
- Supervised by Jeff Templon and yours truly

For whom

- Initially for HEP users, but made applicable to all grid users
- End users that wish to browse datafiles on the grid on their laptop or desktop computers
- Grid users that have many jobs requiring access to the same set of files





Why was it built?

Existing data management tools:

- Require a specific version of Linux – no other platforms are supported
- Provide command-line tools only
- Do not integrate well into the operating system
- Do not scale well when requesting the same datafile many times
- ...

Need I continue?





What is it?

Grid Proxy Cache

- Offers a standard WebDAV interface to users and to grid jobs
- Authentication can be done using username+password (via a MyProxy store) and using grid proxies
- Serves as a proxy to existing LFC and SRM implementations without having to change the LFCs or SRMs themselves
- Current implementation is readonly (it's a cache!)



How does it work?

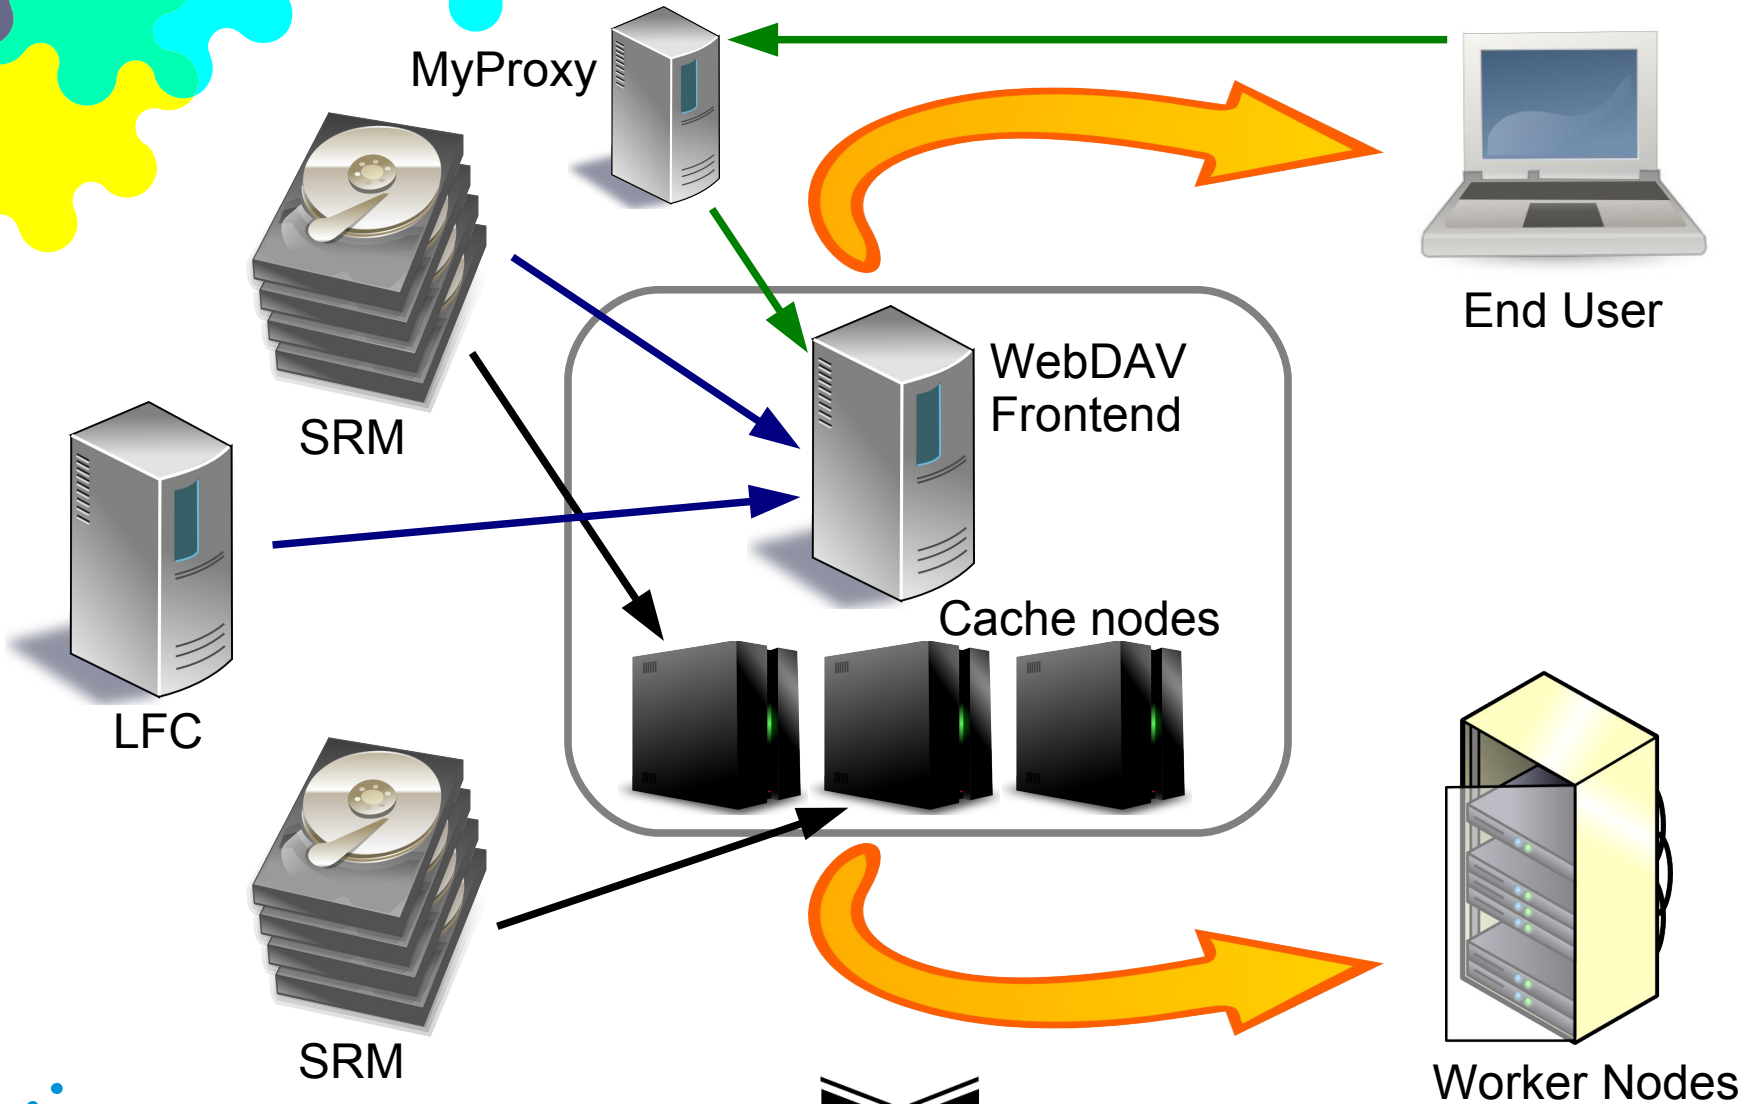
Using mostly existing technologies

- WebDAV server and cache nodes are based on apache+mod_gridsite
- Cache nodes configured for automatic failover and load balancing
- Files retrieved from LFC/SRM using GFAL
- Linux FUSE client 'wdfs'
- WebDAV interface done using FastCGI scripts

Modifications made

- Gridsite's 'htproxypu' enhanced to support limited proxies
- 'wdfs' enhanced to support grid proxies and HTTP REDIRECTs

Architecture





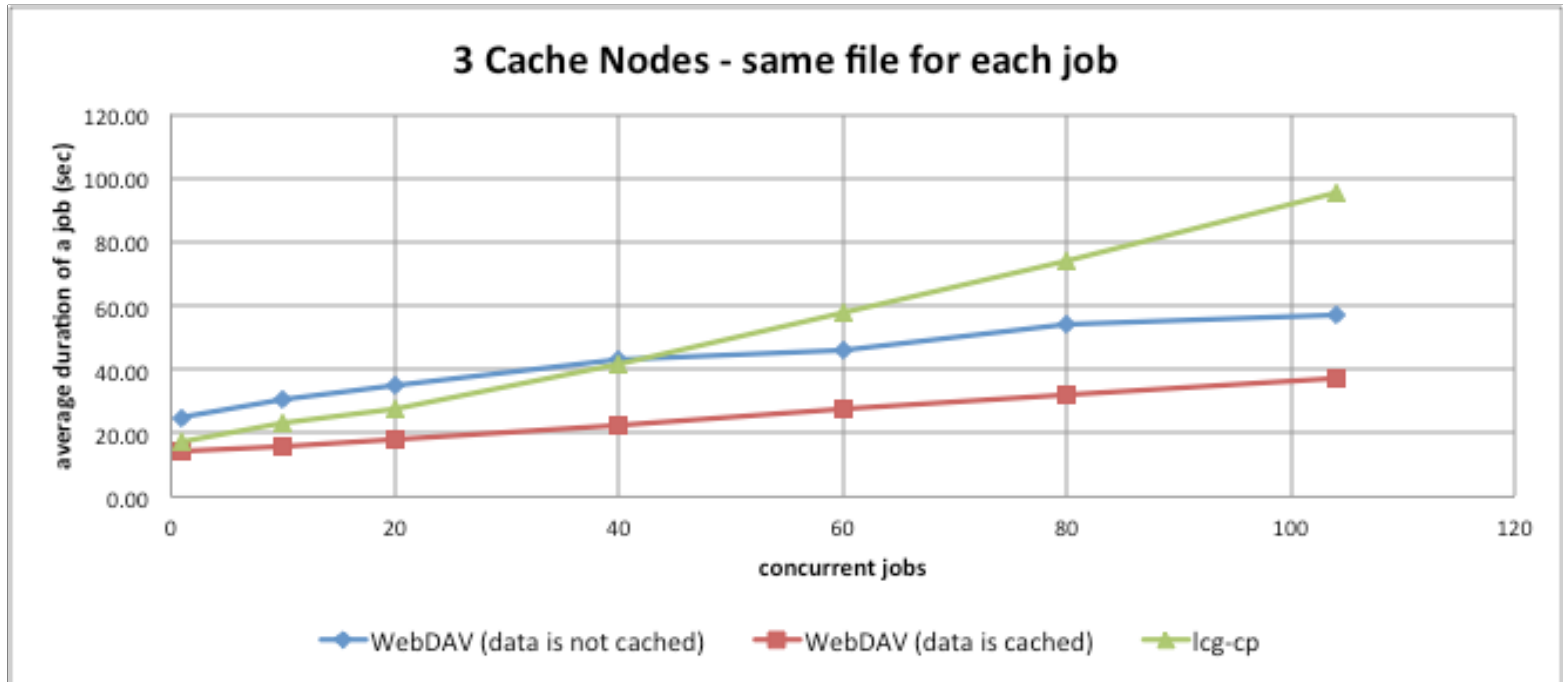
Results (1): Usability

- Works with Windows XP/7, Mac OS X, Linux native WebDAV clients
- Also works with a web browser (and 'curl'/'wget')
- Browsing of various LFCs and SRMs is transparent
- Linux FUSE driver offers full POSIX access ('open'/'fopen') to grid files from within grid jobs

Issues

- Unmounting FUSE mounts sometimes fails → Obstacle for production use
- MyProxy credentials need to be uploaded using Firefox plugin or command-line tool

Results (2): Performance



Issues

- When many jobs all request different files a proxy cache can add extra overhead



Conclusions and future work

Conclusions

- For end user file browsing current solution is production ready
- The system scales well under load
- For grid job access the FUSE issue needs to be solved

Future work

- Investigate possible write access
- Optimize WebDAV throughput
- Once SRMs provide an HTTPS interface SSLProxyCaching might be possible → further speed improvement
- Add a translation layer (e.g. grab ZIP file from SRM, unpack it, then server the contents via WebDAV)