

Infrastructure - Bug #2439

inter-cn replication not working on dev environment

2012-03-06 01:52 - Dave Vieglais

Status:	Closed	Start date:	2012-03-06
Priority:	Urgent	Due date:	
Assignee:	Ben Leinfelder	% Done:	100%
Category:	Metacat	Estimated time:	0.00 hour
Target version:	Sprint-2012.09-Block.2.1	Story Points:	
Milestone:	CCI-1.0.0		
Product Version:	*		

Description

Adding new content to a MN in dev environment results in synchronization picking up the new content, and system metadata for the new object appearing on the three dev CNs as expected.

Attempting to the retrieve the object only works from one CN however.

Example:

<https://cn-dev.dataone.org/cn/v1/object/dave.20120305.00>

OK

<https://cn-dev-2.dataone.org/cn/v1/object/dave.20120305.00>

FAIL

<https://cn-dev-3.dataone.org/cn/v1/object/dave.20120305.00>

FAIL

Not clear if this is a config problem with the Dev environment or something more insidious.

Logs don't show anything of interest that I could find.

Verified behavior with additional content created on both GMN and DEMO2

History

#1 - 2012-03-06 20:44 - Ben Leinfelder

Aside from being molasses slow, I was able to insert an EML document to cn-dev (using the Metacat API via the "dev skin") which was successfully replicated to the other two servers:

<https://cn-dev.dataone.org/cn/v1/object/ben.2.1>

<https://cn-dev-2.dataone.org/cn/v1/object/ben.2.1>

<https://cn-dev-3.dataone.org/cn/v1/object/ben.2.1>

This makes me think the MN->CN create via synchronization has a different insert routine than the standard Metacat insert that somehow circumvents Metacat replication.

Of course rerunning my initial test failed because cn-dev kept churning on the simple insert and I stopped the request.

#2 - 2012-03-07 18:09 - Ben Leinfelder

After the cn-dev* servers were rebuilt and more tests run, I can see that the science metadata for dave.20120307.00 is actually being replicated as a traditional Metacat document, but the ID mapping between PID and Metacat Doc Id is incorrect.

<https://cn-dev.dataone.org/cn/v1/object/dave.20120307.00>

on cn-dev:
PID: dave.20120307.00
Metacat docid: autogen.2012030708481184456.1
identifier table has mapping between PID and Docid

on cn-dev-2:

PID in systemMetadata table: dave.20120307.00

Metacat docid: autogen.2012030708481184456.1

no identifier mapping entry between the PID and Docid - it's missing.

There is an entry that maps PID "autogen.2012030708481184456.1" to Docid "autogen.2012030708481184456.1" - but there is no systemMetadata entry for "autogen.2012030708481184456.1" - and there shouldn't be.

I need to make sure to get the GUID->Docid mapping set correctly when replication runs. It probably means overwriting any existing mapping that was made during the initial insert of the replicated object (EML) before replication information (including the systemmetadata and id mapping info) comes across the replication channel.

replicated EML as seen via Metacat Doc Ids:

<https://cn-dev.dataone.org/knb/metacat/autogen.2012030708481184456.1/default>

<https://cn-dev-2.dataone.org/knb/metacat/autogen.2012030708481184456.1/default>

<https://cn-dev-3.dataone.org/knb/metacat/autogen.2012030708481184456.1/default>

#3 - 2012-03-07 22:35 - Ben Leinfelder

talked it through with Chris and we came to this conclusion:

IdentifierManager.identifierExists(pid) was checking the "identifier" table (where the PID->Docid mappings live) and the "systemmetadata" table (where the pid lives) for the existence of the identifier.

Because HZ immediately replicates the SM entry, the identifier was thought to exist at the target CN but it did not have the identifier mapping yet (comes with traditional replication of the underlying EML document) so the mapping would never be created correctly for the PID->docid.

Now in Metacat replication we only check if the mapping has been created before creating the ID mapping.

Chris will redeploy on cn-devs for more testing.

#4 - 2012-03-14 17:16 - Ben Leinfelder

- Status changed from New to Closed

seems to be have done the trick in the cn-dev* environment