IF Data Handling Client Tools (ifdhc) - Feature #3605

Statistics collection

03/15/2013 05:08 PM - Marc Mengel

Status:	Closed	**Start date:**	03/15/2013
Priority:	Normal	**Due date:**	
Assignee:	Marc Mengel	**% Done:**	100%
Category:		**Estimated time:**	0.00 hour
Target version:	v1_2_5	**Spent time:**	6.00 hours

Description

Can we get better statisitics collection (i.e. i/o stats on cp/gridftp sub processes? data rates?)

History

#1 - 08/15/2013 11:42 AM - Marc Mengel

- Description updated

- Status changed from New to Assigned

- Target version set to v1_2_5

- % Done changed from 0 to 30

We have a rough implementation of this using getrusage(); but this only
works properly on SLF6. And it's still a bit of a lie, because it ignores
buffer-cached reads. But it's better that nothing.

what's there so far is in 0742ecba and 7c105bb0

#2 - 10/01/2013 11:40 AM - Marc Mengel

Updated to try to stat() all our arguments, and if we found files, to keep a sum of the st_size values for the source and destination files separately. If we don't get rusage blocks in and out, we take the st_size sums/512 and report that. That should give us some approximation of block count vs time that we can use
to collect througput numbers.

#3 - 10/03/2013 05:34 PM - Marc Mengel

Okay, now we stat our inputs and outputs **after** we copied the files, and total up any of them we can see, and use the max of inputs/outputs (because often one is remote and we get zero) to report bytes transferred, and we use higher resolution time to do our difference so we get a delta time for quick (i.e. local) copies. Downside is we include fork/excec time in the timing, but for larger transfers that should come out in the wash.

#4 - 10/03/2013 05:34 PM - Marc Mengel

- % Done changed from 30 to 100

#5 - 10/09/2013 10:40 AM - Marc Mengel

- Status changed from Assigned to Closed