

## Metacat - Bug #4616

### Timed Replication takes many hours and drives the load up on KNB

12/09/2009 08:46 AM - Michael Daigle

<b>Status:</b>	Resolved	<b>Start date:</b>	12/09/2009
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Michael Daigle	<b>% Done:</b>	0%
<b>Category:</b>	metacat	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	1.9.2	<b>Spent time:</b>	0.00 hour
<b>Bugzilla-Id:</b>	4616		

#### Description

Timed replication can more than 12 hours lter and almost that long for PISCO.

During that time, the load on the machine can surpass 10 and java can use over 200% cpu.

lter has over 80K records and pisco over 60K. Only a handful of those actually need to be replicated. We need to see why it's taking so long and using so many resources.

#### History

##### #1 - 01/19/2010 10:32 AM - Michael Daigle

After raising the log level on KNB to INFO and tracing through the logs, it was clear that the delay was happening when converting the document StringReader back to a string so the file could be written to disk.

Note, the reason this was done in the first place was that the StringReader for the doc was closed during the xml parsing. The write to disk should only take place after the parse is successfully completed. So the doc was extracted to a string and then written to disk after parse.

Since the doc is currently read into a string when it is pulled from the request object, I changed the code to send that string to DocumentImpl.write and .writeReplication. In that way, a StringReader can be created for the parsing, and another for the write to disk.

##### #2 - 03/27/2013 02:27 PM - Redmine Admin

Original Bugzilla ID was 4616