HTTP Protocol Update

Brian Bockelman,
5 December 2018

Transfer Matrix

- Things have been looking fairly good:
  - Huge improvement in Lancaster DPM — thanks all!
  - Glasgow is perhaps the next step?
  - Something is fishy with INFN StoRM.

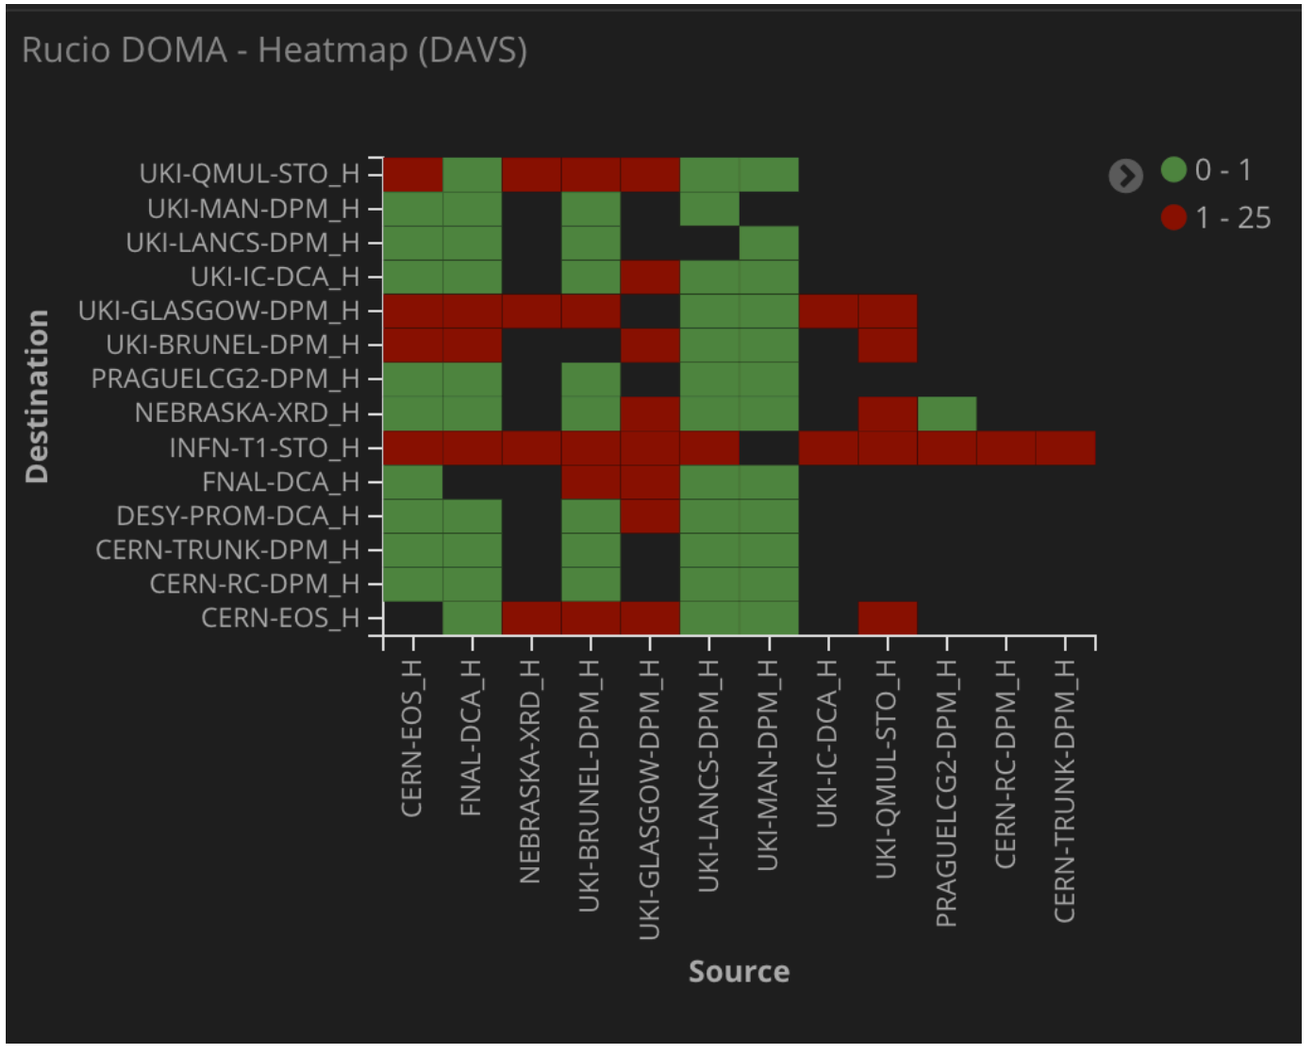

INFN StoRM

- I see the following log lines from FTS:

```
INFO    Wed, 05 Dec 2018 15:17:35 +0100; bytes: 3747840, avg KB/sec:0, inst KB/sec:0, elapsed:2484000
```

- It does seem that the transfer times out - that transfer from FNAL averages 1.5KB/s. Maybe something needs a kick?

- However, “elapsed” time appears to be off by 1000x: potentially StoRM is reporting incorrect units?

Work remaining

- Xrootd 4.9.0-rc1 is tagged and available in the osg-testing repository.
  - Native Macaroon and HTTP TPC support.
  - Late-arriving patch will enable HTTP request pipelining for multi-stream transfers.
- Still working on OAuth2-based token acquisition for tokens issued by the storage server. Not much to report other than we haven't forgotten about it!
- Any updates from Echo?
- Any updates from EOS?

Thoughts on Scale Tests

- Starting to prepare sanity tests for scale tests:
  - Uploaded 240GB of files to FNAL.
  - Prepared a simply FTS copy-job for doing all 240GB in one command.
  - Comparing WebDAV-vs-GridFTP on the same hardware.
- Rather boring results currently:
  - Can increase concurrency until hardware limits are reached.
  - No significant throughput differences between GridFTP and WebDAV.
  - One-off transfers on a link with longer RTT (DESY->Nebraska) also shows no significant performance difference.

Rucio-based Scale Tests

- We proposed the following questions:
    - When is an endpoint “working” for scale tests?
        - **Proposal**: when admins say it is ready and it can demonstrate 7 days of successful transfers (>90% success) for a protocol.
        - May not be “perfect green” because other endpoints can have problems, of course.
    - How do we stress test an endpoint?
        - **Proposal**: Upload one-or-more 1TB datasets to each source endpoint and repeatedly transfer it (“transfer-delete-repeat”).
    - How is an endpoint placed under increased stress?
        - **Proposal**: Let FTS manage concurrency of the Rucio transfers. Establish a baseline for the whole matrix and, as necessary, inject more concurrent source datasets. To be reviewed every 2 weeks at this meeting.
    - What criteria decides if an endpoint is overloaded?
        - **Proposal**: Admin complaints, >10% failure rate, endpoint crashes, or errors FTS decides are due to overload (FTS should self-adopt).