

Infrastructure - Bug #2519

get calls to hzSystemMetadata returns null for valid documents in database

2012-03-22 22:08 - Skye Roseboom

Status:	Closed	Start date:	2012-03-22
Priority:	Normal	Due date:	
Assignee:	Chris Jones	% Done:	100%
Category:	Metacat	Estimated time:	0.00 hour
Target version:	Sprint-2012.11-Block.2.2	Story Points:	
Milestone:	CCI-1.0.0		
Product Version:	*		

Description

There are some objects that appear in the cn-sandbox object list (cn/v1/object) but do not seem to exist in the hazelcast system metadata map.

For example esa.89.1, knb-lter-sbc.28.14, peggym.28.5, nikkis.34.18, knb-lter-sbc.27.8

These document ids were discovered in the hzIdentifiers hzcast set and then confirmed to appear in the object list.

Investigating esa.89.1 - it appears in the backing database. "SELECT * FROM systemmetadata where guid='esa.89.1' does return one row on cn-sandbox-unm-1. But I cannot get it to load into the hazelcast system metadata map or to display on the /cn/v1/meta or object endpoints. It does appear in the object list however.

Not sure if the error is that they appear in the object list and in hzIdentifiers OR if the error is that they do not appear in system metadata hz map.

The sandbox CNs were scrubbed on Mar 21, 2012 11:47:43 PM

esa.89.1 was synchronized on cn-sandbox-unm-1 at 2012-03-22 00:09:17,320 from urn:node:mnStageUCSB1.

Replication occurs and the document appears on urn:node:mnStageORC1.

Synchronization runs again and picks up the record on cn-sandbox-orc-1. However, instead of going through the normal update process as it would with a replica that it finds in the hzSystemMetadata map, it receives a null object and so determines that esa.89.1 should be created again. However, upon the call to CNCORE.create(), a IdentifierNotUnique error is returned.

As noted above, esa.89.1 is in the database in systemmetadata, in xml_documents and the sciMetadata can be found in /var/metacat/Documents as 'autogen.2012032116091821001.1'. The record was found on all the CNs. But failed to load on all the CNs.

hz_peek confirms that hzSystemMetadata does not contain the Identifier key for esa.89.1. However a call to D1NodeService.getSystemMetadata will throw an NotFound error in D1NodeService.isAuthorized at line 938 indicating that HazelcastService.getInstance().getSystemMetadataMap().get(pid) returns null.

History

#1 - 2012-03-22 22:54 - Robert Waltz

- Subject changed from Missing documents that appear in object list to get calls to hzSystemMetadata returns null for valid documents in database
- Category set to Metacat
- Assignee set to Chris Jones
- Target version set to Sprint-2012.11-Block.2.2

#2 - 2012-04-20 15:26 - Chris Jones

- Status changed from New to Closed

There were a few situations where metacat was returning null system metadata from the backing store, some due to access control code. Ben and I tracked down what we think are all of the issues and have made minor changes to the HazelcastService to deal with the issues. Also, when a given CN drops from the Hazelcast cluster and re-enters, it now evaluates the identifiers table and the identifiers found in the systemmetadata table. If there are missing system metadata records, they are re-populated from the hzSystemMetadata map so they are consistent across all CNs.