

Pre-meeting on "busy" storage services (11/02/2009)

Participants

J.P. Baud, F. Donno, A. Frohner, E. Lanciotti, G. Lo Presti, R. Mollon, A. Sciabà, D. Smith, P. Tedesco

Conclusions

Suggesting that the client specifies *desiredTotalRequestTime* and the server should do its best to abort a request, if the time is reached and the client did not abort it explicitly (i.e. the client has crashed). Returning *remainingRequestTime* is desirable but less important.

Suggested a back-off algorithm on both synchronous and asynchronous methods: **if an SRM_INTERNAL_ERROR is received, the client shall use an exponential retry period on the same operation.**

Suggested an asynchronous polling method:

1. send an *srmPrepareToGet/Put*
2. in a loop with an exponential retry period until something changes: *srmGetRequestSummary* (lighter on the SRM)
3. use *srmGetStatus* to get the details of the change.

On the metascheduling level: it is very hard to give a meaningful estimate on how "busy" is the SE, so we give up on that for now.

-- AndreaSciaba - 12 Feb 2009

This topic: LCG > CCRC08SSWGStorageBusyMeeting090211

Topic revision: r1 - 2009-02-12 - AndreaSciaba



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback