



GridPP

UK Computing for Particle Physics

GridPP File Transfers: dCache to DPM Woes

Graeme Stewart
University of Glasgow



GridPP PMB/DB, RAL January 2006



UNIVERSITY
of
GLASGOW



File Transfers

- Idea was to perform file transfers as part of SC4 preparations:
 - Tier 2 to/from Tier 1
 - Target rate 300Mb/s
 - Transfers of 1TB
 - Inter-Tier2
 - Target rate 100Mb/s
- Gain experience with tuning SRMs and network, uncover problems early.





- Immediately obvious that transfers from RAL dCache to Glasgow DPM was poor - 2Mb/s
- Investigated with help of FTS team. Pinned down to underlying gridftp transfer being throttled.
- Able to reproduce problems using `lcg-rep`, which confirmed this.





Transfer Times between SE Types

lcg-rep 1GB file, 10 GridFTP streams
Transfer time in seconds

Source	Destination			
	RAL dCache	CERN Castor	Gla Classic	Gla DPM
RAL dCache	0	167	241	3907
CERN Castor	87	0	177	212
Gla Classic	27	157	0	42
Gla DPM	29	157	136	0


- Notice the poor transfer time from dCache to DPM.
- Results confirmed as a general problem by testing with other dCaches (Edinburgh) and DPMs (Edinburgh, Durham).





Transfer Time Graph

lcg-rep 1GB file, 10 GridFTP streams





- Issue raised with dCache (GC) and DPM (GS, JKF) in December.
- Suggestions from DPM developers of workarounds - multistreams, kernel tweaks, etc. Tried unsuccessfully.
- DPM developers believe problem may disappear by implementing srmCopy (but I'm skeptical).





- Urgent issue to resolve for GridPP (13 DPMs installed/planned) and the whole of LCG
- It is being looked at, but...
- Do we need to apply pressure at a higher level?

See also:

<http://ppewww.ph.gla.ac.uk/~fergusjk/lcgRepTesting.txt>

