

Infrastructure - Story #8172

investigate atomic updates for some solr updates

2017-09-01 19:35 - Rob Nahf

<b>Status:</b>	In Progress	<b>Start date:</b>	2017-09-01
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Rob Nahf	<b>% Done:</b>	30%
<b>Category:</b>	d1_indexer	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	CCI-2.4.0		
<b>Story Points:</b>			
<b>Description</b>			
Atomic updates came to solr with v4.0. (We're currently at 5.x)			
Atomic updates are supposed to be more efficient, and could help us with the race condition in <a href="#">#7771</a> . (multiple tasks reading a solr record and then modifying it in divergent ways via overwriting existing values.			
atomic add and remove modifiers allow addition and removal of multivalued fields, which is where our race conditions arise.			
<b>Related issues:</b>			
Related to CN Index - Task #7771: Use multiple threads to index objects		Closed	2016-05-04
Related to Infrastructure - Story #8702: Indexing Refactor Strategy		New	2018-09-24

History

#1 - 2017-09-01 19:36 - Rob Nahf

- Related to Task #7771: Use multiple threads to index objects added

#2 - 2017-10-24 16:46 - Dave Vieglaiss

- Target version changed from CCI-2.3.5 to CCI-2.3.6

#3 - 2017-10-30 19:05 - Jing Tao

- Target version changed from CCI-2.3.6 to CCI-2.3.7

#4 - 2017-12-15 19:45 - Rob Nahf

- Target version changed from CCI-2.3.7 to CCI-2.4.0
- % Done changed from 0 to 30
- Status changed from New to In Progress

Work about halfway done in v2.4-SNAPSHOT (trunk). That's to say that the implementation is working but a bit poorly designed, due to the existing structure of the code base (the iterative nature of sub-processor parsing and merging in SolrIndexService.processTask method ) The work is extensive enough that it should be in a new branch.

#5 - 2018-01-17 18:39 - Dave Vieglaiss

- Sprint set to Infrastructure backlog

#6 - 2018-09-24 15:28 - Rob Nahf

- Related to Story #8702: Indexing Refactor Strategy added