

Infrastructure - Bug #3264

Production data problems after 1.0.4 release

2012-09-26 16:49 - Robert Waltz

Status:	Closed	Start date:	2012-09-26
Priority:	Normal	Due date:	
Assignee:	Chris Jones	% Done:	100%
Category:	Environment.Production	Estimated time:	0.00 hour
Target version:		Story Points:	
Milestone:	None		
Product Version:			

Description

There are documents in in xml_documents on UCSB that only show up in xml_revisions on ORC and UNM.

There are documents in in xml_documents on ORC and UNM that only show up in xml_revisions on UCSB.

see the attached file

History

#1 - 2012-09-26 16:49 - Robert Waltz

- Target version changed from Sprint-2012.37-Block.5.3 to Sprint-2012.39-Block.5.4

#2 - 2012-10-02 00:23 - Chris Jones

- Status changed from New to In Progress

- Target version deleted (Sprint-2012.39-Block.5.4)

- Milestone changed from CCI-1.1 to None

The symptoms described here are not the expected behavior across Metacat's that perform two-way replication. In fact, one would expect that the xml_documents tables and xml_revisions tables on each Coordinating Node would stay synchronized over time. However, this issue is minor in that, for CNs, all operations take place on a pid rather than a docid, and therefore inconsistencies between documents found in xml_documents and those found in xml_revisions are functionally not a problem whatsoever.

That said, we still need to investigate why this is happening. The only code in Metacat that transfers documents from xml_documents into xml_revisions are calls like update() and delete(). These methods fire off ForceReplicationHandler events to subsequently update remot Metacats. These replication events may be failing do to timeouts or some other issue, and I'm seeing many types of ERROR messages in /var/metacat/logs/replicate.log* having to do with calls in ForceReplicationHandler. Each class of error probably needs to be tracked down to understand why we're seeing inconsistencies across tables. Many may be temporary, and so getting a grasp on this will take a while.

Overall, the differences are limited to a few dozen docids. Given that, and the fact that there's no functional problem on the CN, I'm setting this to low priority for now and will put it on the back burner. It will be good to track down the cause though.

#3 - 2017-03-28 17:52 - Dave Vieglais

- Status changed from In Progress to Closed

- % Done changed from 0 to 100

Files

ErrorsOnProd20120926.txt

4.06 KB

2012-09-26

Robert Waltz