



Quantifying XRootD Overhead

De Witt, Shaun (STFC); Lahiff, Andrew (STFC)

Understanding The Test Architecture

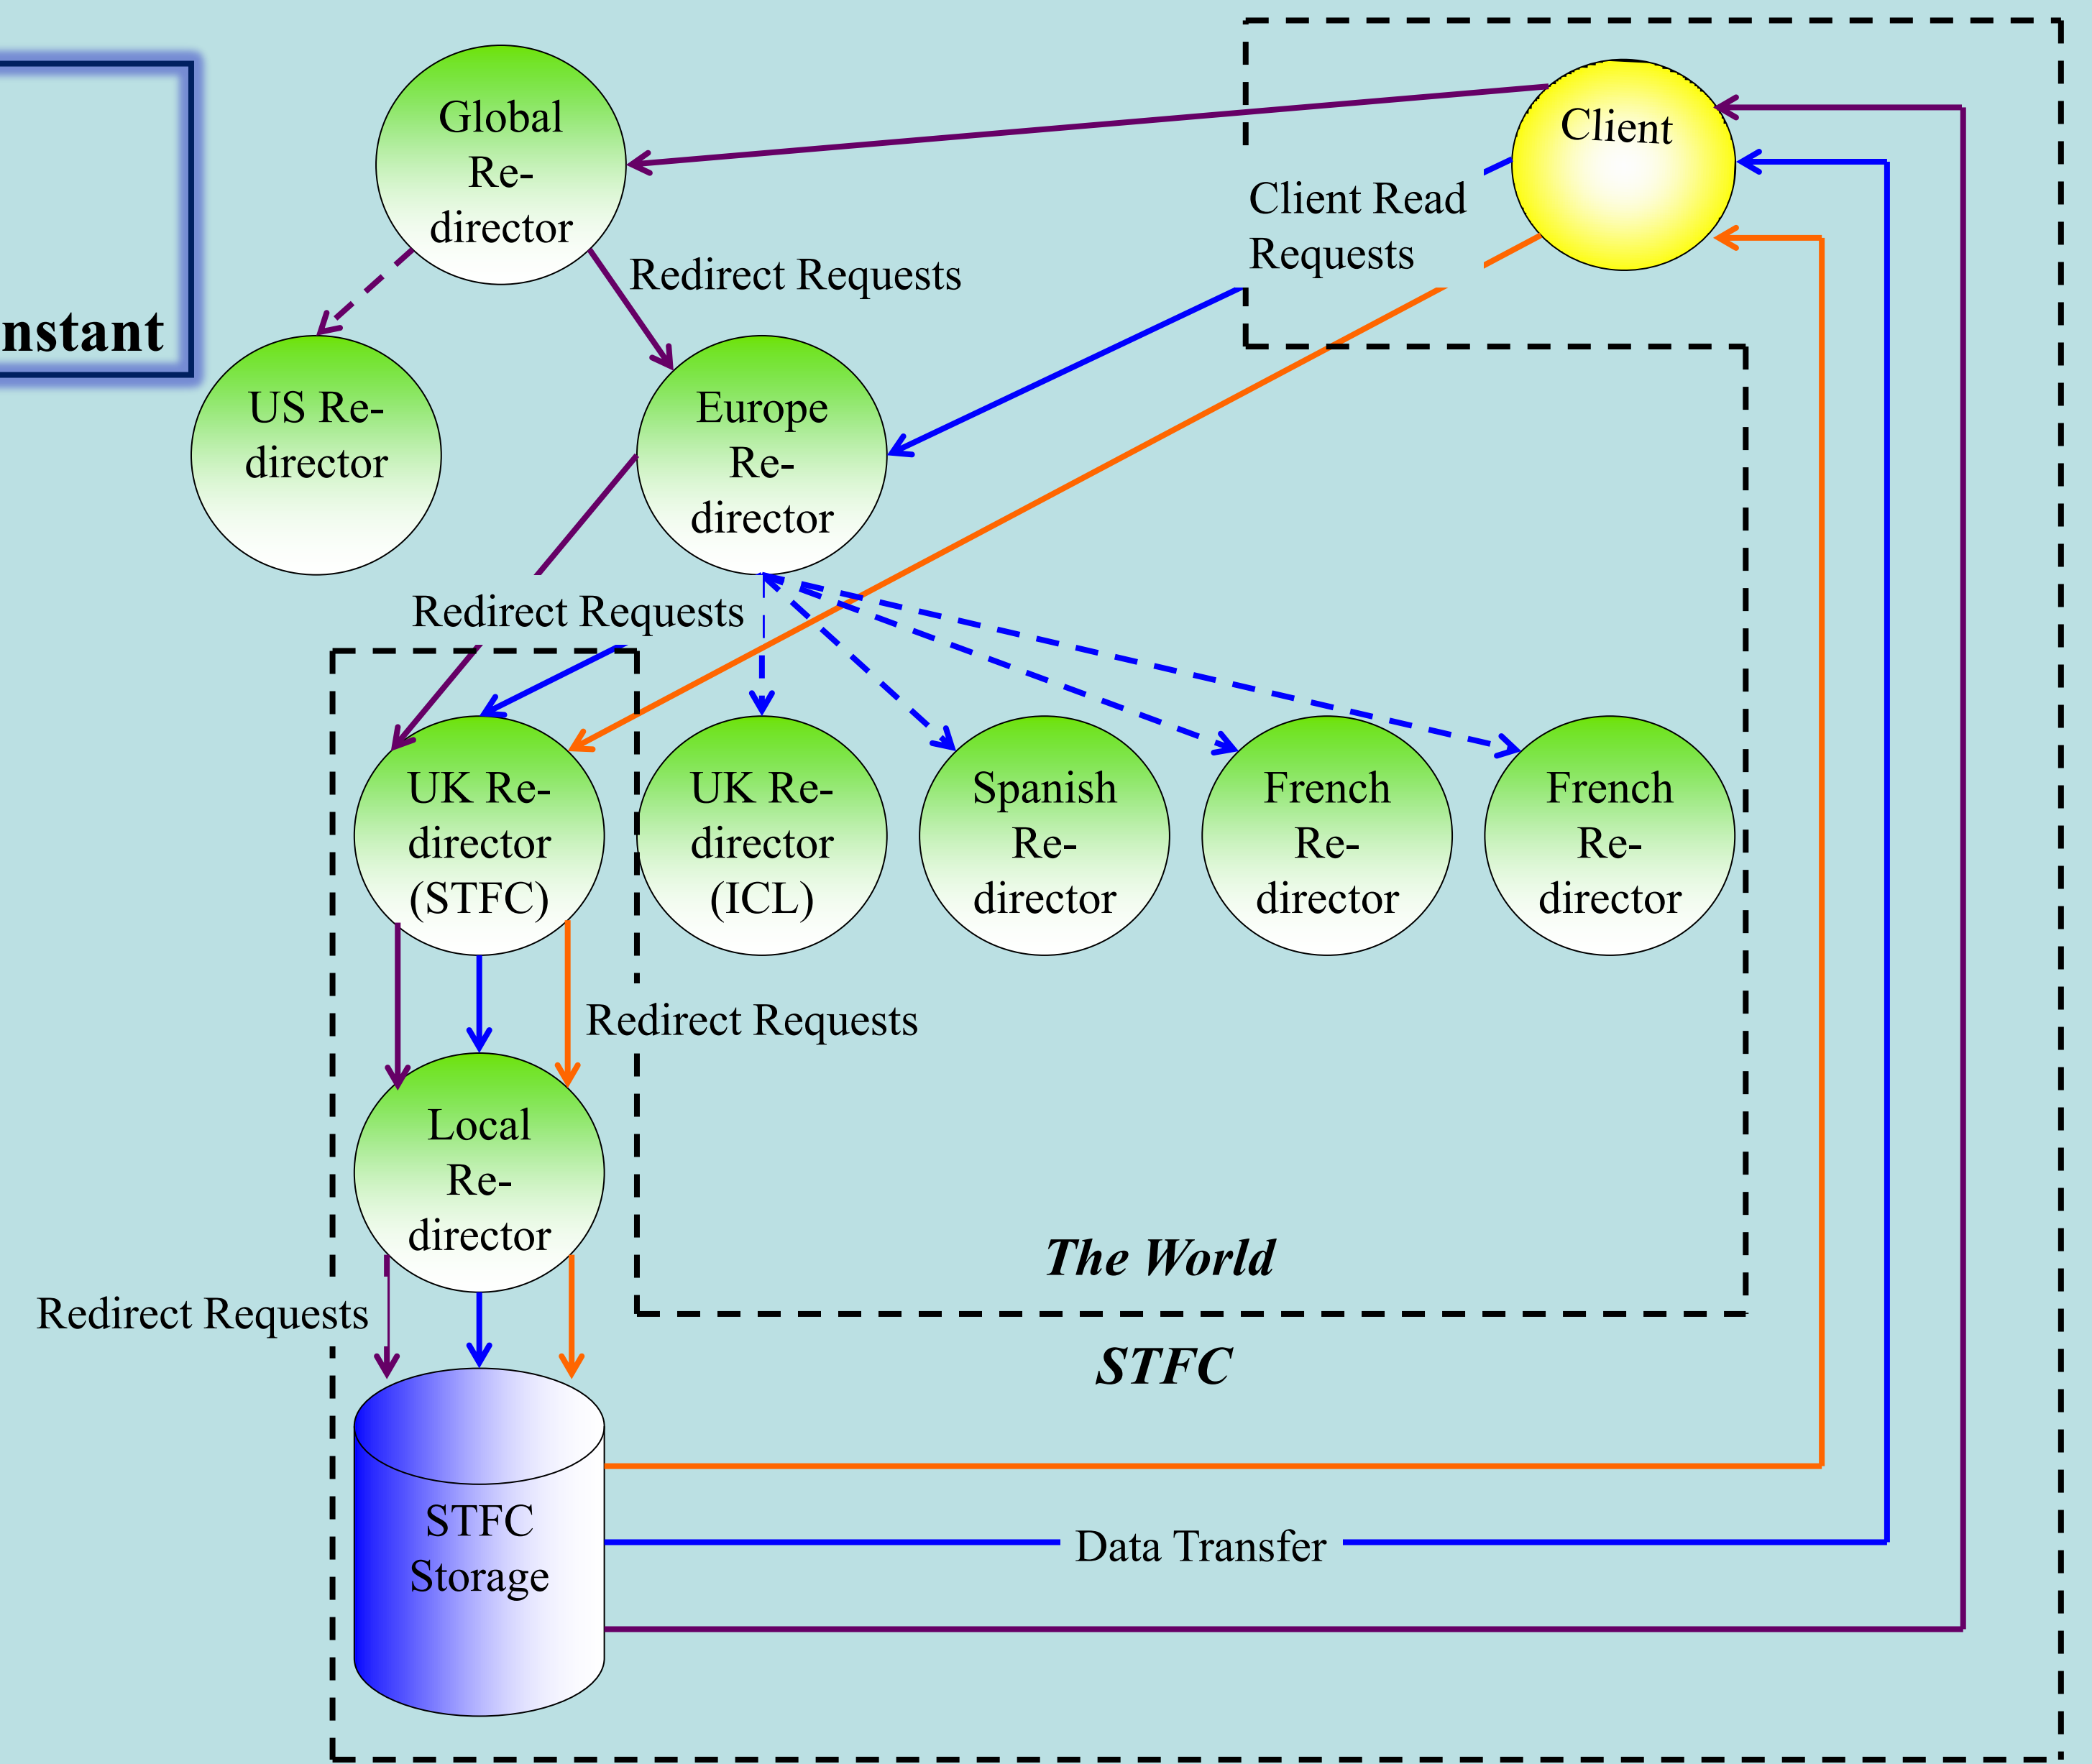
- **Client and Data** located at STFC
- **Client** users xroot protocol (*xrdcp*) for all tests
- Data Transfer is **always** within STFC
 - Transfer Times should be (approximately) **constant**

Scalability Tests

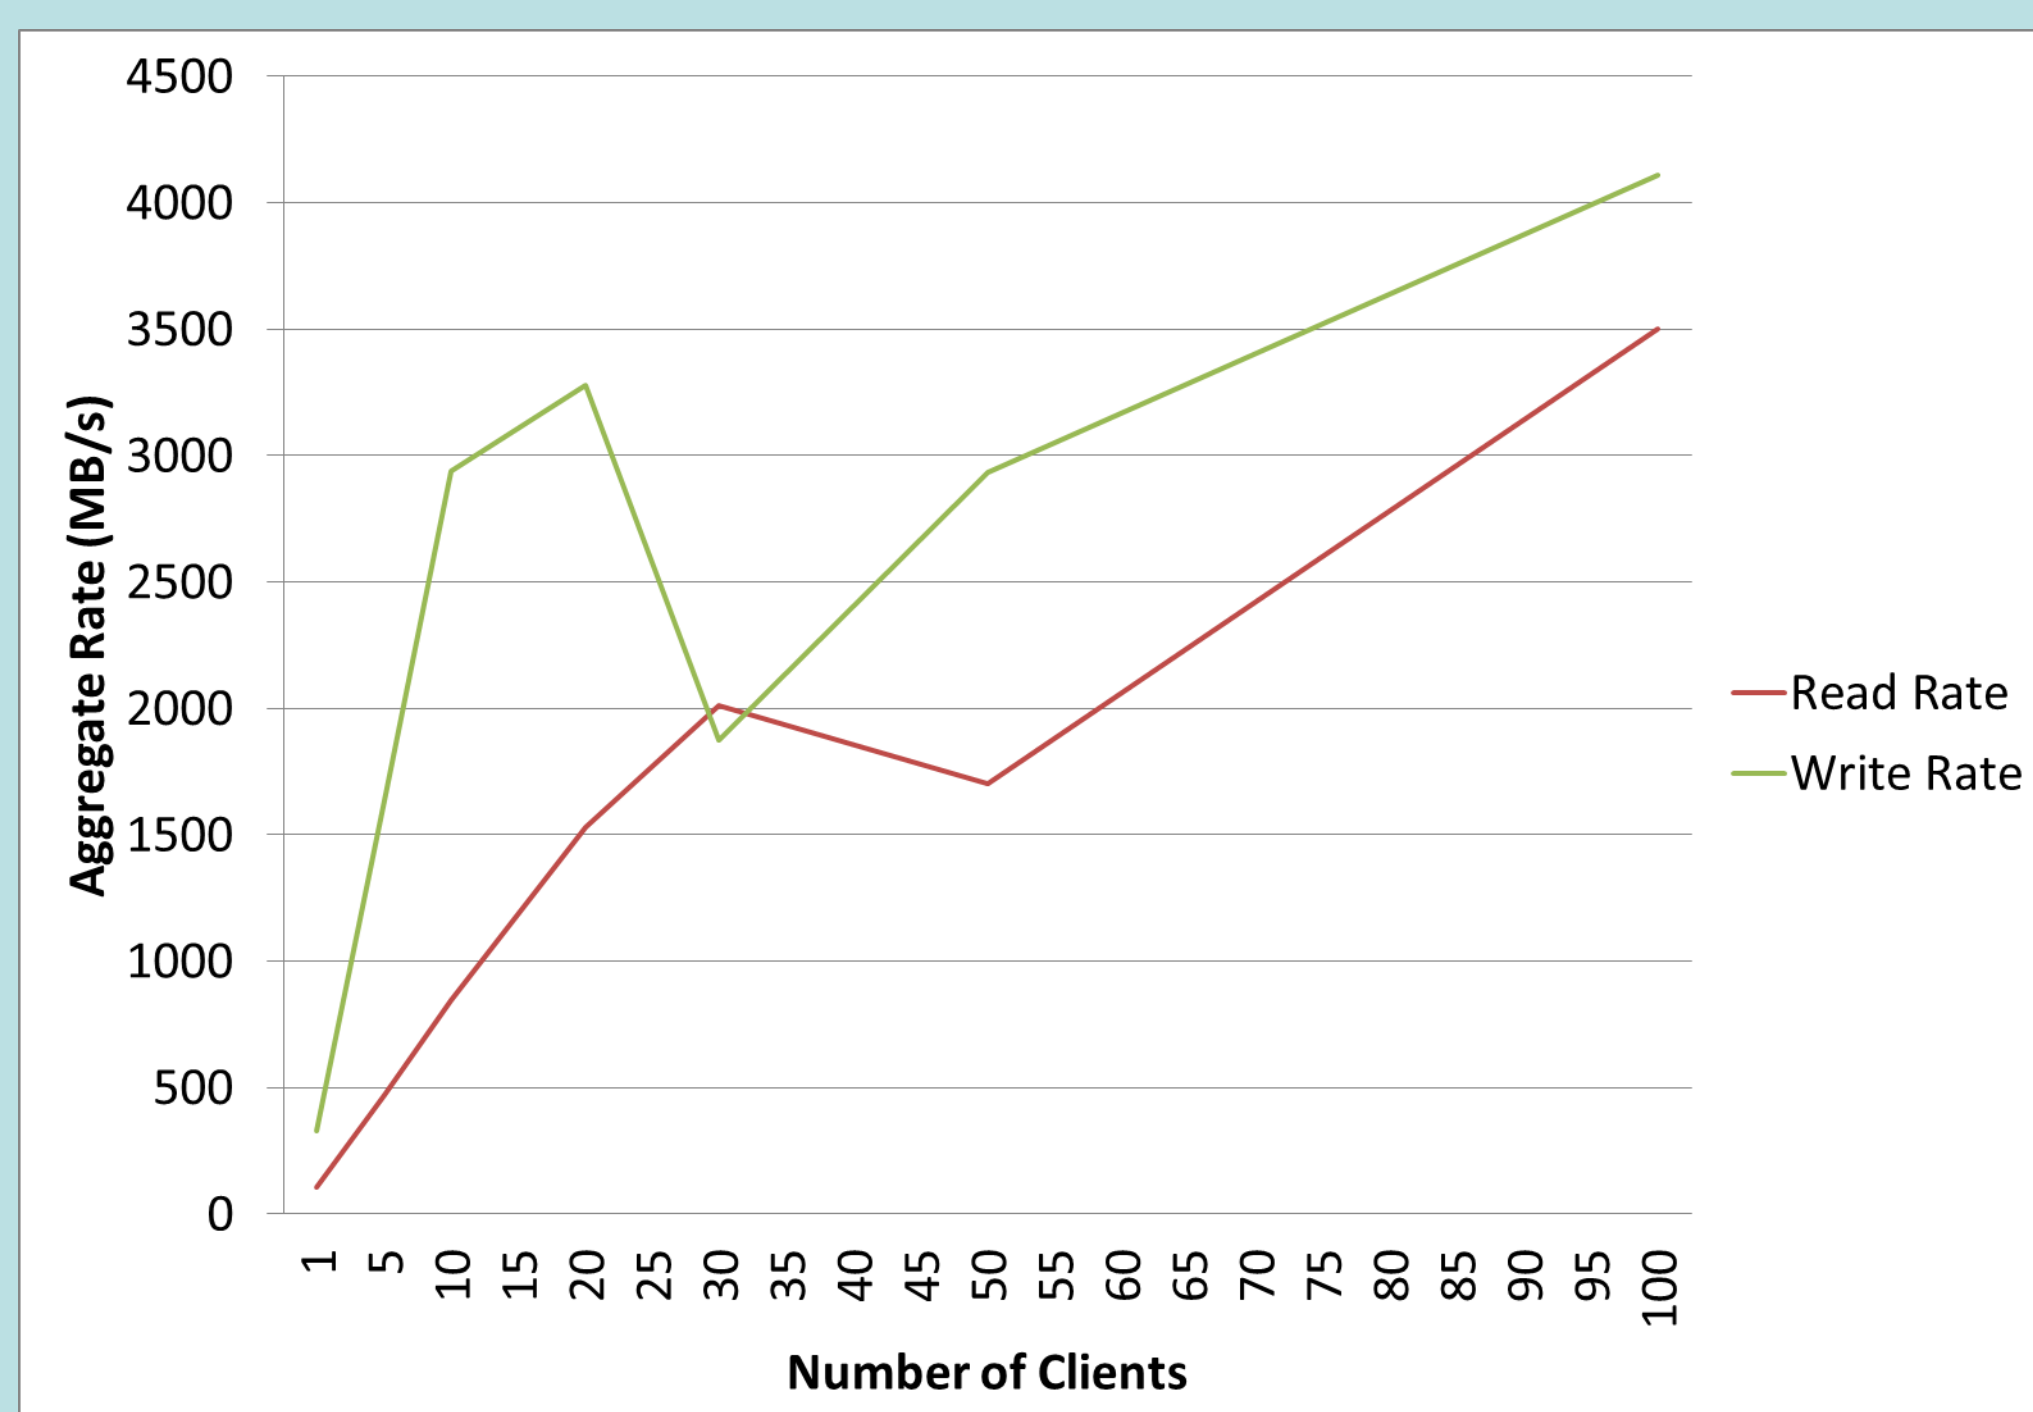
- Tests performed using test storage system, but production worker nodes
- Two tests of concurrent usage of XRootD (reading/writing)
- Write tests write unique 2GB files into storage
- Read tests select 1 of 100 unique 2GB files for reading
- Increase the number of worker nodes
- **All tests use the same redirector**

Performance Tests

- Uses CMS production infrastructure
 - Using load test files (all the same size)
- 100 consecutive reads, each of a random file
- Tests repeated using local, European and Global redirectors
- Single (dedicated) client running test
- Recall **transfer is local regardless of redirector**



Scalability Test Results

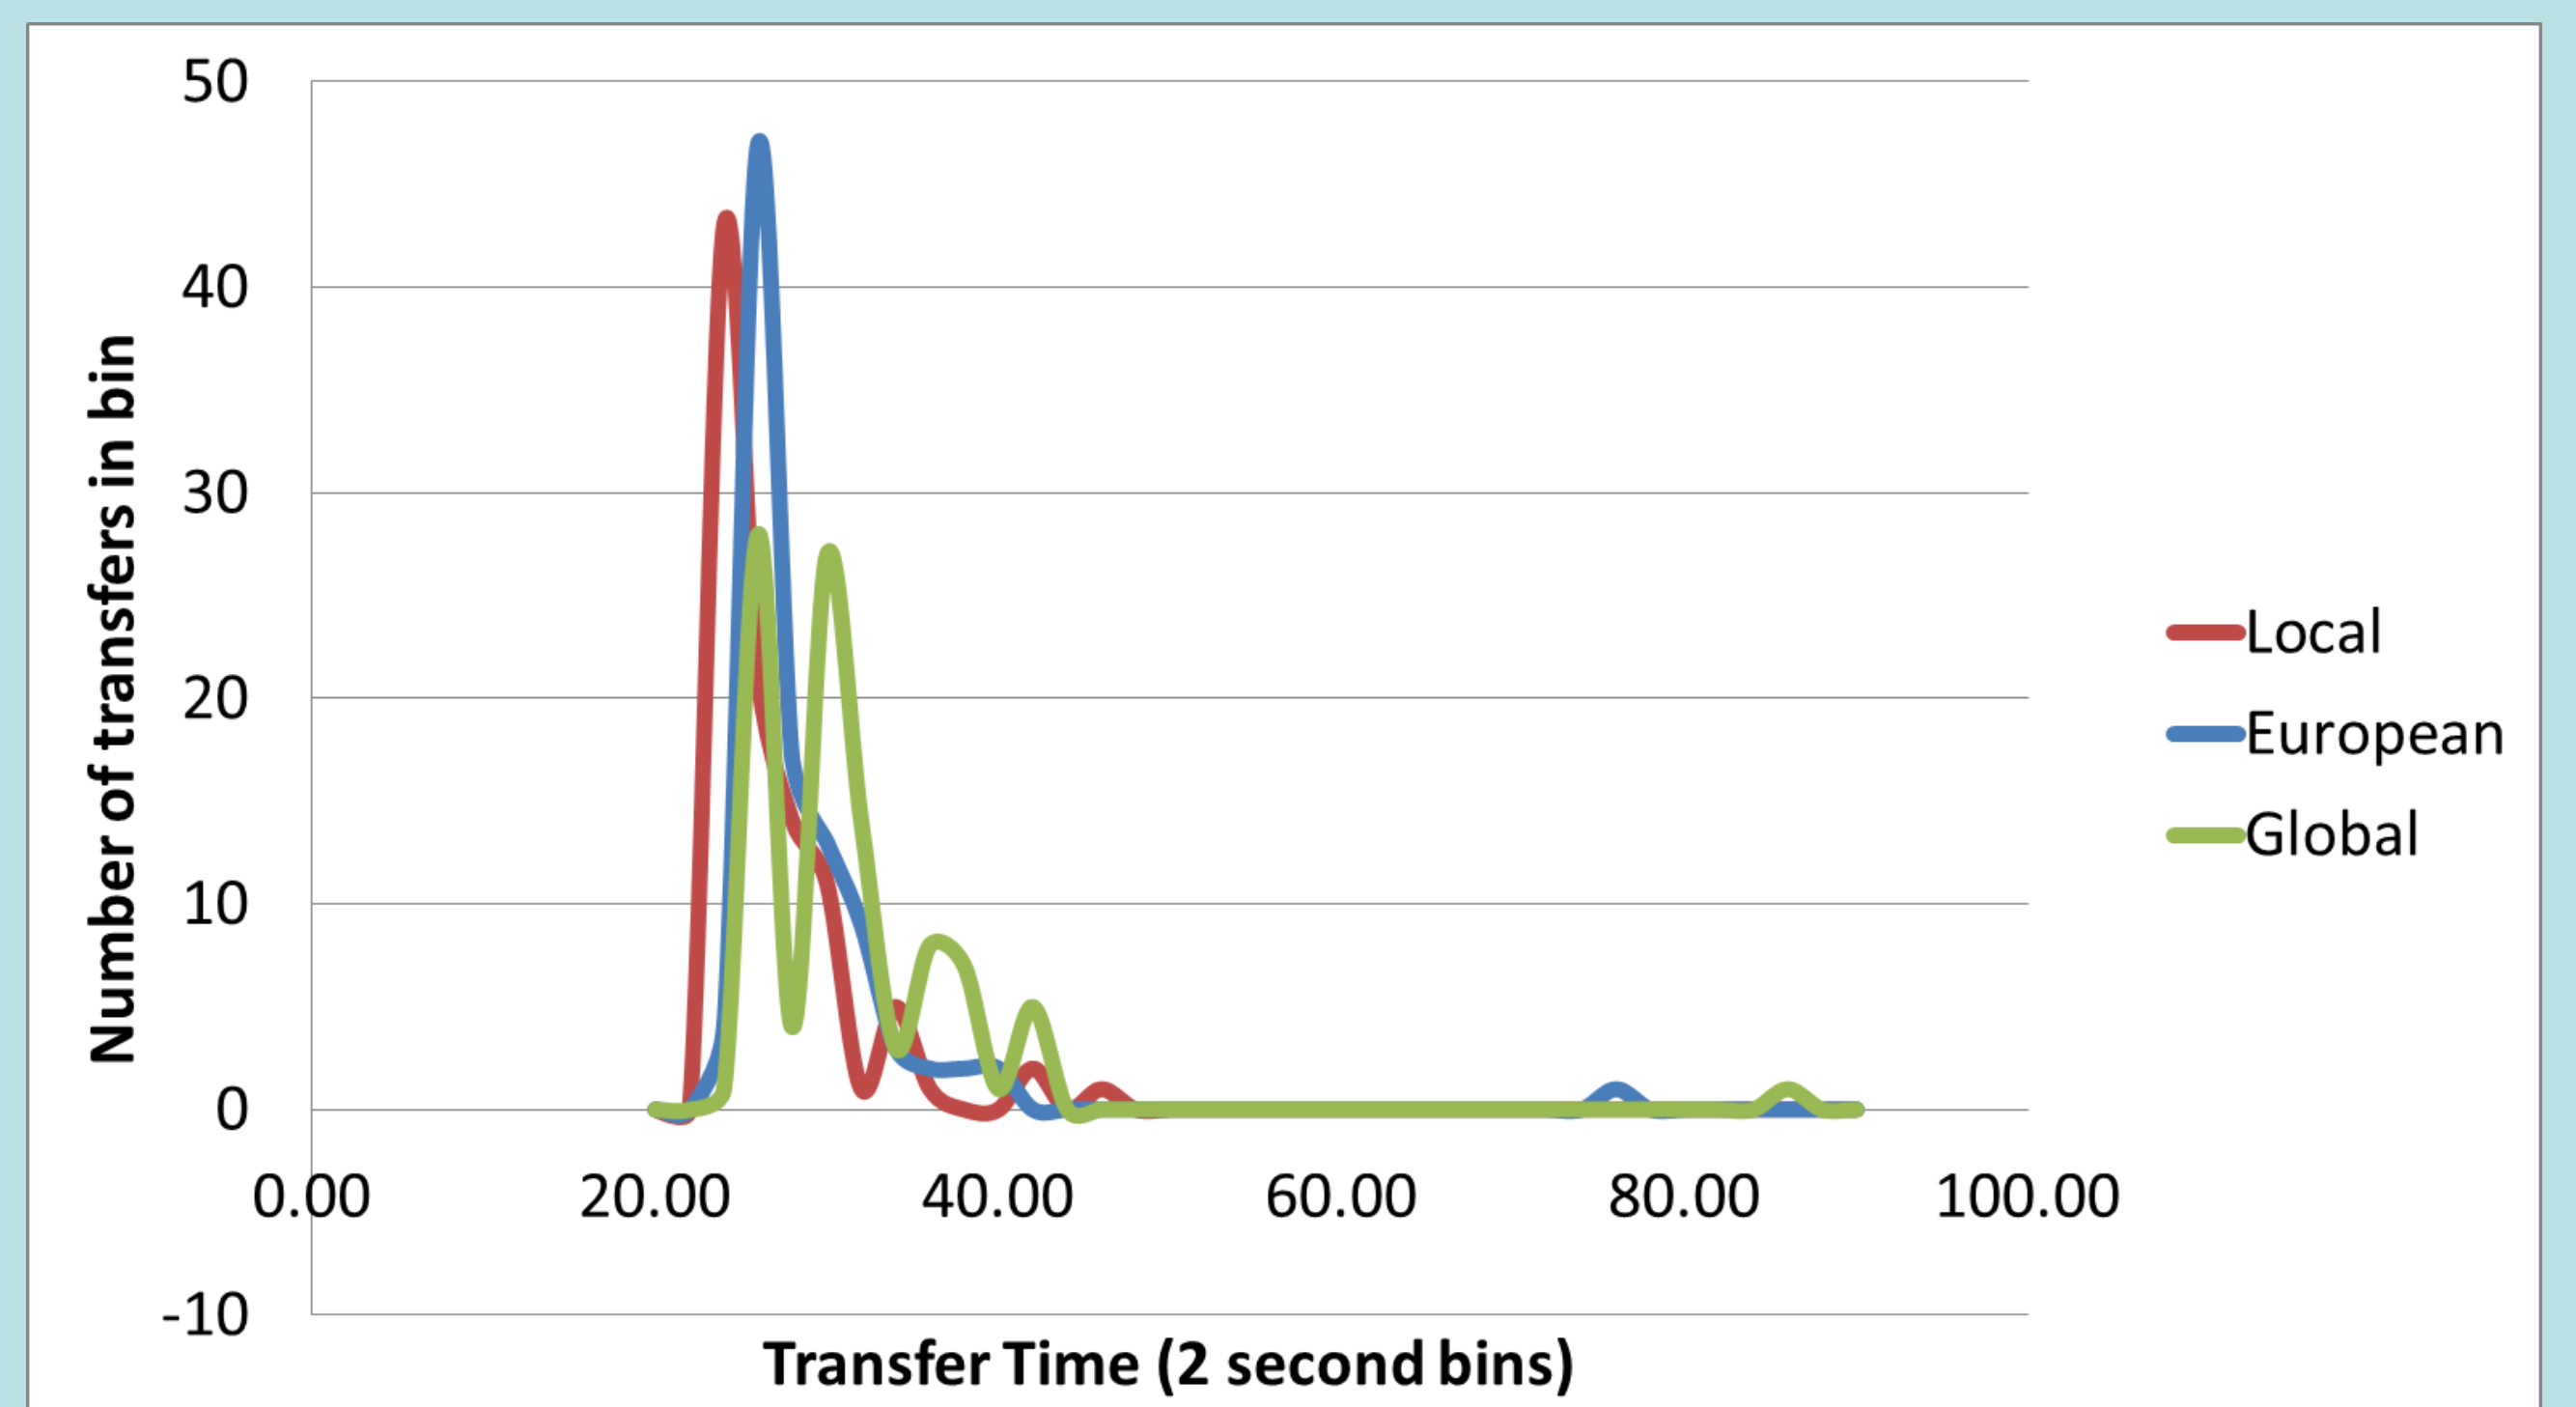


Summary

- Each test repeated 5 times
- **Performance scales ~linearly with number of clients**
- Rates increase at 30MB/s/client
- Write rates increase at 26 MB/s/client
- No significant increase in memory usage by the redirector observed

Note: Tests used production batch system and may have had to share bandwidth

Overhead Test Results



Summary

- Performance overhead **0.3-1.5 seconds/’hop’**
- Significantly better than SRM (1-10 seconds at STFC)
- **Within WLCG on-line data should never be more than 10 seconds away**

