

GridFTP redirection

- Currently, all GridFTP requests are streamed through the Head node
 - Same for PUT
- This is obviously not ideal



GridFTP redirection

- Based on the internal redirection mechanism in Globus' GridFTP server
 - Currently, the redirection happens before asking for a filename, as this was originally intended for load balancing
 - Andrey has found that, with some modifications, it is possible to do the redirection when a file is requested
 - Internals are not well documented, so it is taking time

GridFTP redirection

- Changes are only needed in the DSI plugin
 - Not in the GridFTP server
 - Neither in the clients
 - Any existing client (i.e. UberFTP) will follow transparently the redirection

FTS3 and xrootd

- xrootd 3.3 will include third party copy support on both server and client side
- We already have an **experimental branch** that uses this new functionality
- The server **must** be upgraded to 3.3 for this to work
- External plug-ins (i.e. dpm-xrootd) will need to be adapted

FTS3 and xrootd

- When is it coming?
 - Hopefully, not too long after 3.3 is released
 - Which happened yesterday!
 - Will it make into EPEL?
- What do we already have?
 - Third party copy
 - Checksum validation
 - Size validation

FTS3 and xrootd

- What do we need?
 - Working endpoints (preferably from different vendors) for functional and stress tests
-

FTS3 and http

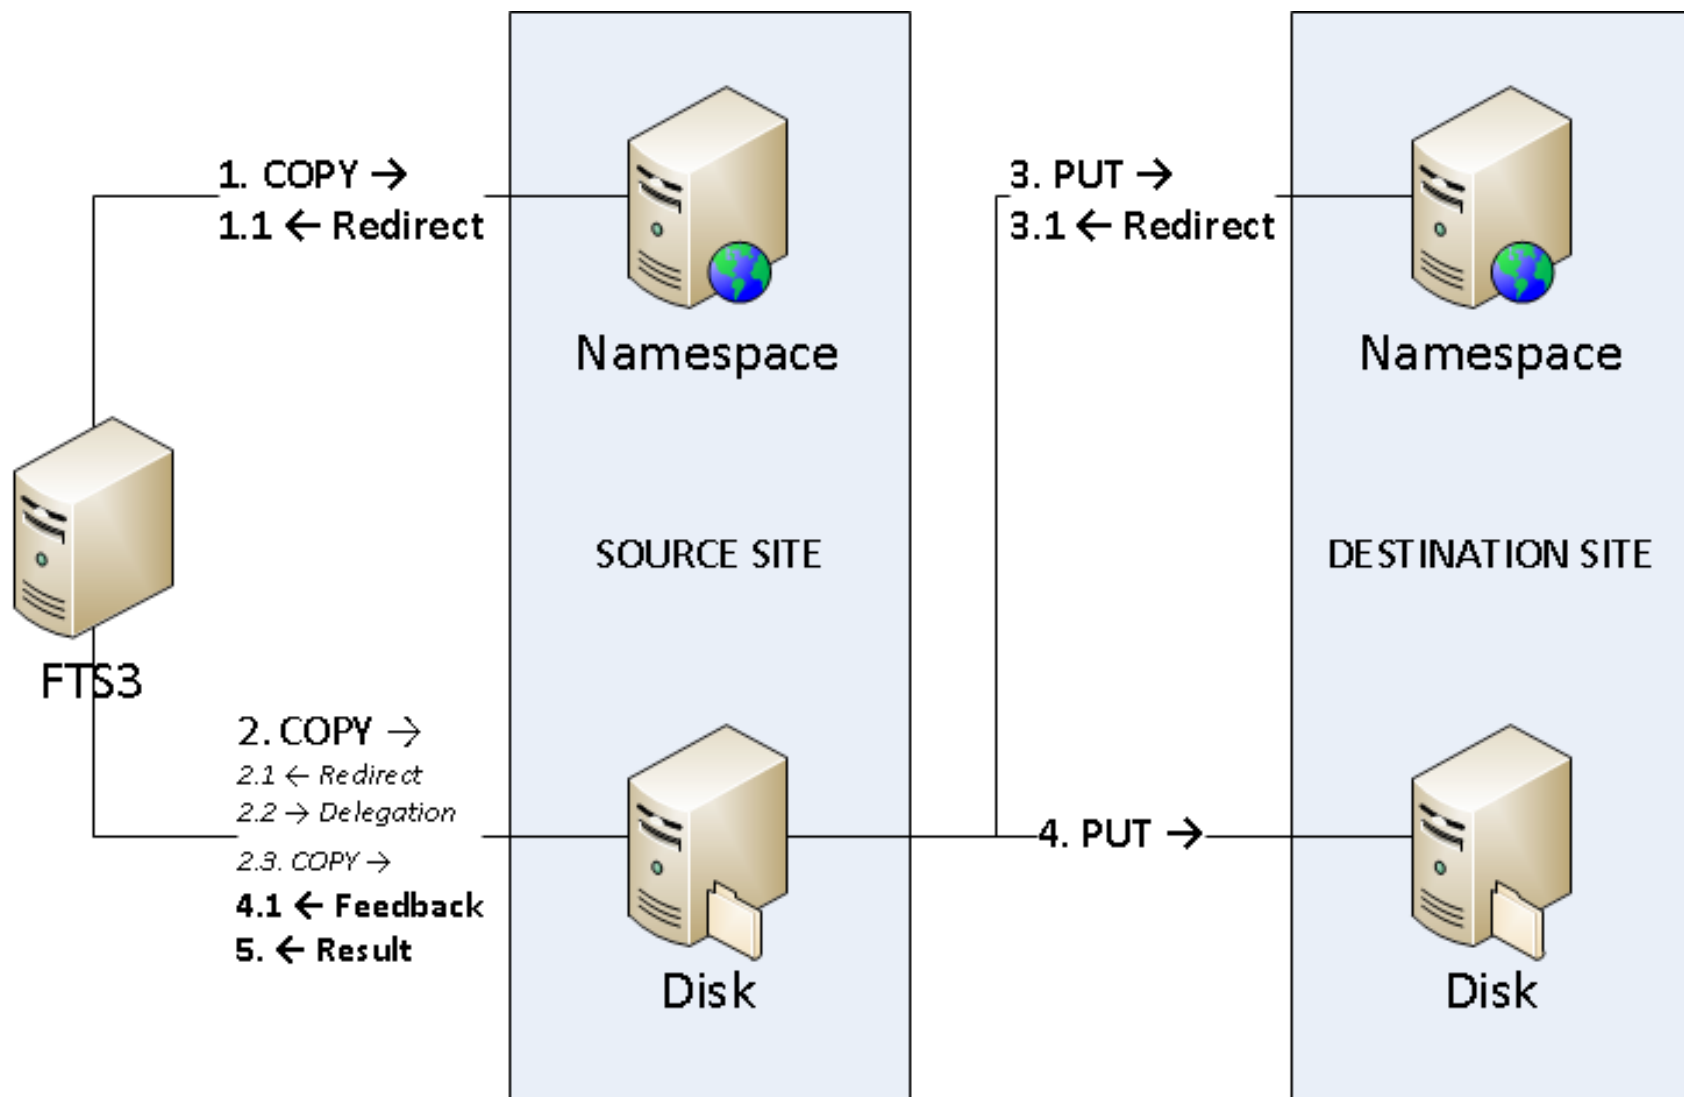
- No pure HTTP standard out there
 - Vendors as Google Cloud Storage or Amazon S3 use their own, not compatible, API's
- WebDAV RFC defines a COPY method, together with a destination header, so we use it instead
 - Although it doesn't allow remote copies, it can be built over it

FTS3 and http

- Since COPY is executed on the source, only a *push* model can be used
 - On the bright side, the destination could be *any* http(s) endpoint that supports PUT
 - The source SE orchestrates
- Delegation is an issue
 - This bit is not standard, and will not work with out-of-the-box clients
 - A mechanism agreed between DPM and dCache

FTS3 and http (Delegation)

- The endpoint redirects to itself (307) with a *X-Delegate-To* header **only if** needed.
- The client does a Gridsite delegation to the URL specified with *X-Delegate-To*, and follow the redirect
- The endpoint performs the copy
 - If no delegation was done, abort (the client may have not recognized the header)



FTS3 and http

- GridFTP performance markers will be used to give feedback to the client
 - It is an already known format
- The connection **must** stay alive during the copy
 - If the connection is closed, the copy will be canceled

FTS3 and http

- When is it coming?
 - Soon after DPM and/or dCache have a working implementation
- What do we already have?
 - A working prototype
 - No checksum validation, since http endpoints do not support this yet (DPM working on it)
- What do we need?
 - A working implementation for both third party copy **and** checksum querying though http
 - At least a couple of working endpoints
 - Again, preferably from different vendors

FTS3, xrootd and http

- One last comment:
 - FTS3 uses GFAL2
 - GFAL2 is plug-in based
 - So xrootd and/or http third party copies can be enabled without upgrading FTS3 itself.