

# Data Access – Performance in Remote I/O

Martin Hellmich

( on behalf of the LCGDM team )





- HTTP/DAV vs XROOTD vs RFIO
  - Soon adding NFS 4.1 / pNFS to the comparison

**LAN / Chunk Size: 10240-20480 / File Size: 2208309933**

Protocol	N. Reads	Read Size	Read Time
HTTP	500	22,773,112	0.43
HTTP	1000	46,027,143	0.87
XROOT	500	22,773,112	0.39
XROOT	1000	46,027,143	0.78
RFIO	500	22,773,112	136.35
RFIO	1000	46,027,143	274.89

**200x slower**



- HTTP/DAV vs XROOTD vs RFIO
  - Soon adding NFS 4.1 / pNFS to the comparison

**LAN / Chunk Size: 128-102400 / File Size: 2208309933 / 5000 Reads**

Protocol	Max. Vector	Read Size	Read Time
HTTP	8	1,166,613,844	12.55
HTTP	16	2,156,423,936	21.02
HTTP	24	3,211,861,800	29.97
HTTP	32	4,226,877,829	38.60
HTTP	64	8,535,839,293	75.20
XROOT	8	1,166,613,844	12.64
XROOT	16	2,156,423,936	22.11
XROOT	24	3,211,861,800	31.59
XROOT	32	4,226,877,829	41.14
XROOT	64	8,535,839,293	79.91



- HTTP/DAV vs XROOTD vs RFIO
  - Soon adding NFS 4.1 / pNFS to the comparison

**WAN / Chunk Size: 10240-20480 / File Size: 2208309933**

Protocol	N. Reads	Read Size	Read Time
HTTP	500	22,773,112	193.88
HTTP	1000	46,027,143	383.32
XROOT	500	22,773,112	143.37
XROOT	1000	46,027,143	283.63
RFIO	500	22,773,112	
RFIO	1000	46,027,143	





- We've also run a set of Hammercloud tests
  - Using ATLAS analysis jobs
  - These are only a few of all the metrics we have

	Remote HTTP	Remote HTTP (TTreeCache)	Staging HTTP	Remote XROOT	Remote XROOT (TTreeCache)	Staging XROOT	Staging GridFTP
Events Athena(s)	11.8	27.9	37.9	10.1	34.8	35.4	38.7
Event Rate(s)	11.5	26.8	24.2	9.9	33.1	17.7	24.2
Job Efficiency	1.0	0.994	1.0	1	0.99	0.999	1.0



- Big thanks to **ShuTing** and **ASGC**
  - For doing a lot of the testing and providing the infrastructure
- First recommendation is to phase down RFIO
  - No more development effort on it from our side
- HTTP vs XROOTD
  - Performance is equivalent, up to sites/users to decide
  - But we like standards... there's a lot to gain with them
- Staging vs Direct Access
  - Staging not ideal... requires lots of extra space on the WN
  - Direct Access is performant if used with ROOT TTreeCache