Infrastructure - Task #1784

Task # 1782 (Closed): CN Replication components should be separated for scalability

Implement CNReplication in Metacat's CNodeService using Hazelcast

2011-09-11 19:47 - Chris Jones

Status:	Rejected		**Start date:**	2011-09-11
Priority:	Normal		**Due date:**	
Assignee:	Chris Jones		**% Done:**	0%
Category:	d1_replication		**Estimated time:**	0.00 hour
Target version:				
Milestone:	CCI-0.6.4		**Story Points:**	
Product Version:	*			

Description

Metacat's edu.ucsb.nceas.metacat.dataone.CNodeService currently implements the CNReplication API for the local Metacat instance, but doesn't communicate across the CN Hazelcast cluster . Change CNodeService to use the hzSystemMetadata map (as a member) and hzNodes map (as a client) while responding to CNReplication requests. The local Metacat's system metadata table will be updated via put() calls to the map using the edu.ucsb.nceas.metacat.dataone.hazelcast.SystemMetadataMap map-store implementation. Updates to the remote CN Metacat's tables need to happen via a Hazelcast Executor service.

CNodeService currently uses edu.ucsb.nceas.metacat.replication.ForceReplicationSystemMetadataHandler to kick off system metadata replication handling for Metacat (in a serial fashion). This should be replaced|refactored|subclassed to use the Hazelcast ExecutorService to submit CNReplicationTask objects to each CN in parallel:

Set = Hazelcast.getCluster().getMembers();

Each CNReplicationTask should be submitted to all cluster members except the member that owns the key (pid) of the SystemMetadata object since that member updates it's system metadata table through the MapStore implementation. For dcience metadata, all members should be updated. to find out which member owns a key:

PartitionService partitionService = Hazelcast.getPartitionService();
Partition partition = partitionService.getPartition(key);
Member ownerMember = partition.getOwner();

Remove that owner from the Set of Members when submitting the distributed task.

See "Distributed Execution":http://www.hazelcast.com/documentation.jsp#ExecutorService

History

#1 - 2012-01-03 16:46 - Chris Jones

- Status changed from New to Rejected

This functionality occurs via a change listener in Metacat that monitors the system metadata map, and so the executor service isn't needed.