



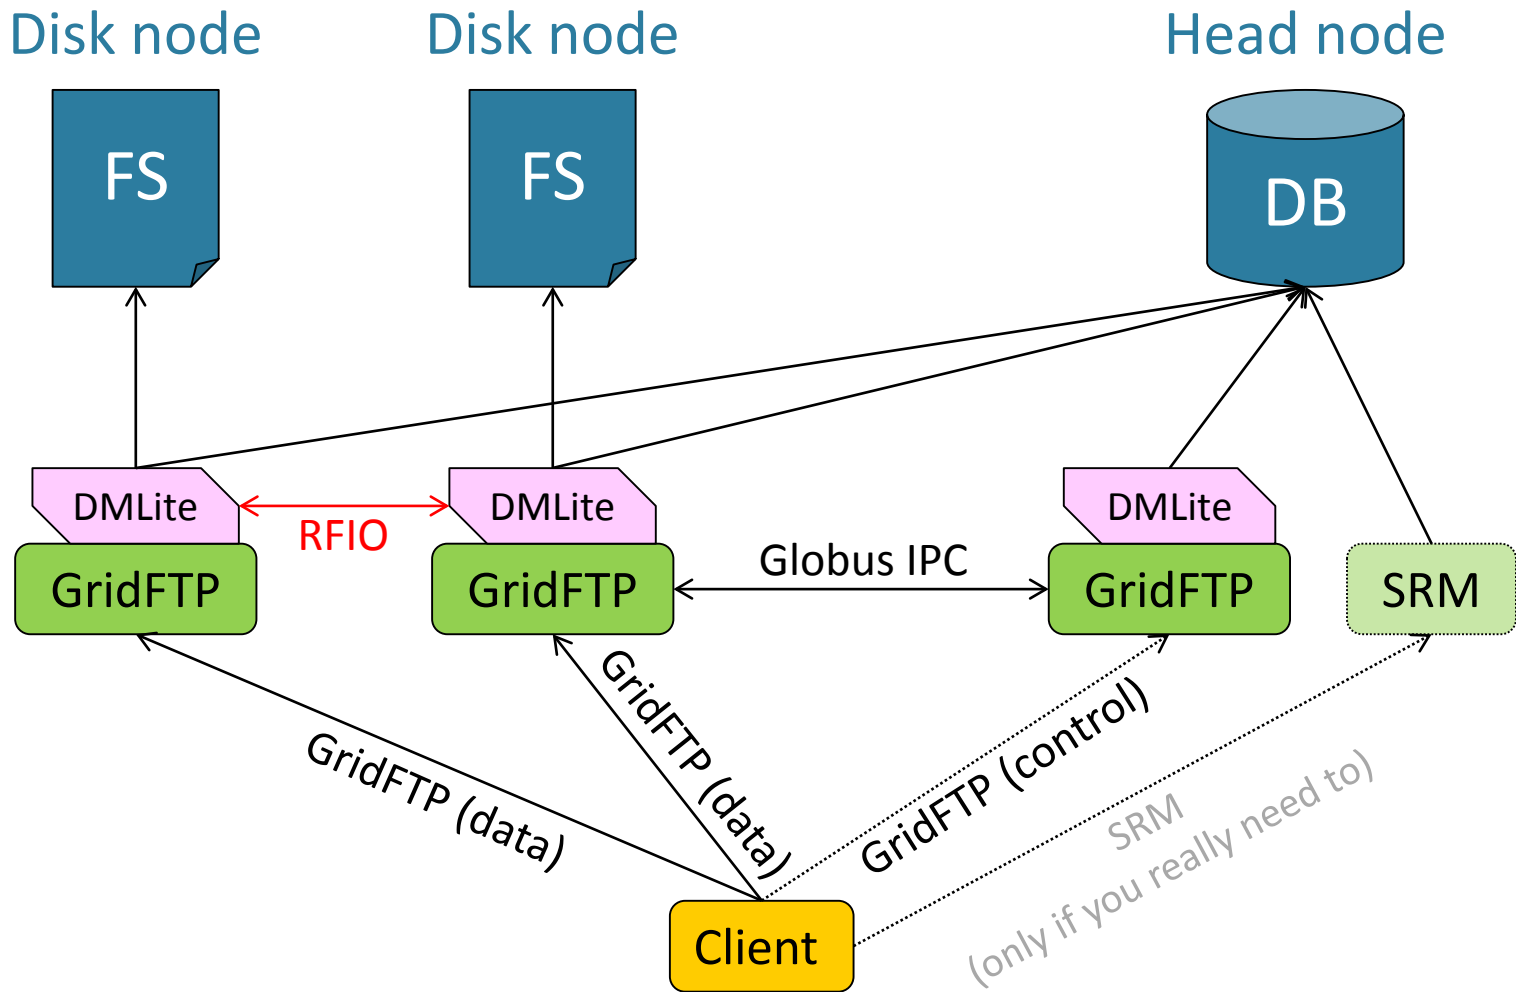
GridFTP frontend for DMLite status update

Andrey Kiryanov
DPM Workshop, CERN

7/12/2015



GridFTP with redirection



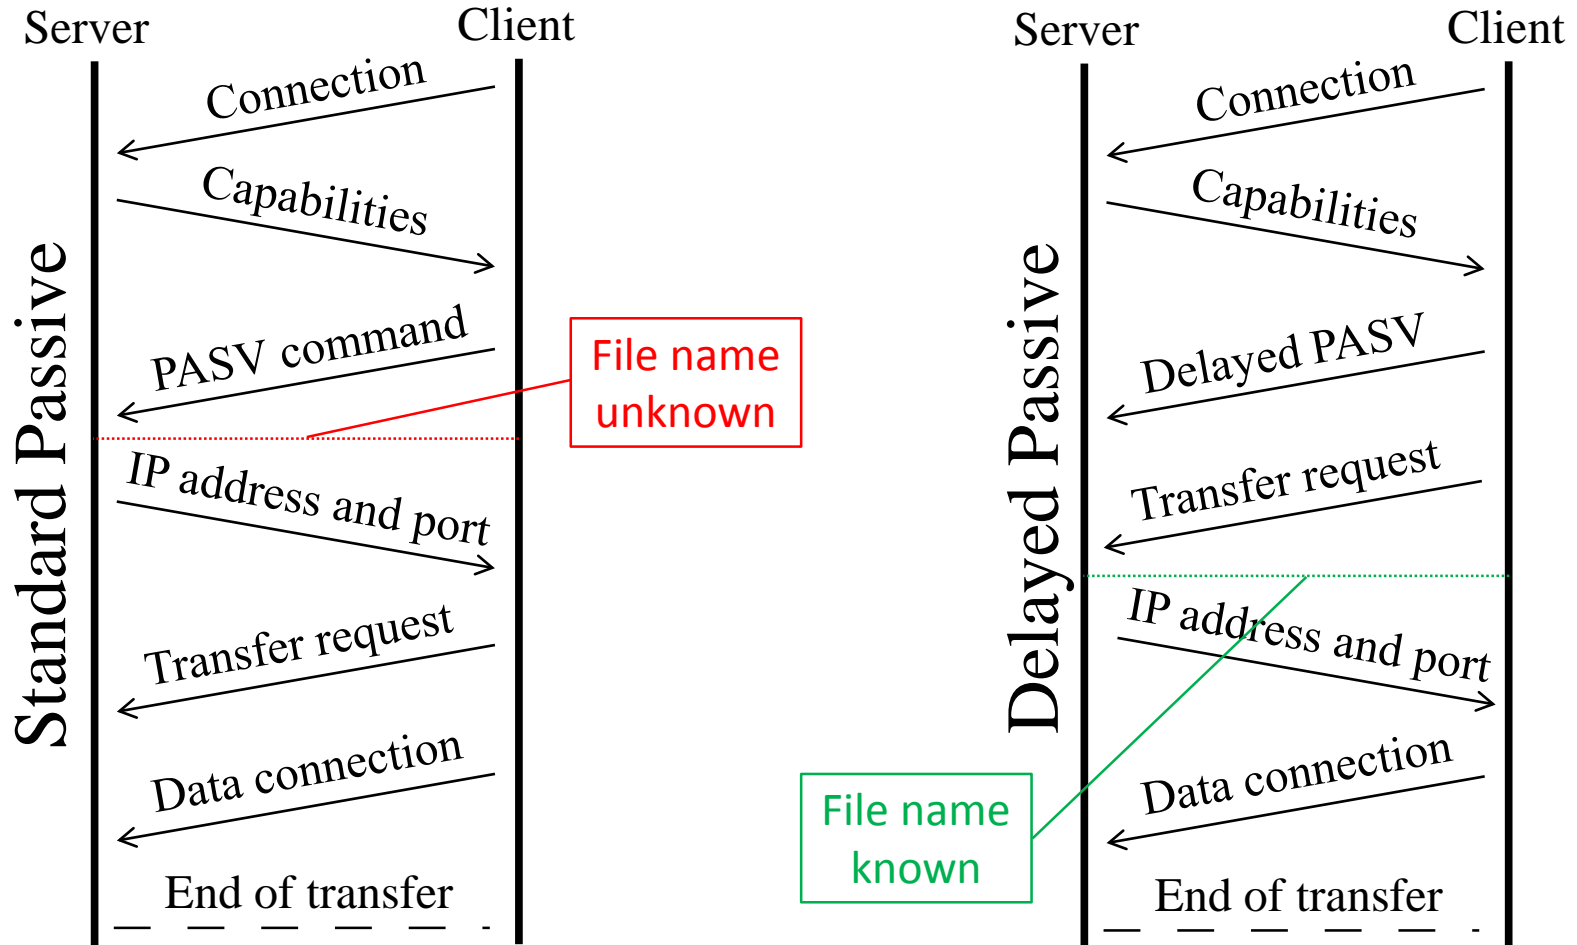
Manual configuration reference

- **Disk nodes:** add **data_node 1** option to `/etc/gridftp.conf` and restart GridFTP
 - Don't be surprised that you won't be able to directly access these nodes with GridFTP clients anymore
- **FTP head node:** add **remote_nodes <list>** option to `/etc/gridftp.conf` and restart GridFTP
 - **<list>** is a comma-separated list of disk node FQDNs with ports (e.g. `disk1.foo.bar:2811,disk2.foo.bar:2811`)
- **SRM head node:** add **DPM FTPHEAD <head>** option to `/etc/shift.conf` and restart SRM
 - **<head>** is FQDN of the FTP head node

Redirection facts

- Needs *Puppet* for configuration (The YAIM is dead, hail the YAIM)
- Head node cannot be used as storage
- No direct control connections to disk nodes
- Clients do not need to contact SRM, but if they do, a proper TURL pointing to a head node is given
- Clients need to support Delayed Passive mode for optimal performance, the ones that fail to do so will face the dark side of RFIO
 - Avoid using *lcg-util* or *uberftp* as both of them will result in performance penalty even with SRM

Standard Passive vs. Delayed Passive



Recent problems and fixes

- Early versions of DPM-DSI 1.9.5 had a race bug in command processing which sometimes caused next command to be accepted before finishing putDone()
 - ✓ Fixed in 1.9.5
- In DPM-DSI 1.9.5 checksum calculation was performed on a head node, involving RFIO reads from disk nodes
 - ✓ Fixed in 1.9.7
- Globus releases new versions of GridFTP server too often. We have to re-compile/re-release DPM-DSI with every new version because we depend on internal ABI.
 - **Cannot do much about it. Need to extend Globus API.**

Current status

- We have GridFTP with redirection on our testbeds under continuous testing for almost a year now
- It is fully supported by FTS3 and *gfal2-util* CLI
- It is seen as the most useful way of deploying GridFTP with DMLite-based storage
- It renders SRM (along with its overhead) unnecessary for most typical use cases

HTTP/S3 support

Latest GridFTP supports 3rd-party copy to/from HTTP/S3 with extended set of commands:

- HTTP CONFIG CA_CERTS=<base64>
 - Alternative CA bundle for HTTPS
- HTTP UPLOAD OFFSET=<int>; LENGTH=<int>; PATH=<path>; REQUEST=<req>
- HTTP DOWNLOAD OFFSET=<int>; LENGTH=<int>; PATH=<path>; REQUEST=<req>
 - 3rd party transfer with provided HTTP request <req>

No client uses it yet and documentation is almost missing

That's all
Thank you!