

Remote data transfer

it can be easy

Leonardo Sala, PSI

Why? Examples

- Today: it's quite common to have measurements done at LCLS
- Tomorrow: we do hope it will be common to have many external users at FELs
- We are talking about O(TB) for experiment

How?

- *ssh: no-go*. Bad performances for non-negligible network latencies
- But! Big data transfers is daily routine for *LHC*
 - Actually, it was part of my previous job
 - **GridFTP** protocol: developed for efficient WAN transfers of large data around the world
 - **Can be complex**: X509 certs, dedicated WLCG SW...
- GridFTP can be used in a user-friendly way, too. *GlobusOnline* is a good example of this

A quick example



Manage Data | Support | leonsala

start transfer | view transfer activity | manage endpoints | dashboard

Transfer Files

Get Globus Connect
Turn your computer into an endpoint.

Endpoint: Go

Path: Go

select all | none | up one folder | refresh list

HEWL-1	Folder
HEWL-1-Periodic	Folder
test_10GB	Folder
test_1GB	Folder
test_6MB	Folder
test_6MB_live	Folder
fillPeriodicDatasetFromRealData.sh	314 b
fillPeriodicDatasetFromRealData.sh~	319 b
testfile_1GB	2 MB

Endpoint: Go

Path: Go

select all | none | up one folder | refresh list

testfile_1GB_1	905 MB
testfile_1GB_10	1000 MB
testfile_1GB_2	938 MB

This is a Win/MacOsX/Linux client you can install on your machine to simply do GridFTP transfers

more options

Label This Transfer

This will be displayed in your transfer activity.

Live!

Is it new?

Well, at least some knowledge at DESY is already available!



globus online  Manage Data | News & Events | About

Home » News & Events » Blog

blog TIDBITS & EXABYTES

Friday, November 2nd, 2012 comments off



Rajkumar Kettimuthu.

October 2012 User of the Month: Dmitry Ozerov


I am pleased to announce that the user of the month for October 2012 is Dmitry Ozerov from DESY in Germany.

Dmitry has a background in particle physics research and is currently working as an IT data management specialist in the Scientific Data Management and Grid Computing team at DESY, Hamburg. The team manages 7.5 petabytes of disk and 4 petabytes of tape storage, and the batch facility for the scientists from the High Energy Physics and Photon Science.

Dmitry is helping scientists from the [Center for Free-Electron Laser Science](#) to transfer a large amount of data from the LCLS in Stanford, USA to his home institute in Hamburg, Germany. Dmitry moved a whopping 276 terabytes of data in less than a month! He achieved transfer rates as high as 450Mbps over the transatlantic link, with round trip time greater than 150 ms.

Is it fast?

Quick answer: yes! In case of high-latency networks, GridFTP improvement over SSH (even with HPN patch) is impressive



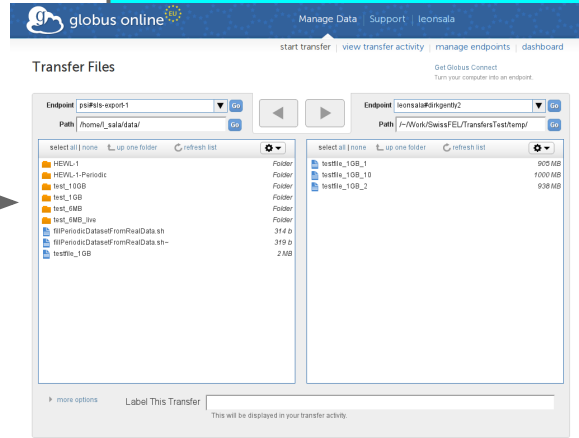
Test details:

- All transfers done within the PSI network
- Latencies are simulated in the Linux kernel
- Packet loss not taken into account
- 1 Gbit link

Behind the scenes



OpenID



The screenshot shows the Globus Online web interface for managing data transfers. The page title is "Transfer Files" and it includes navigation links for "start transfer", "view transfer activity", "manage endpoints", and "dashboard". The interface is split into two panels for source and destination endpoints.

Endpoint 1 (Source): `ssslab-exp01-1`
Path: `/home/ssala/data/`

Item	Type
MEWL-1	Folder
MEWL-1-Periodic	Folder
test_1GBB	Folder
test_1GB	Folder
test_0MB	Folder
test_0MB_dir	Folder
RPPeriodicDataFromRealData.st	314 D
RPPeriodicDataFromRealData.st-	319 D
testfile_1GB	2 MB

Endpoint 2 (Destination): `ssslab01gen012`
Path: `~/Work/SwissFEL/transfer/testtemp/`

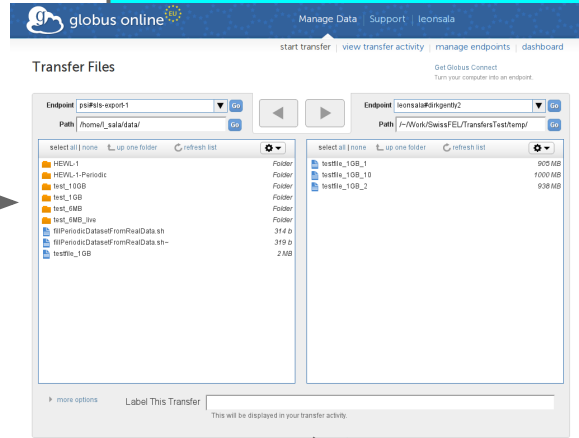
Item	Type	Size
testfile_1GB_1	File	900 MB
testfile_1GB_10	File	1000 MB
testfile_1GB_2	File	938 MB

At the bottom, there is a "Label This Transfer" field and a note: "This will be displayed in your transfer activity."

Behind the scenes



OpenID



request

credential

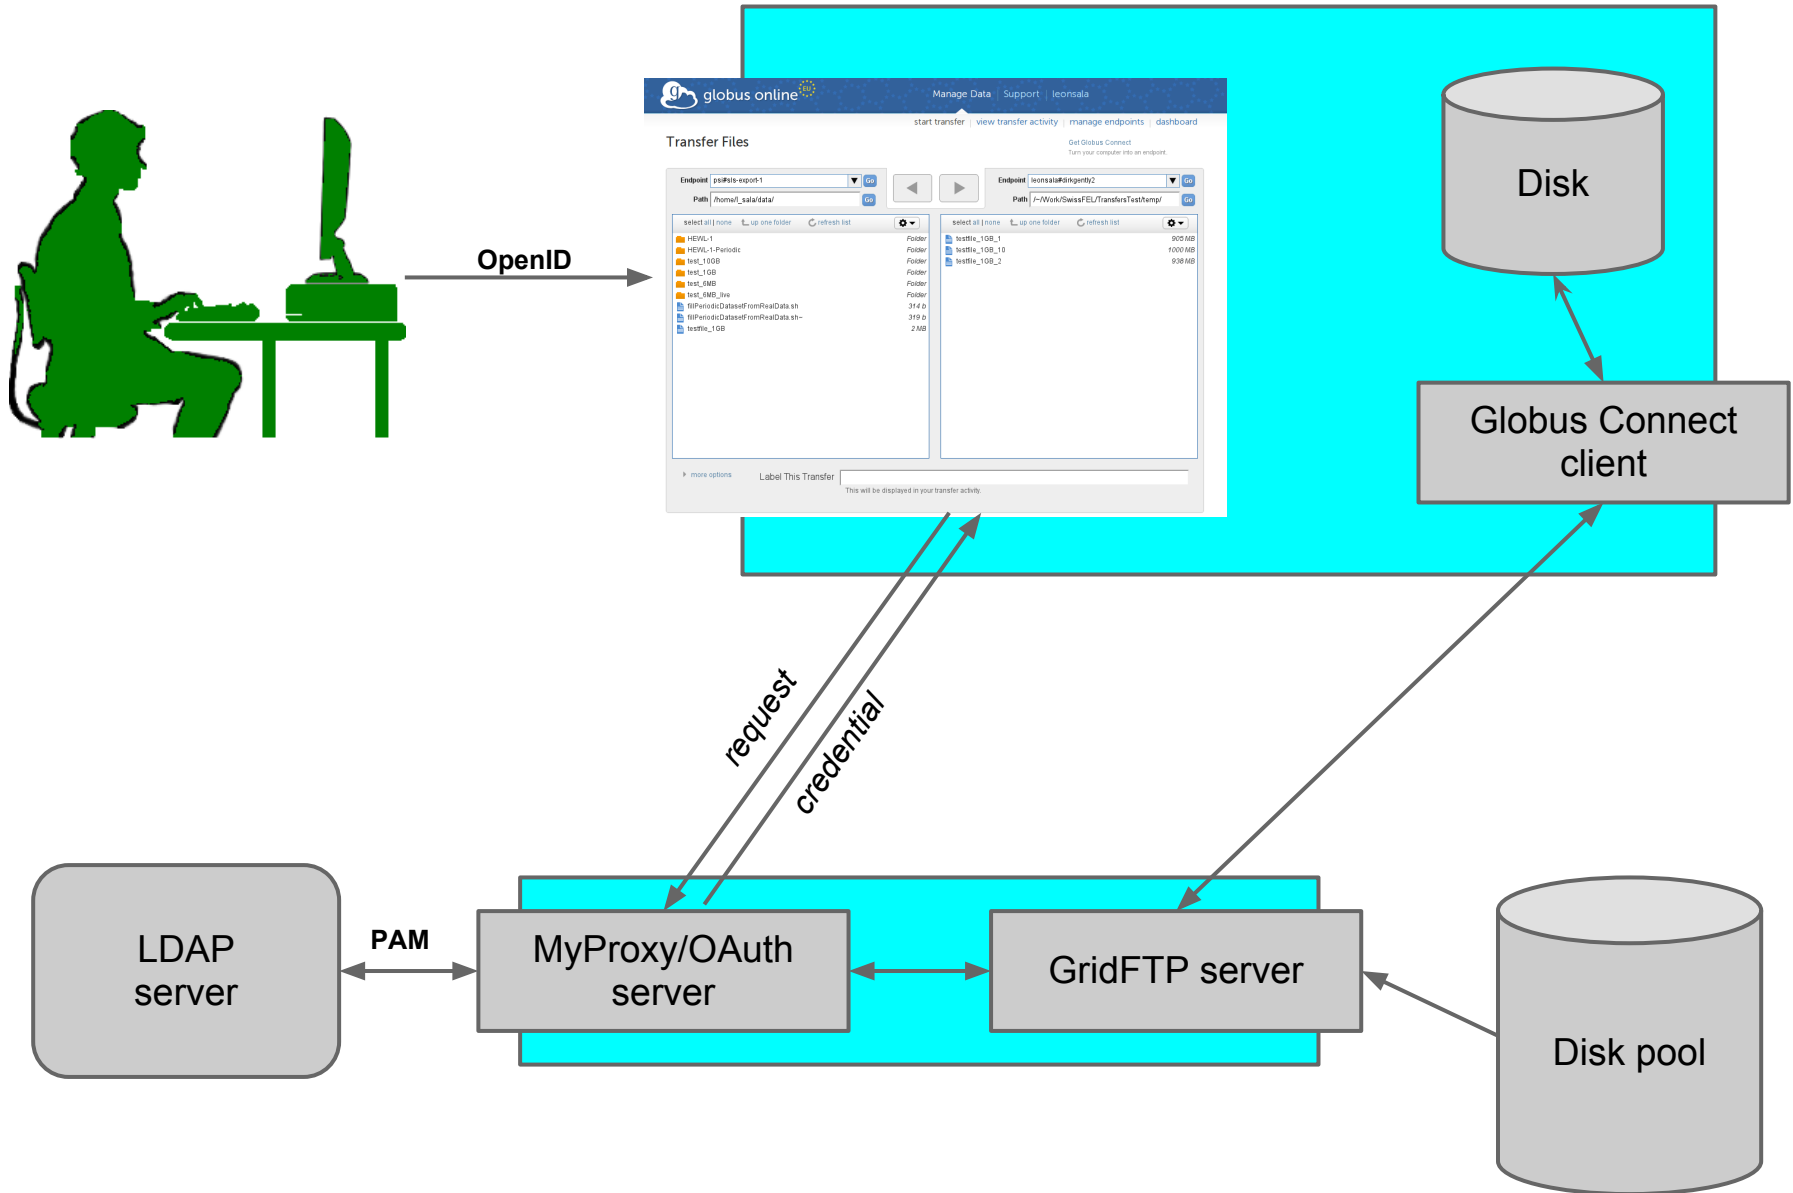
LDAP
server

PAM

MyProxy/OAuth
server

OAuth enables us to have Globus Online delegate the authentication to our own internal authentication service

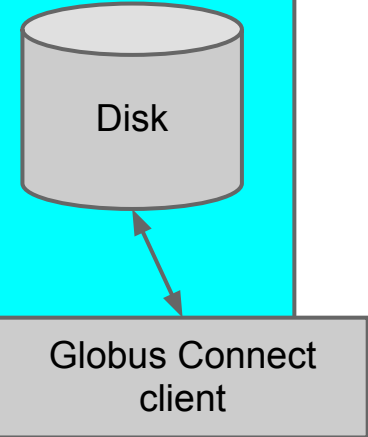
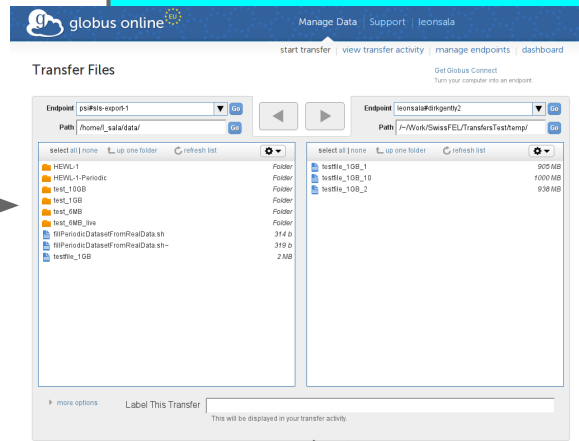
Behind the scenes




Behind the scenes



OpenID



If you want to connect two GridFTP servers, substitute the GC client with another GridFTP/MyProxy server



User tests

At PSI we have installed a GridFTP+OAuth MyProxy server, and enabled few selected users for testing

- E.g.: transfers from LCLS: ~20MB/s. Using SSH is ~4-5 times slower
- Good feedback from users:
 - Service is easy and intuitive
 - No IT knowledge required *to transfer data*
- Getting full throughput require anyhow work on networks (this is what LHC has done)

Wrap-up

- "Big data" transfers are already here
 - it took years to LHC to be fully efficient on this
 - we can profit from this experience
- "Big data" will be here!
 - and won't go away
 - it will be an important part of the analysis workflow
- Solutions are already there
 - I've shown the one I found: quicker, easier, production-ready
 - I believe that the underlining technology will stay

Backup slides

How?

- ssh: no-go. In case of network latency, performances are quite bad
- Do not need to reinvent the wheel: tools are already available
 - GridFTP is one of them
 - developed within the Grid community, daily used by LHC
- To overcome complexity connected with Grid*, we can profit from an existing web-based service: GlobusOnline