

Infrastructure - Task #2209

Story # 2093 (Closed): Implement Log Aggregation

Design recovery mechanism for revived CN

2012-01-11 21:13 - Robert Waltz

Status:	Closed	Start date:	2012-01-11
Priority:	Normal	Due date:	
Assignee:	Robert Waltz	% Done:	100%
Category:	d1_log_aggregation	Estimated time:	0.00 hour
Target version:	Sprint-2012.05-Block.1.3	Story Points:	
Milestone:	CCI-1.0.0		
Product Version:	*		
Description			
<p>If a CN joins a cluster, it will need to build &/or rebuild its index.</p> <p>I think i should look at the last identifier that the CN has indexed. And then start up and start receiving entries. It should not immediately processing MNs for log entries. It should only listen until the first entries received after startup. After the first new entries are received then we know the gap that needs to be recovered.</p> <p>To assist ordering, the unique identifier should be a Double composed of the system's datetime as a long + . + Hazelcast.getAtomicNumber.incrementAndGet()</p> <p>This should guarantee uniqueness, and a certain degree of sequentiality (depending upon system clock skew it can not be guaranteed). because sequentiality may not be absolute, any recovery must start at a sequences before the start of the gap and end after the gap. Thus, the unique identifier field should in solr terms overwritable.</p> <p>Of course clock skew between the CN's can not be too great anyhow or else other facets of CN operation will be affected.</p>			

History

#1 - 2012-01-31 18:26 - Robert Waltz

- Status changed from New to Closed