


GridFTP performance

Tomas Lindén

Helsinki Institute of Physics

CMS programme

28.08.2005



M-grid admin meeting Oulu, Finland

GridFTP

What is *GridFTP*?

- *GridFTP* is the name of a variation of the File Transfer Protocol (FTP) which has Grid Secure Infrastructure (GSI) enabled authentication using X.509 certificates.
- There is no client application called *GridFTP*.

ARC distributes several *GridFTP* clients

- **ngcopy**
- **gsincftp**
- **globus-url-copy**

ngcopy

- Can transfer a directory
- More userfriendly than gobus-url-copy
- Does not support third party transfers
- Does not support parallell file transfers

gsincftp

- Very useful for interactive transfers and to debug URLs
- Can complete filenames by pressing TAB
- Unsupported
- Does not support parallell file transfers
- Bug: Does not support filesizes larger than 2 GB

globus-url-copy

- Does support parallell file transfers
- Does support third party transfers
- Not really an end user command and very awkward to use
- Does not support wildcards
- Does not support recursive copies
- Does not create needed intermediate directories
- Bug: The verbose option `-vb` gives no output for parallell streams, when pushing data

On a WAN-connection the network latency can drastically diminish the bandwidth if the application waits for acknowledgement of each small packet.

The latency problem has been solved for **globus-url-copy** and other tools by

- Increasing the size of the tcp-buffers and the size of the data that is acknowledged during the transfer
- Allowing several simultaneous tcp/ip streams within one logical transfer

Using **globus-url-copy** on a WAN can improve the bandwidth significantly:
Example a 256 MB file transferred 2005.08.25 from lxplus057.cern.ch to silo2.physics.helsinki.fi (50–60 ms RTT from ping) over a Gb/s link

protocol	<i>scp</i>	<i>globus-url-copy</i>	<i>globus-url-copy</i>
options	-c blowfish		-p 16 -tcp-bs 2000000
bandwidth MB/s	0,5	0,5	10,5

This example shows an speedup factor of more than 20.

Sometimes 20 MB/s has been reached between CERN and Helsinki, still without really searching for the most optimal parameters.

The variable **GLOBUS_TCP_PORT_RANGE** needs to be set by hand for parallell streams to work according to the local firewall configuration.

On a LAN connection the latency is smaller so there is not such a big difference between different file transfer tools.

It seems that pushing data on a LAN with **globus-url-copy** is faster than pulling data.

Even better results can be obtained by tuning the operating systems tcp/ip stack and the switches and routers parameters in addition to using parallell file transfers.

With *sepeli* integrated into M-grid and the CSC supercomputing infrastructure, CSCs users can hopefully soon use parallell file transfers.