

Kepler - Bug #4325

Workflow Run Manager - deleted runs sometimes reappear after Kepler relaunch

08/20/2009 05:14 PM - Derik Barseghian

Status:	Resolved	Start date:	08/20/2009
Priority:	Normal	Due date:	
Assignee:	Derik Barseghian	% Done:	0%
Category:	workflow run manager	Estimated time:	0.00 hour
Target version:	wrp-modules-2.0.0	Spent time:	0.00 hour
Bugzilla-Id:	4325		
Description			
To reproduce: <ul style="list-style-type: none">- Launch Kepler- Run a workflow twice.- Export the 2nd run.- Run the workflow again.- Delete third run from WRM- Close, then relaunch Kepler. #3's back The problem is related to exporting a run.			

History

#1 - 08/21/2009 05:53 PM - Derik Barseghian

Also: If instead of deleting the third run, you delete the second run (the one you exported), and then close and relaunch kepler, you get a `FileNotFoundException` from `CacheManager's getCacheObjectIterator`. This is because the file is still in `cacheContentTable`. Watching the DB after a delete operation, the DB row never actually gets deleted.

#2 - 08/21/2009 06:08 PM - Derik Barseghian

If you:

- Launch Kepler
- Run a workflow twice.
- Export the 2nd run.
- Delete the 2nd run
- Close and relaunch, the 2nd run is gone and no errors.

#3 - 08/21/2009 06:23 PM - Derik Barseghian

Now I'm unable to reproduce the error from the procedure in comment [#2](#), even though I'd done it a few times. Maybe something more insidious happening here...

#4 - 09/02/2009 05:38 PM - Derik Barseghian

One way to sometimes get this to occur seems to be quitting kepler immediately after clicking ok in the delete runs dialog. You needn't have exported any of the runs. The preparedStatements will return how many rows they've "deleted" and the delete method always seems to complete before kepler-shutdown, yet when you connect to the database post-kepler shutdown (connecting to the db file instead of server) the run rows still exist in `workflow_exec`.

#5 - 10/07/2009 04:43 PM - Daniel Crawl

HSQL delays writing to the file system after updates occur to the database. According to the docs, http://hsqldb.org/doc/guide/ch09.html#set_write_delay-section, the default delay is 20 seconds, but my `hsqldb.script` says it's 10 seconds. We should probably decrease this.

#6 - 10/07/2009 05:05 PM - Derik Barseghian

Dan and I discussed this, it seems like a likely culprit. I'm going to try to change the write delay to see if it fixes this bug, and if/how bad it hurts performance. We figure a better solution is probably to do a clean shutdown of the server when Kepler quits, if it's the last Kepler instance running. To know if it's the last Kepler running, we might e.g. store and check, on shutdown, a `numberKeplersRunning` variable in the database.

#7 - 10/07/2009 06:57 PM - Derik Barseghian

Decreased `write_delay` to 100ms in `r20931`. Needs more testing to check for any performance hit, but so far so good. I can't get runs to reappear no matter how fast I quick after doing a delete (good deal). If this does impact performance, I suspect we still want to make `write_delay` much shorter than what it was, 10s, so that the db loses less in cases of kepler-crash.

#8 - 10/16/2009 11:22 AM - Derik Barseghian

This is fixed, but we need to test to see if the fix degraded performance (so far it seems fine) before closing.

Dan, do you have any particular workflows you use for testing provenance performance?

#9 - 11/30/2009 05:49 PM - Derik Barseghian

No reports of degraded performance, going to close this.

#10 - 03/27/2013 02:26 PM - Redmine Admin

Original Bugzilla ID was 4325