

Kepler - Bug #5790

Not all remote run-kars appear in Workflow Run Manager

01/17/2013 06:31 PM - Derik Barseghian

Status:	Resolved	Start date:	01/17/2013
Priority:	Normal	Due date:	
Assignee:	Derik Barseghian	% Done:	0%
Category:	workflow run manager	Estimated time:	0.00 hour
Target version:	reporting-2.4.0	Spent time:	0.00 hour
Bugzilla-Id:	5790		

Description

I believe the Workflow Run Manager is not displaying all the remote run-kars that it should. E.g. "04-HelloWorld-execlD6_20120830-204114819PDT,kar" (as seen in Component Search). My guess would be something's going wrong in the rather hideous ProvenanceStore.downloadAllWorkflowRuns method I wrote. This method avoids the very slow process of downloading all run-kars to get at their contained WorkflowRun files, and instead reassembles, in a very fragile way, WorkflowRuns from fragments returned by a query on the KarXMLs in metacat. At the time I didn't know of a way to query for/extract from results the entire WorkflowRun xmls in a cleanly delimited way, but iirc later Matt and/or Ben told me there is. Look into this...

History

#1 - 01/24/2013 04:30 PM - ben leinfelder

A way to extract full fragments of XML in a query result!? I would love to know how.

#2 - 01/25/2013 09:48 PM - ben leinfelder

Talked with Derik about this today and suggested issuing multiple queries, each targeting a specific part of the runkar moml @name/@value pairs using a predicate xpath condition in the query. You'd then aggregate these multiple search results to build the runkar preview. I think this is the best way to guarantee query consistency since we can't rely on the @name/@value pairs being in any particular order let alone corresponding to one another.

#3 - 01/28/2013 05:54 PM - Derik Barseghian

After some testing and further discussion w/ Jing and Ben, it turns out the condition section of the query is just for determining which docids to return, with returnfields being returned from those, so predicates there won't help my usecase. However you **can** use predicates in returnfields (something we didn't think we supported), which is great. I'm in the process of refactoring my query to do this.

#4 - 01/29/2013 02:41 PM - Derik Barseghian

fixed at r31387

#5 - 03/27/2013 02:31 PM - Redmine Admin

Original Bugzilla ID was 5790