



Mechanisms for IDS Data Transfer

Steve Fisher <dr.s.m.fisher@gmail.com>

Kevin Phipps <kevin.phipps@stfc.ac.uk>

Frazer Barnsley <frazer.barnsley@stfc.ac.uk>

Condie McKenzie-Smith <condie.mckenzie-smith@stfc.ac.uk>

March/April 2015

The problem

- The IDS works nicely for “reasonable” download sizes but transfers taking more than a few minutes can be problematic:
 - if the network is not working well
 - or if you want to unplug your laptop and take it home with you
- Zip files are generated on the fly. Most of the time this is good but to restart after a large offset requires a lot of computation and waste reading of the data that has already been transferred.
- Prefer to avoid adding complexity to the IDS.
- People like file system view of the world.

Various solutions

- FUSE
 - Needs installation on user's machine (in user space)
 - Offers file system view of ICAT/IDS
- WebDAV
 - Web application (installed on server)
 - Offers file system view of ICAT/IDS
- GridFTP
 - Would need to get files from IDS to some area local to the IDS (using link mechanism) then use some variant of gridFTP
 - Needs installation on user's machine (as "root")
- SmartClient
 - Client on user's machine and a personal server talking to the IDS
 - Needs installation on user's machine

FUSE

What is it?

- Filesystem in Userspace
- Includes a kernel module and a library
- Python wrapper makes it very easy to expose almost anything as a file system just need to implement a few functions

FUSE

Implementations for ICAT/IDS

- As it “easy” to do there are quite a number of partial implementations including experiments exposing parameters as part of the file system.
- A practical one is been developed for ISIS but should be of general applicability.
 - A configuration file allows you to restrict the view of the file system to just your data and to specify the hierarchy you want to see.
 - Currently the following hierarchies are supported:
 - facility, cycle, instrument, investigation, dataset, datafile...
 - facility, investigation, dataset, datafile...

FUSE

Pros and cons

Pros

- Fast
- Simple file system view of the world
- FUSE kernel module supported by OS on Linux

Cons


- Needs installing on client machine though can be installed by ordinary user
- Hard to manage slow archive storage as reads and writes can block for a long time
- You have the power of “rm -rf”

WebDAV

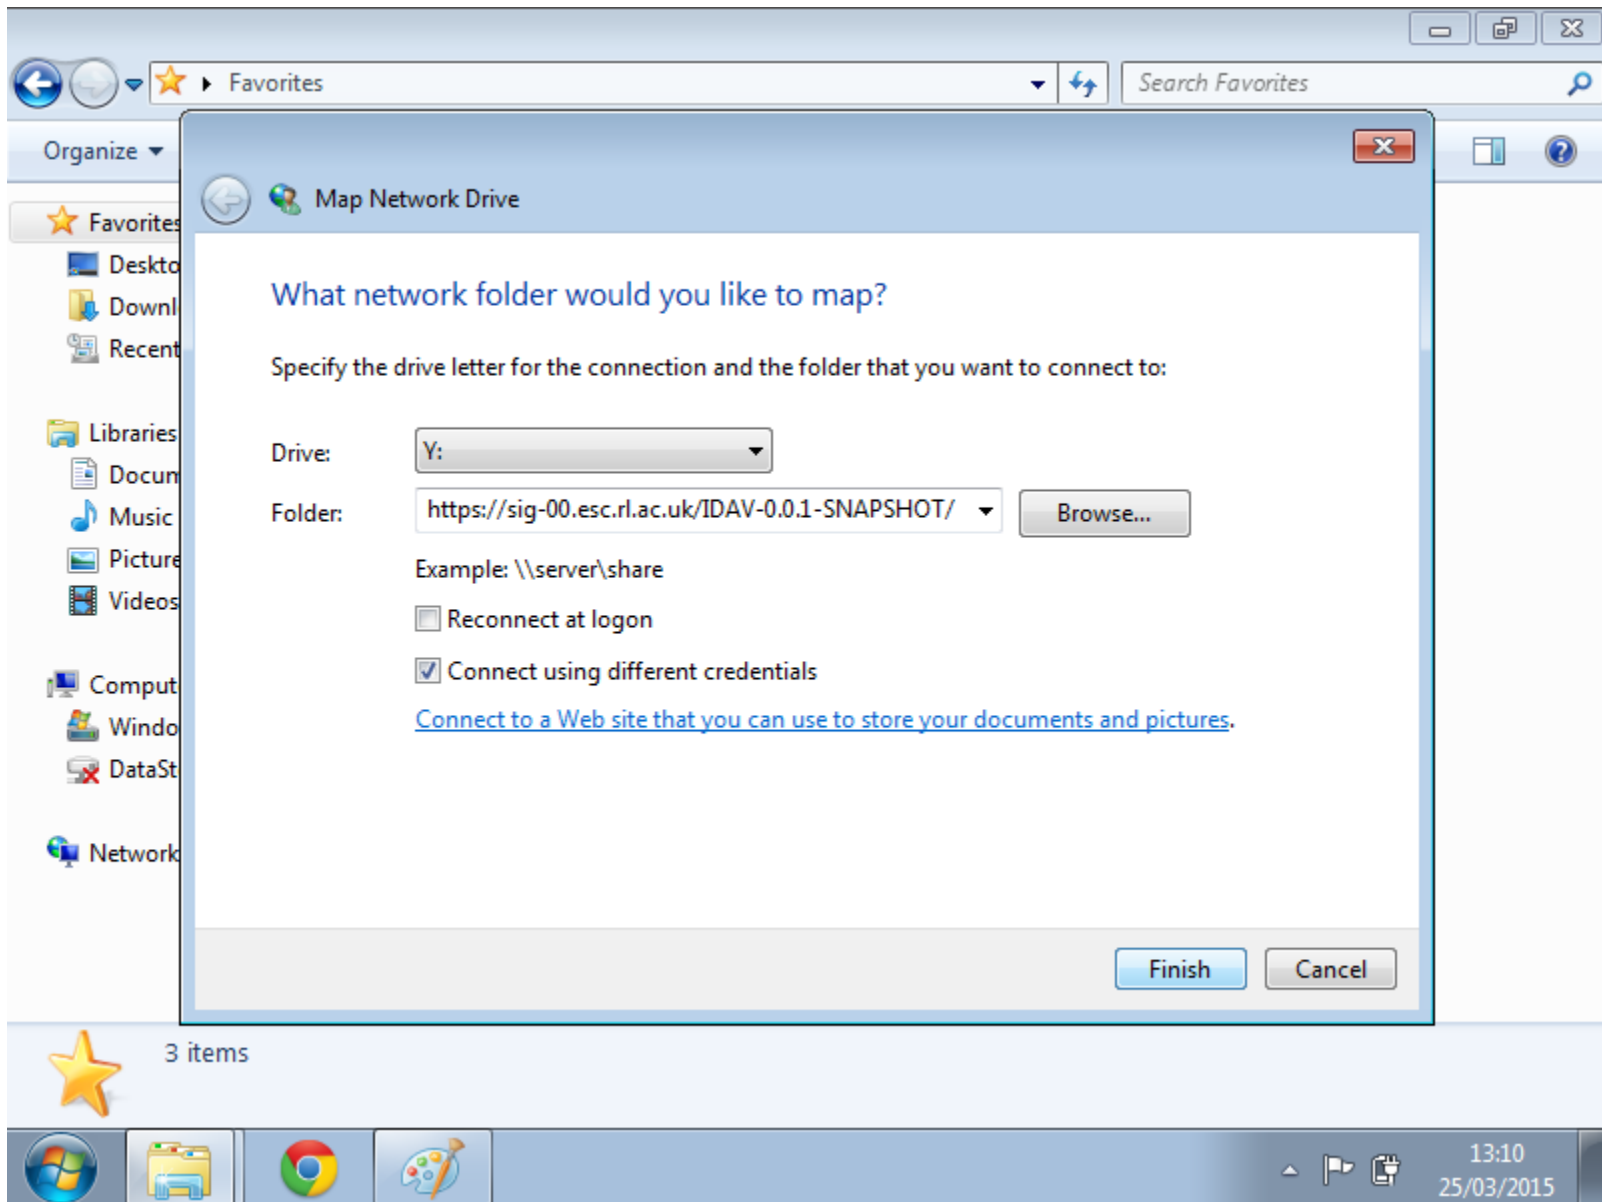
what is it?

- Web Distributed Authoring and Versioning
- An extension of the HTTP protocol
- Mature specification mostly developed between 1996 and 2007
- Allows a file system to be accessed via HTTP
- Built in support in the main operating systems

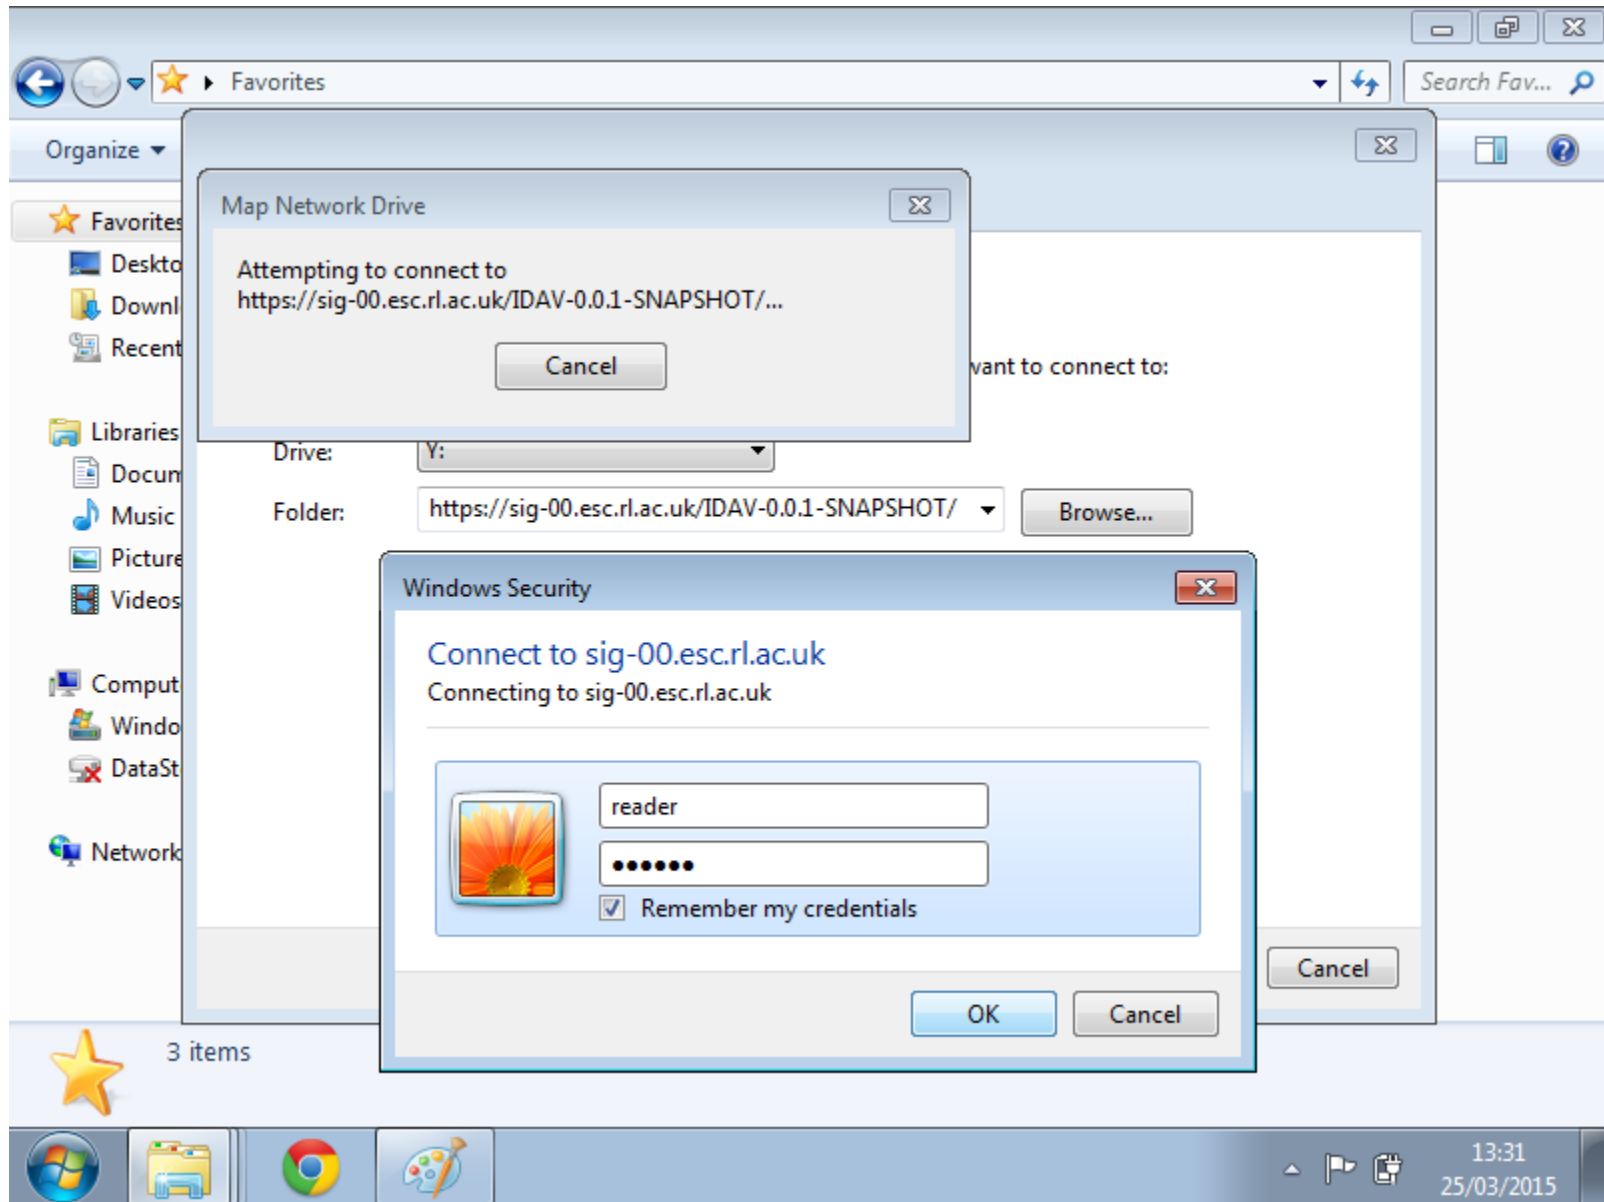
Map Network Drive to WebDAV (Windows)




Enter the URL to the WebDAV webapp



Enter ICAT login credentials



Browse ICAT as a file system



WebDAV with ICAT/IDS

Why?

- A more user friendly route into ICAT/IDS for Artemis laser in the Central Laser Facility
- Data are collected in ad-hoc file structures
- Artemis operator drags and drops folders and files into the relevant experiment folder
- Users can mount ICAT/IDS as a file system remotely

WebDAV with ICAT/IDS

Status

- Fully functional read/write prototype
- Built on 'webdav servlet' project in SourceForge
- Added ~1000 lines of ICAT implementation
- Runs as a webapp in Glassfish
- Tested using Windows Explorer
- Some Linux testing using Nautilus and via a davfs mount
- No Mac testing yet
- Currently only supports Facility, Investigation, Dataset, Datafile ...

WebDAV with ICAT/IDS

Pros and cons

Pros

- Simple drag and drop upload and download
- Built in OS support
- Other 3rd party clients available
- Better sub-Datafile level browsing than TopCAT

Cons

- Not optimised for speed
- Delete your whole ICAT in one mouse click / button press! (with enough permissions)
- Hard to manage slow archive storage

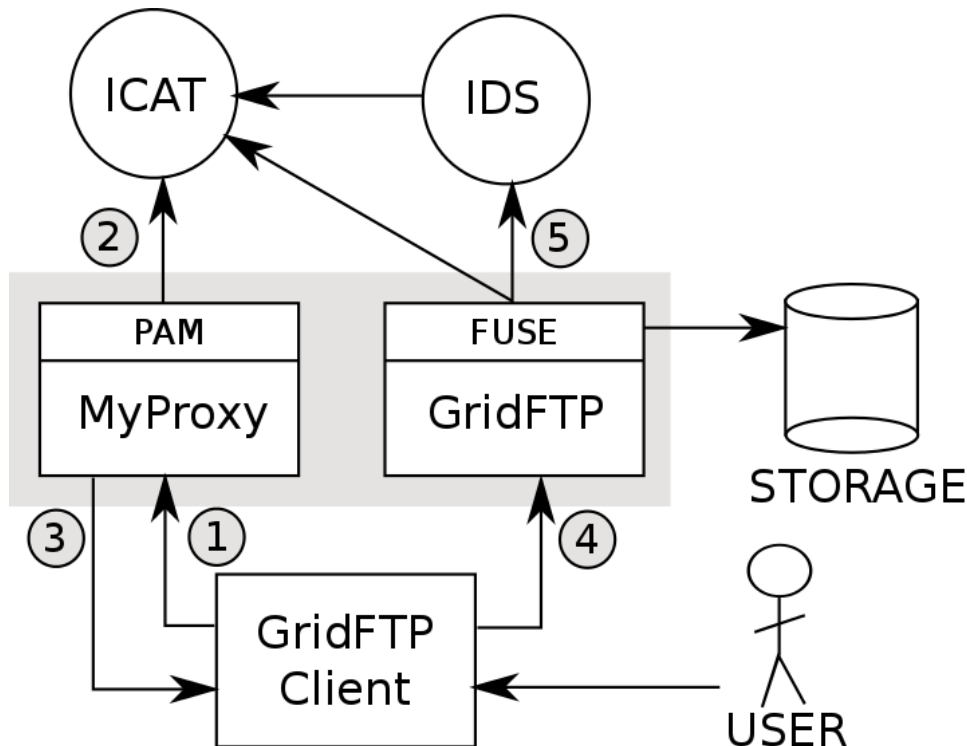
GridFTP

Overview

- Extension of the standard File Transfer Protocol (FTP)
- Designed to enable high performance file transfer of large files
- Used extensively by LHC and many other large scientific facilities

- Start and stop transfers
- Parallel/striped TCP transfers
- UDP transfers
 - mitigates problems of latency and packet loss associated with TCP

GridFTP Architecture



1. User authenticates to MyProxy using ICAT credentials
2. MyProxy authenticates to ICAT using custom ICAT PAM module
3. MyProxy returns proxy certificate to user's GridFTP client
4. User's upload / download requests sent to GridFTP server along with proxy certificate
5. User presented with personal file system view via ICAT FUSE (variant). Files accessed via IDS getLink pointing to locally mounted storage.

GridFTP

Pros and cons

Pros

- Stable GridFTP / MyProxy install using yum package manager
- Popular in scientific communities
- Secure
 - Uses X.509 Public Key Infrastructure
 - Certificate complexity hidden from user


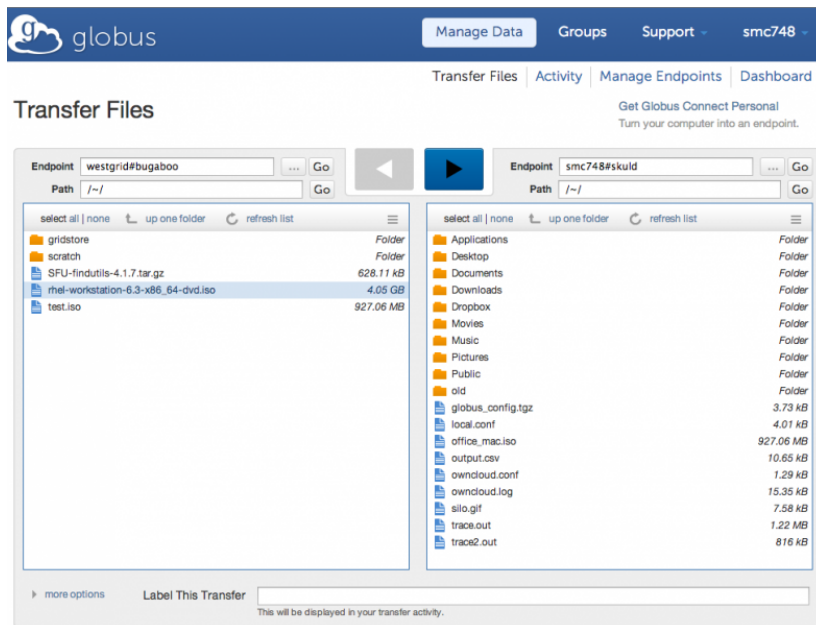
Cons

- Main interaction through command line - globus-url-copy
- No stable GUI client
- Hard to manage slow archive storage

GridFTP

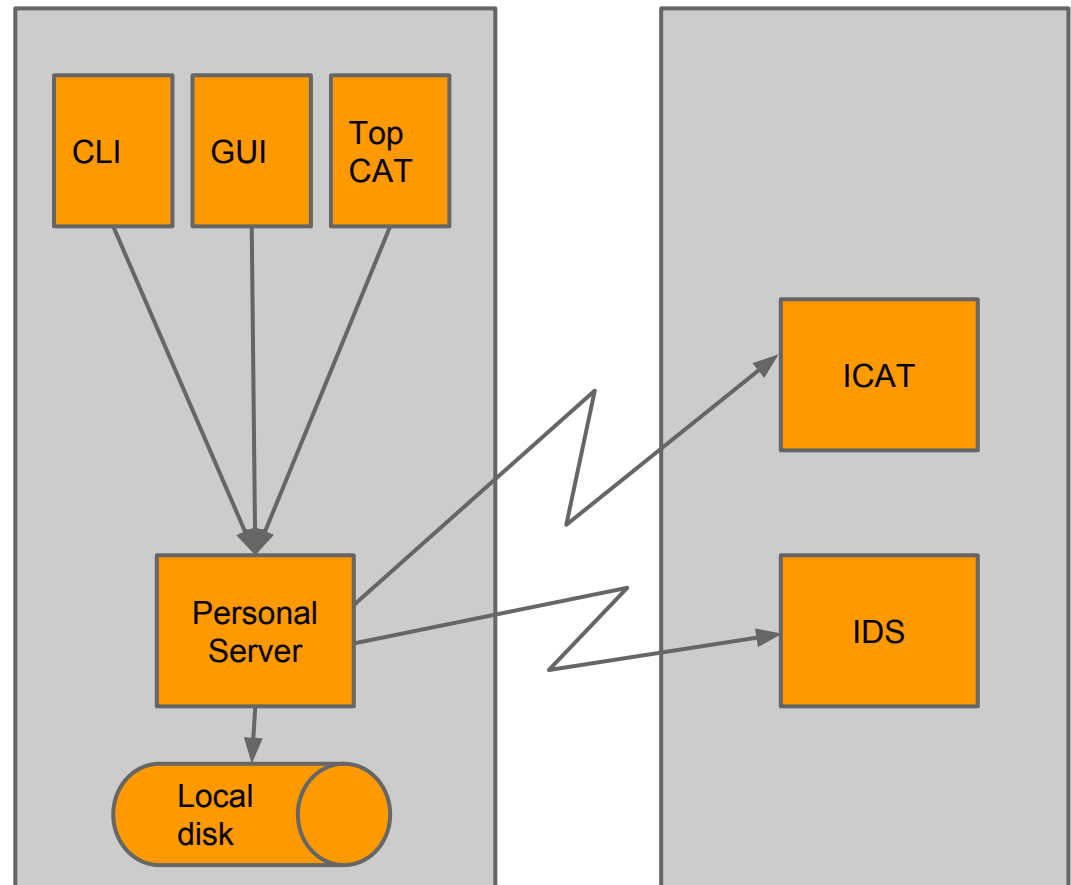
Globus Online

- Non profit organisation provides GridFTP service for scientific communities
- Requires users to have additional Globus Online account
- Web interface a little too generic
- Requires users to download java client



SmartClient

- Client on user's machine and a personal server talking to the IDS
- Client could be command line, GUI or inside TopCAT
- All calls to the client make sure that the personal server is running.
- The personal server accepts the same calls as the IDS to get data



SmartClient

Personal Server

- Belongs to you
 - Only serves your requests
 - Has your uid/gid
- Uses a directory holding files with status information
 - PID and port of server
 - a file for each request
- For each get request:
 - personal server sends request to IDS to bring data back online if necessary
 - expands download request into multiple IDS get requests which are then stored replacing the original request
 - then executes some maximum number of requests in parallel

SmartClient

Personal server requests

Get Investigation with id 17 and store in \$HOME

Get Datafile with id 23 and store in \$HOME

Get Datafile with id 27 and store in \$HOME

Get Datafile with id 31 and store in \$HOME

Get Datafile with id 67 and store in \$HOME

Get Datafile with id 95 and store in \$HOME

Get Datafile with id 12 and store in \$HOME

Get Datafile with id 42 and store in \$HOME

Get Datafile with id 95 and store in \$HOME

Get Datafile with id 12 and store in \$HOME

Get Datafile with id 42 and store in \$HOME

SmartClient

Personal server features

- If an individual transfer fails will continue from where it left off
- Can use as many parallel transfer streams as you choose
 - UDP would require an IDS extension
- Can easily be extended to support upload
- Will provide status information
- If server is stopped it will continue where it left off
- The client could write to crontab (on unix) if permitted for that user to start the server (if not running) at regular intervals otherwise first use of client starts it up
- If useful (e.g. latency is significant) could zip up numerous small files

Personal Server Status

- So far an outline design and some experiments
- The server will be based on `com.sun.net.httpserver.HttpServer`
- Will use JavaFX (now part of Java 8)
 - Includes nice graphics for the GUI
 - Supports rich client packaging to build `.exe`, `.msi`, `.dmg`, `.pkg`, `.rpm` and `.deb`
 - Packages will include a copy of jre to make them self contained
- Will use CORS (Cross-Origin Resource Sharing) so that TopCAT front-end can communicate directly with Personal server

SmartClient

Pros and Cons

Pros

- It solves the problem of data being offline which is an unpleasant characteristic of those solutions offering a file view

Cons

- Though very easy to do it must be installed on client machines