

## Infrastructure - Story #8447

### synchronization queue equity and monitoring

2018-03-01 16:46 - Rob Nahf

<b>Status:</b>	Closed	<b>Start date:</b>	2018-03-01
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Rob Nahf	<b>% Done:</b>	100%
<b>Category:</b>	d1_synchronization	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	CCI-2.3.9		
<b>Story Points:</b>			
<b>Description</b>			
<p>The recent initial sync from a new, large member node (with 409k items) brought out some equity issues with synchronization. Namely, that the queue is so long that new sync tasks can be 2 days back in line.</p> <p>We configured the queue to be much shorter (about an hour's worth of tasks), but maintaining an equitable sync priority might be challenging for nodes that synchronize less frequently than the one with a large number of sync tasks.</p> <p>Ideally, a multi-queue system would be adopted, with tasks "fanning-in" from each queue. This could be complicated development however.</p> <p>Other quicker improvements include:</p> <ol style="list-style-type: none"><li>1. limiting the max harvest to a percent of remaining capacity, instead of a fixed number. This might help get around the sync-frequency problem.</li><li>2. adding a metric that identifies the leading position of a member node's items.</li><li>3. limiting a node's harvest by sync frequency - we'd have to be able to convert the Node synchronization schedule into a frequency.</li></ol>			
<b>Related issues:</b>			
Related to Infrastructure - Bug #8468: synchronization requeueing for tempora...		<b>Closed</b>	<b>2018-03-02</b>

#### Associated revisions

##### Revision 19161 - 2018-03-15 17:10 - Rob Nahf

refs: #8447, #8466. Addressed unit test failures. Completed code changes, and ready for integration tests.

##### Revision 19161 - 2018-03-15 17:10 - Rob Nahf

refs: #8447, #8466. Addressed unit test failures. Completed code changes, and ready for integration tests.

#### History

##### #1 - 2018-03-01 17:27 - Rob Nahf

Hazelcast doesn't support dynamic queue creation, so a Map isn't possible. We could create a bunch of anonymous sync-queues, and create a Map that maps a Node to one of the anonymous ones. This introduces a scalability issue - when the number of Nodes > number of queues.

A simpler idea is to create a syncPriorityQueue and a syncBulkQueue for harvests under and over a certain size that's configurable.

##### #2 - 2018-03-02 00:30 - Rob Nahf

- Related to Bug #8468: synchronization requeueing for temporary unavailability of nodeComms causes massive delays for package added

##### #3 - 2018-03-02 21:48 - Dave Vieglais

- Target version changed from CCI-2.3.8 to CCI-2.3.9

##### #4 - 2018-03-16 04:24 - Rob Nahf

have sync running in sandbox, and demonstrated that a second harvest starts getting processed right away.

From the process-metric log, below, you can see that each MN is being drained from equally (the totals are decreasing by the same amount)

```
{"event":"synchronization queued","message":"Total Sync Objects Queued: 3765","threadName":"SynchronizationQuartzScheduler_Worker-2","threadId":61,"dateLogged":"2018-03-16T04:19:51.066+00:00"}
{"event":"synchronization queued","nodeId":"legacy","message":"Sync Objects Queued: 0","threadName":"SynchronizationQuartzScheduler_Worker-2","threadId":61,"dateLogged":"2018-03-16T04:19:51.066+00:00"}
{"event":"synchronization queued","nodeId":"urn:node:mnTestNCEI","message":"Sync Objects Queued: 3282","threadName":"SynchronizationQuartzScheduler_Worker-2","threadId":61,"dateLogged":"2018-03-16T04:19:51.066+00:00"}
{"event":"synchronization queued","nodeId":"urn:node:mnSandboxUCSB1","message":"Sync Objects Queued: 482","threadName":"SynchronizationQuartzScheduler_Worker-2","threadId":61,"dateLogged":"2018-03-16T04:19:51.066+00:00"}
{"event":"synchronization queued","message":"Total Sync Objects Queued: 3682","threadName":"SynchronizationQuartzScheduler_Worker-3","threadId":62,"dateLogged":"2018-03-16T04:20:51.071+00:00"}
{"event":"synchronization queued","nodeId":"legacy","message":"Sync Objects Queued: 0","threadName":"SynchronizationQuartzScheduler_Worker-3","threadId":62,"dateLogged":"2018-03-16T04:20:51.071+00:00"}
{"event":"synchronization queued","nodeId":"urn:node:mnTestNCEI","message":"Sync Objects Queued: 3240","threadName":"SynchronizationQuartzScheduler_Worker-3","threadId":62,"dateLogged":"2018-03-16T04:20:51.071+00:00"}
{"event":"synchronization queued","nodeId":"urn:node:mnSandboxUCSB1","message":"Sync Objects Queued: 441","threadName":"SynchronizationQuartzScheduler_Worker-3","threadId":62,"dateLogged":"2018-03-16T04:20:51.071+00:00"}
```

#### #5 - 2018-03-22 16:41 - Rob Nahf

- % Done changed from 0 to 100

- Status changed from New to Closed

tested and deployed on March 21.