

Metacat - Bug #6035

The system metadata table in a Metacat was messed up with the value of archived

07/10/2013 11:19 AM - Jing Tao

Status:	Rejected	Start date:	07/10/2013
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	2.1.0	Spent time:	0.00 hour
Bugzilla-Id:			
Description			
In the mn-demo-4.test.dataone.org, I queried the systemmetadata table and got: metacat=> select count(*) from systemmetadata where obsoleted_by is not null and archived=false; count ----- 4644 (1 row) You see, there are 4644 documents which are not archived but have the obsoleted_by value. This is a contradiction. Also, the xml_documents and xml_revisions table are messed up: metacat=> select count(*) from xml_documents; count ----- 4998 (1 row) metacat=> select count(*) from xml_revisions; count ----- 0 We need figure out why the metacat has this funny result.			

History

#1 - 07/10/2013 01:12 PM - ben leinfelder

- Target version set to 2.1.0

The archive flag need not be set to true when there is an obsoletedBy value present. Either the obsoletedBy OR the archived=true flag indicate that the object is not to be included in search results. It is fine if both are set, but that is not a requirement.

As for the xml_revisions table, if all objects have been added/updated via the DataONE API, then there are no true revisions in the old Metacat sense of it. We generate a new "autogen" id for each insert/update and there is no connection between one autogen id and the next (update).

This might pose an issue for Metacat deployments that still utilize the pathquery service fro searching but use the D1 API for inserting and updating since it will appear that all revisions of all objects are "current".

#2 - 07/19/2013 05:19 PM - Jing Tao

- Status changed from New to Rejected