

Infrastructure - Bug #8641

Story # 8639 (New): Replication performance is too slow to service demand

Any change to SystemMetadata causes a new replication task to be generated without consideration of the object replicability

2018-07-04 12:28 - Dave Vieglaiss

<b>Status:</b>	New	<b>Start date:</b>	2018-07-04
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Chris Jones	<b>% Done:</b>	0%
<b>Category:</b>	d1_replication	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>		<b>Story Points:</b>	
<b>Milestone:</b>	None		
<b>Product Version:</b>	*		
<b>Description</b>			
The hazelcast event listener implemented by ReplicationEventListener basically does:			
<pre>ReplicationEventListener.entryUpdated()     if isAuthoritativeReplicaValid()         createReplicationTask()</pre>			
isAuthoritativeReplicaValid() checks whether the replication status for the Authoritative MN is complete.			
Hence, any update or add event on the systemmetadata map in Hazelcast will trigger addition of a replication task if the authoritative MN has a completed replica, even if replication is not allowed for the object. This causes a significant number of entries to be added to the replication task queue even though those tasks will never do anything as they will be later rejected.			
It would be appropriate in entryUpdated() to also check whether replication of the object is allowed. The overhead would be minimal since a copy of the system metadata is already available in entryUpdated(). The same logic should also be added to entryAdded().			
ReplicationManager implements boolean isAllowed(SystemMetadata sysmeta) which should do the job.			