

Infrastructure - Story #2622

d1_replication should prioritize MN replication tasks based on load, failures, and bandwidth factors

2012-04-20 23:23 - Chris Jones

Status:	Closed	Start date:	2012-04-20
Priority:	Normal	Due date:	
Assignee:	Chris Jones	% Done:	100%
Category:	d1_replication	Estimated time:	0.00 hour
Target version:	Sprint-2012.19-Block.3.2		
Story Points:			
Description			
<p>The current implementation of ReplicationManager evaluates ReplicationPolicies for objects and only prioritizes target member nodes based on the preferred and blocked lists. Otherwise, all MNs have been treated as equals. We've seen performance problems on single MNs that tend to cause performance problems on the CNs when sending MNRReplicationTasks to the ExecutorService. Task execution looks to slow down significantly when threads in the thread pool are held up by non-performant MNs.</p> <p>To alleviate this, we need to more intelligently evaluate the capabilities of an MN as a target, and prioritize targets that are performant.</p> <p>The strategy is to 1) Make MN implementations more resilient (i.e. queue replicate() requests), and 2) throttle requests to MNs based on a few different performance metrics. These are outlined at:</p> <p>http://epad.dataone.org/20120420-replication-priority-queue</p> <p>At first we will only throttle based on a limit of pending replication requests, but down the road will also evaluate the failure factor and the bandwidth factor.</p>			
Subtasks:			
Task # 2623: Modify Metacat's MNodeService.replicate() to queue requests			Closed
Task # 2624: Modify ReplicationManager.createAndQueueTasks() to limit replication tasks...			Closed

History

#1 - 2012-04-21 00:22 - Matthew Jones

- Tracker changed from Task to Story

#2 - 2012-04-23 17:55 - Dave Vieglais

- Position set to 1

- Target version changed from Sprint-2012.15-Block.2.4 to Sprint-2012.17-Block.3.1

#3 - 2012-05-07 16:58 - Dave Vieglais

- Position deleted (6)

- Target version changed from Sprint-2012.17-Block.3.1 to Sprint-2012.19-Block.3.2

- Position set to 2

#4 - 2012-05-21 16:22 - Chris Jones

- Status changed from New to Closed

This is implemented except for the bandwidth factors. We'll need to decide how each MN is scored based on bandwidth (calculated or assigned).