

Infrastructure - Task #3357

Story # 3352 (Closed): Production CNs are out of sync in object count

Address ISet iterator bug that only iterates over a subset of the ISet

2012-10-23 17:27 - Chris Jones

Status:	New	Start date:	2012-10-23
Priority:	Normal	Due date:	
Assignee:	Ben Leinfelder	% Done:	0%
Category:	Metacat	Estimated time:	0.00 hour
Target version:	Sprint-2012.50-Block.6.4	Story Points:	
Milestone:	None		
Product Version:	*		
Description When bringing the CNs back into sync, Metacat's HazelcastService compares the local set of known pids with the shared set in the Hazelcast ISet structure (hzIdentifiers). However, a call to ISet.size() gives a higher number (328K) than the iterator provides (!208K). This is likely a bug in the Hazelcast implementation, and may be addressed in newer HZ versions, but for now we need a workaround. Perhaps doing a putAll() into another structure will get all of the listed pids, and then iterating over that structure locally?			

History

#1 - 2012-10-24 00:18 - Ben Leinfelder

Current plan of attack:
Investigate using ISet in place ISet with 300K+ records. If that works, we can convert Metacat and indexer to use this structure. If not...tbd.

#2 - 2012-10-24 15:53 - Ben Leinfelder

I have been unable to replicate this behavior while testing locally (350K ISet records). I did momentarily see inconsistencies in the count when I had my heap memory too small -- we could double check that other nodes are not throwing exceptions like that when we request the iterator from the production ISet.
I want to try a few more things (adding duplicate Identifier objects, perhaps?).

#3 - 2012-10-25 20:15 - Ben Leinfelder

With 350K test Identifiers on cn-dev, I still cannot replicate this issue -- .size() and counting the iterator items both give the expected number.

#4 - 2012-12-12 16:54 - Chris Jones

- Target version changed from Sprint-2012.41-Block.6.1 to Sprint-2012.50-Block.6.4