

Infrastructure - Task #1789

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

Integrate replication queues and maps into Hazelcast processing cluster

2011-09-11 21:12 - Chris Jones

<b>Status:</b>	Closed	<b>Start date:</b>	2011-09-11
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Robert Waltz	<b>% Done:</b>	100%
<b>Category:</b>	d1_replication	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>		<b>Story Points:</b>	
<b>Milestone:</b>	CCI-0.6.4		
<b>Product Version:</b>	*		

**Description**

The ReplicationService class was previously only connecting to a single Hazelcast cluster. It is now defined as ReplicationManager, and will connect to the storage cluster initiated by Metacat as a HazelcastClient and to the processing cluster initiated by the batch daemon as a HazelcastInstance cluster member. The queues and maps defined in "this configuration file": [https://repository.dataone.org/software/cicore/trunk/cn/d1\\_replication/src/main/resources/hazelcast.xml](https://repository.dataone.org/software/cicore/trunk/cn/d1_replication/src/main/resources/hazelcast.xml) should be integrated into the default configuration file for the processing cluster.

Specifically, add the hzReplicationTasks queue, the hzReplicationTasksMap map (that backs the queue), and the hzPendingReplicationTasks map. The hzSystemMetadata map will be added to the Metacat storage cluster, so it doesn't need to be transferred.

History

#1 - 2011-09-11 21:12 - Chris Jones

- Subject changed from Integrate replication queus and maps into Hazelcast processing cluster to Integrate replication queues and maps into Hazelcast processing cluster

#2 - 2011-09-28 19:22 - Robert Waltz

- Status changed from New to Closed