

I

XrootD and GridFTP development plans at RAL

Alastair Dewhurst, George Vasilakakos, Ian
Johnson

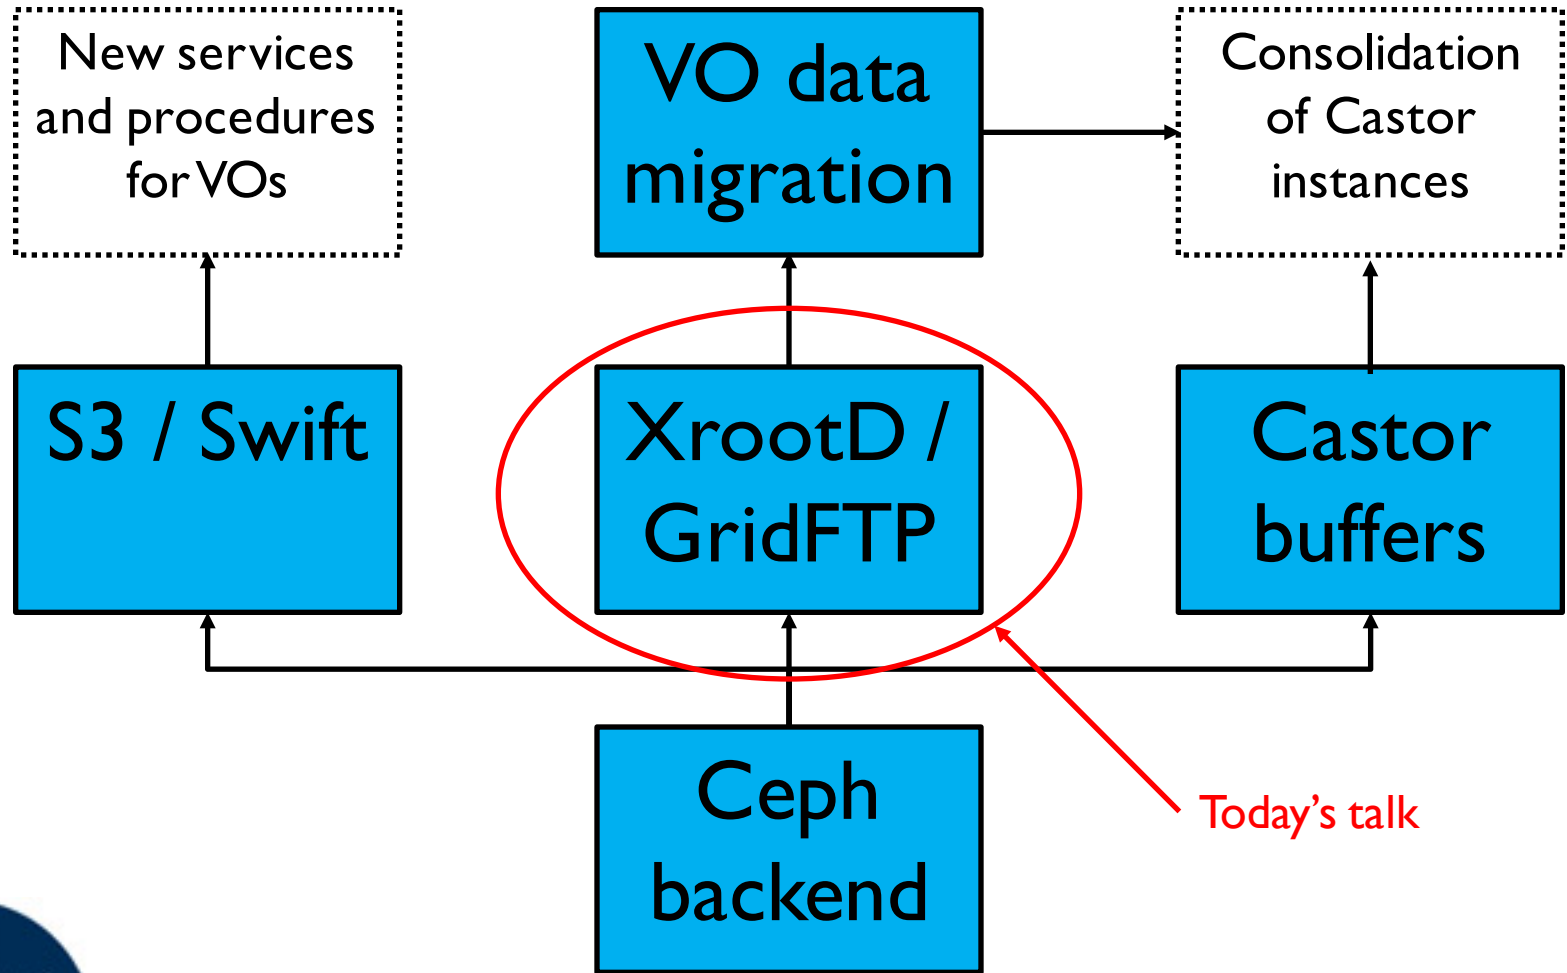


Outline

- Motivation
- XrootD authorization
- GridFTP
 - Authorization
 - Performance
- Leading /
- Summary



The Echo Project



Today's talk

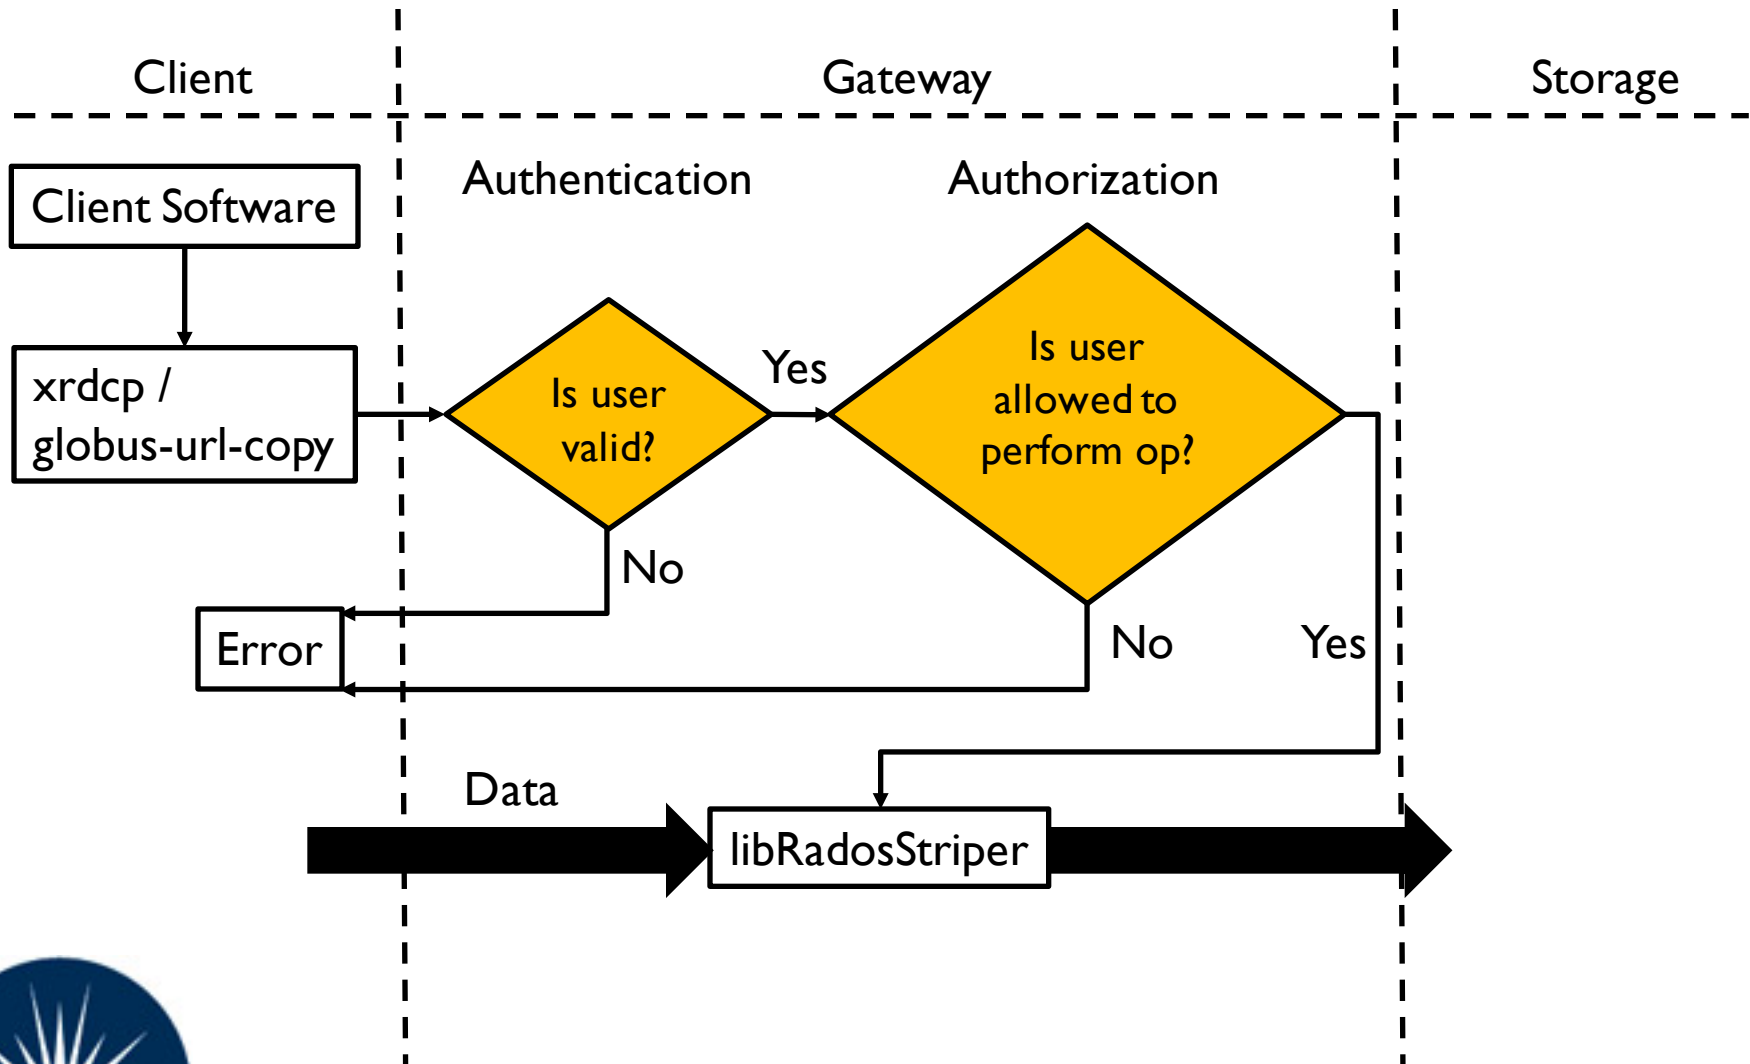


Motivation

- The LHC VO are not currently able to run all their workflows through an SE providing only S3/Swift access.
- Minimum requirement to provide GridFTP and XrootD access to storage.
 - Aim to create thinnest layer on top of Ceph (why CephFS + GridFTP servers are not favoured).
 - Only intend to provide needed functionality for existing use cases.
- Relying heavily on work done by Sebastien Ponce and Brian Bockelman.



Authorization architecture

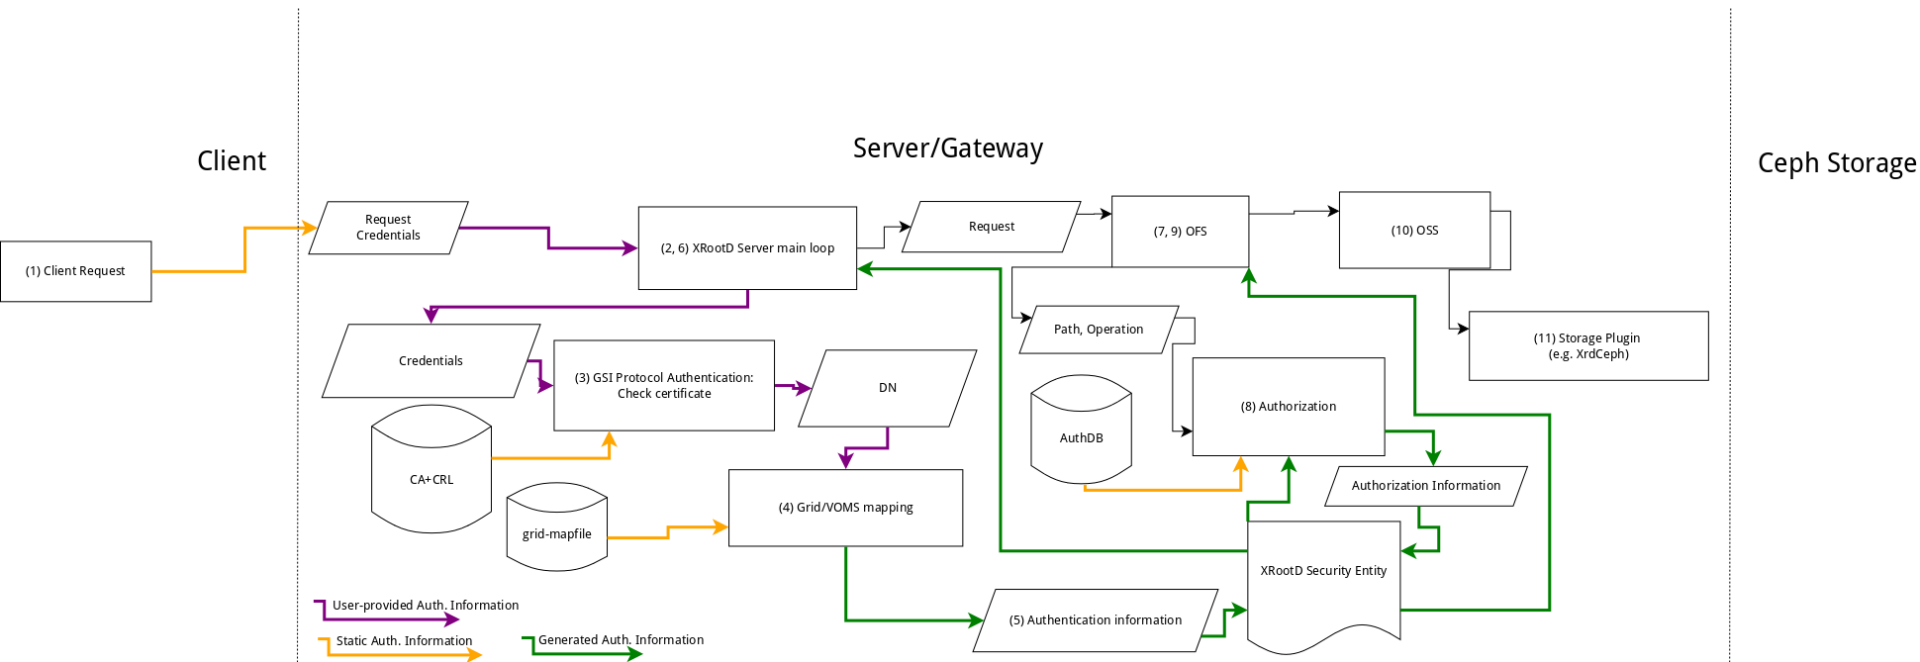


For XrootD the Gateway could be the WN



XrootD authorization (I)

- "Just" need to develop authorization code for XrootD.
- Grid-map file + AuthDB are existing components we are trying to reuse.



XrootD authorization (2)

Xrootd.conf
all.export *?
all.export /

authDB
u atlasprod /test a

Created pool called /test

grid-mapfile
"/C=UK/O=eScience/OU=CLRC/L=RAL/CN=alastair dewhurst" atlasprod

```
-bash-4.1$ xrdcp 128m root://ceph-gw1.gridpp.rl.ac.uk/test:/128m
```

```
[0B/0B][100%][=====][0B/s]
```

```
Run: [ERROR] Server responded with an error: [3010] Unable to create test:/128m; permission denied
```

```
-bash-4.1$ xrdcp 128m root://ceph-gw1.gridpp.rl.ac.uk//test:128m
```

```
[0B/0B][100%][=====][0B/s]
```

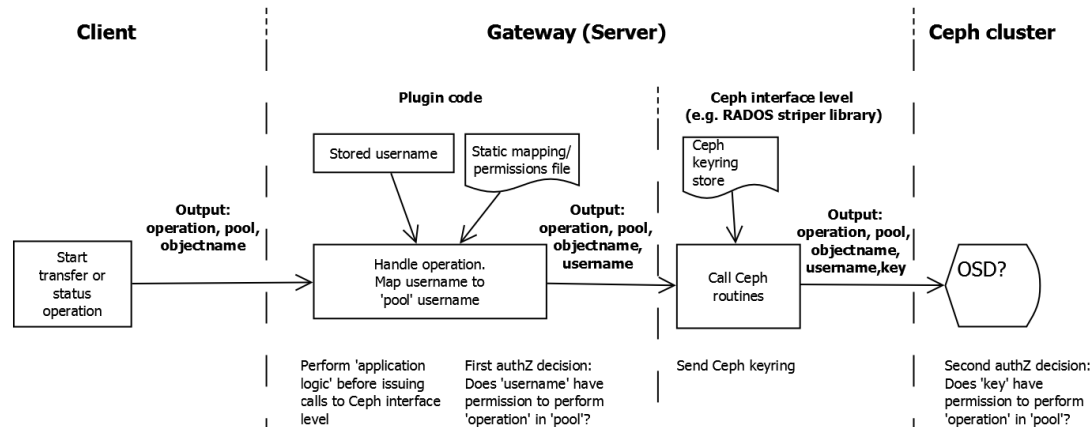
```
Run: [ERROR] Server responded with an error: [3005] Unable to open /test:128m; invalid argument
```



GridFTP authorization

- Will need to build GridFTP authorization.
- If possible re-use code.
- Must share Grid-map file and authDB.
- GridFTP authorization development depends on how we get XrootD authorization to work.

2b. Proposed authorization - to be performed for every operation



Updated plan: Only use one Ceph keyring



GridFTP Performance (I)

	Rados get/put	Globus-url-copy	FTS transfers
Reads	~80MB/s	~80MB/s	6-7MB/s
Writes	~80MB/s	~80MB/s	6-7MB/s

- Current plugin based on proof of concept work by Sebastien.
- Acceptable performance if using Globus-url-copy directly (Streaming).
- Poor performance using FTS (block mode).
 - Tried to get larger chunks to be sent but this is not viable over WAN.



GridFTP Performance (2)

10



- Ian Johnson is working on this currently.
- Based on Brian Bockelman's GridFTP code for HDFS
 - https://github.com/bbockelm/gridftp_hdfs/tree/master/src



Leading /

- There are several different issues with leading /.

1. Unable to delete paths (with XrootD) not starting with a /

- <https://github.com/xrootd/xrootd/issues/314>

2. XrootD authDB requires / else path is considered a template.

- Work around means creating a pool starting with a /.

3. GridFTP plugin adds a leading /.

- Ian Johnson has removed this from plugin.



Conclusion

- A lot of effort has gone into understanding XrootD
 - Still trying to decide on best solution.
- GridFTP plugin development ongoing.
- Concern about leading / preventing objects being accessible via both protocols.

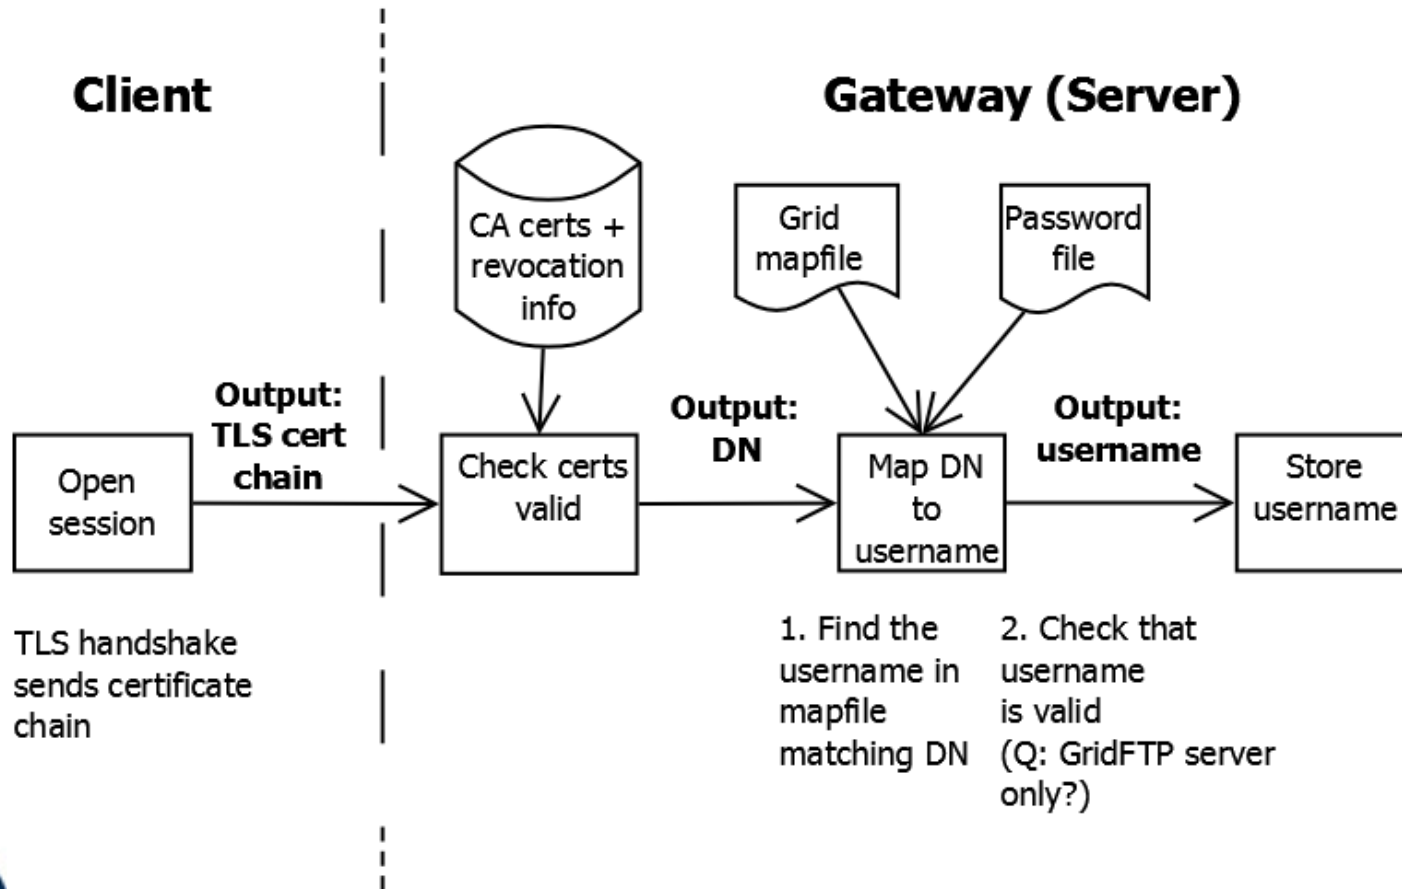


Backup



Backup

1. GSI Authentication - performed at session start-up



Backup

2a. Current GridFTP authorization - performed for every operation

