

Kepler - Bug #5410

provenance data is not flushed to HSQL database before Kepler quit.

05/24/2011 11:09 AM - jianwu jianwu

Status:	Resolved	Start date:	05/24/2011
Priority:	Normal	Due date:	
Assignee:	Daniel Crawl	% Done:	0%
Category:	provenance	Estimated time:	0.00 hour
Target version:	reporting-2.3.0	Spent time:	0.00 hour
Bugzilla-Id:	5410		
Description			
Currently, provenance data writing to HSQL database has write delay (100 ms by default). It causes some data is not update or written because the data is not flushed before Kepler quit.			
To reproduce it,			
1) Run workflow in Kepler GUI or non-GUI mode.			
2) Quit Kepler if Kepler is run in GUI mode right after workflow is executed.			
3) Use HSQL client to connect hsql provenance database.			
4) Check 'workflow_exec' table: the record the execution might not be written in the database or the execution status might be always 'running'.			
It should work for both GUI mode or non-GUI mode. Not sure whether it works for other databases, such as Oracle and mySQL.			

History

#1 - 05/24/2011 11:30 AM - Derik Barseghian

I'm a little surprised you're seeing this with write_delay at 100ms. Have you tried 0 to see if it's truly the problem? I set write_delay to 100ms at r20931 (see bug#4325) to avoid a similar problem, and that was fast enough. Lowering it is probably fine, but keep an eye out for performance issues - e.g. does workflow execution speed slow down as it waits for writes to provenance?

#2 - 05/24/2011 11:35 AM - jianwu jianwu

I discussed this issue with Dan. He confirmed the bug. My test is done in non-GUI mode, so Kepler exit right after workflow execution is done.

I know set write delay to be 0 will slow down workflow execution, maybe quite a lot. So a better way is to keep some write delay and flush data before Kepler exit.

#3 - 06/16/2011 11:00 AM - Daniel Crawl

Fixed in r27741.

#4 - 03/27/2013 02:30 PM - Redmine Admin

Original Bugzilla ID was 5410