

## Infrastructure - Story #8175

### Solr crashing with out of memory error

2017-09-05 16:33 - Dave Vieglais

<b>Status:</b> Closed	<b>Start date:</b> 2017-09-05
<b>Priority:</b> Immediate	<b>Due date:</b>
<b>Assignee:</b> Dave Vieglais	<b>% Done:</b> 100%
<b>Category:</b> d1_indexer	<b>Estimated time:</b> 0.00 hour
<b>Target version:</b> CCI-2.3.7	
<b>Story Points:</b>	
<b>Description</b>	
Solr has been crashing and/or thrashing with out of memory errors on CN-UCSB-1 resulting in poor or no response when using search.dataone.org.	
There are several related issues to be addressed here:	
<ol style="list-style-type: none"><li>1. Poorly written queries generated by the search UI are causing excessive demand on the solr index</li><li>2. Solr is running with 2GB allocated to the JVM. This appears to be insufficient.</li><li>3. Should design and implement a load balancing solution so a single instance of Solr is not used all the time</li><li>4. Memory use is always high on cn-ucsb-1 where processing (and so log aggregation) is running. Review the log aggregation process for memory leaks and performance.</li><li>5. The configuration for Solr is using mostly default parameters. This should be reviewed and revised where appropriate.</li></ol>	
<b>Subtasks:</b>	
Task # 8176: Adjust memory allocated to Solr	<b>Closed</b>
Task # 8177: Ensure that inefficient queries are refactored in the search UI	<b>Closed</b>
Task # 8178: Design and implement load balancing for the search index	<b>Rejected</b>
Task # 8179: review log aggregation code for memory leaks and efficiency	<b>Closed</b>
Task # 8180: Review solr configuration, schema for efficiency	<b>Rejected</b>
<b>Related issues:</b>	
Related to Infrastructure - Story #7939: Indexing is too slow, especially wit...	<b>Rejected</b> <b>2016-11-25</b>

### History

#### #1 - 2017-09-05 16:33 - Dave Vieglais

- Related to Story #7939: Indexing is too slow, especially with large packages added

#### #2 - 2017-10-30 19:05 - Jing Tao

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

#### #3 - 2017-12-26 18:10 - Dave Vieglais

- Sprint set to Infrastructure backlog

#### #4 - 2017-12-26 18:18 - Dave Vieglais

- Sprint changed from Infrastructure backlog to CCI-2.3.7

#### #5 - 2017-12-26 18:21 - Dave Vieglais

- % Done changed from 0 to 30

- Status changed from New to In Progress

PBI status auto updated to In Progress because at least one task wasn't on New status

#### #6 - 2019-01-22 19:05 - Dave Vieglais

- % Done changed from 30 to 100

- Status changed from In Progress to Closed