

Infrastructure - Bug #7570

Search UI requests cause Out of Memory Error

2016-01-05 15:02 - Chris Jones

Status:	In Progress	Start date:	2016-01-05
Priority:	Normal	Due date:	
Assignee:	Jing Tao	% Done:	30%
Category:	d1_indexer	Estimated time:	0.00 hour
Target version:	CCI-2.4.0	Story Points:	
Milestone:	None		
Product Version:			

Description

When deploying the 1.9.0 Search UI, we saw out of memory errors in Solr on the CN. The requests executed by the UI in sequential order are:

https://search.dataone.org/cn/v2/query/solr/?fl=id,seriesId,title,origin,pubDate,dateUploaded,abstract,resourceMap,beginDate,endDate,read_count,i.geohash_9,datasource,isPublic,northBoundCoord,southBoundCoord,eastBoundCoord,westBoundCoord&q++++-obsoletedBy:*+formatType:METADATA++&sort=dateUploaded+desc&rows=25&start=0&facet=true&facet.sort=index&facet.field=geohash_1&facet.field=geohash_2&facet.field=geohash_3&facet.field=geohash_4&facet.field=geohash_5&facet.field=geohash_6&facet.field=geohash_7&facet.field=geohash_8&facet.field=geohash_9&facet.mincount=1&facet.limit=-1&wt=json

[https://search.dataone.org/cn/v2/query/solr/?q++++-obsoletedBy:*+formatType:METADATA+++beginDate:\[1800-01-01T00:00:00Z TO NOW\]&rows=1&fl=beginDate&sort=beginDate+asc&wt=json](https://search.dataone.org/cn/v2/query/solr/?q++++-obsoletedBy:*+formatType:METADATA+++beginDate:[1800-01-01T00:00:00Z TO NOW]&rows=1&fl=beginDate&sort=beginDate+asc&wt=json)

[https://search.dataone.org/cn/v2/query/solr/?q++++-obsoletedBy:*+formatType:METADATA+++endDate:\[1800-01-01T00:00:00Z TO NOW\]+-obsoletedBy:*&rows=1&fl=endDate&sort=endDate+desc&wt=json](https://search.dataone.org/cn/v2/query/solr/?q++++-obsoletedBy:*+formatType:METADATA+++endDate:[1800-01-01T00:00:00Z TO NOW]+-obsoletedBy:*&rows=1&fl=endDate&sort=endDate+desc&wt=json)

https://search.dataone.org/cn/v2/query/solr/?q++++-obsoletedBy:*+formatType:METADATA++&rows=0&facet=true&facet.sort=count&facet.mincount=1&facet.limit=-1&facet.field=keywords&facet.field=origin&facet.field=family&facet.field=species&facet.field=genus&facet.field=kingdom&facet.field=phylum&facet.field=order&facet.field=class&facet.field=site&facet.field=attributeName&facet.field=attributeLabel&wt=json&json.wrf=jQuery19109610712484859317_1452005239758&_=1452005239759

Here's the stacktrace reported back to the Search UI:

```
"error":{"msg":"Error while processing facet fields: java.lang.OutOfMemoryError: Java heap space","trace":"org.apache.solr.common.SolrException: Error while processing facet fields: java.lang.OutOfMemoryError: Java heap space
at org.apache.solr.request.SimpleFacets.getFacetFieldCounts(SimpleFacets.java:659)
at org.apache.solr.request.SimpleFacets.getFacetCounts(SimpleFacets.java:285)
at org.apache.solr.handler.component.FacetComponent.process(FacetComponent.java:102)
at org.apache.solr.handler.component.SearchHandler.handleRequestBody(SearchHandler.java:255)
at org.dataone.solr.handler.component.SolrSearchHandler.handleRequestBody(SolrSearchHandler.java:101)
at org.apache.solr.handler.RequestHandlerBase.handleRequest(RequestHandlerBase.java:143)
at org.apache.solr.core.SolrCore.execute(SolrCore.java:2064)
at org.apache.solr.servlet.HttpSolrCall.execute(HttpSolrCall.java:654)
at org.apache.solr.servlet.HttpSolrCall.call(HttpSolrCall.java:450)
at org.apache.solr.servlet.SolrDispatchFilter.doFilter(SolrDispatchFilter.java:227)
at org.apache.solr.servlet.SolrDispatchFilter.doFilter(SolrDispatchFilter.java:196)
at org.eclipse.jetty.servlet.ServletHandler$CachedChain.doFilter(ServletHandler.java:1652)
at org.dataone.solr.servlet.SessionAuthorizationUtil.handleNoCertificateManagerSession(SessionAuthorizationUtil.java:131)
at
org.dataone.solr.servlet.SearchServiceSessionAuthorizationFilter.handleNoCertificateManagerSession(SearchServiceSessionAuthorizationFilter.java:45)
at org.dataone.solr.servlet.SessionAuthorizationFilterStrategy.doFilter(SessionAuthorizationFilterStrategy.java:264)
at org.eclipse.jetty.servlet.ServletHandler$CachedChain.doFilter(ServletHandler.java:1652)
at org.eclipse.jetty.servlet.ServletHandler.doHandle(ServletHandler.java:585)
at org.eclipse.jetty.server.handler.ScopedHandler.handle(ScopedHandler.java:143)
at org.eclipse.jetty.security.SecurityHandler.handle(SecurityHandler.java:577)
at org.eclipse.jetty.server.session.SessionHandler.doHandle(SessionHandler.java:223)
at org.eclipse.jetty.server.handler.ContextHandler.doHandle(ContextHandler.java:1127)
```

```
at org.eclipse.jetty.servlet.ServletHandler.doScope(ServletHandler.java:515)
at org.eclipse.jetty.server.session.SessionHandler.doScope(SessionHandler.java:185)
at org.eclipse.jetty.server.handler.ContextHandler.doScope(ContextHandler.java:1061)
at org.eclipse.jetty.server.handler.ScopedHandler.handle(ScopedHandler.java:141)
at org.eclipse.jetty.server.handler.ContextHandlerCollection.handle(ContextHandlerCollection.java:215)
at org.eclipse.jetty.server.handler.HandlerCollection.handle(HandlerCollection.java:110)
at org.eclipse.jetty.server.handler.HandlerWrapper.handle(HandlerWrapper.java:97)
at org.eclipse.jetty.server.Server.handle(Server.java:497)
at org.eclipse.jetty.server.HttpChannel.handle(HttpChannel.java:310)
at org.eclipse.jetty.server.HttpConnection.onFillable(HttpConnection.java:257)
at org.eclipse.jetty.io.AbstractConnection$2.run(AbstractConnection.java:540)
at org.eclipse.jetty.util.thread.QueuedThreadPool.runJob(QueuedThreadPool.java:635)
at org.eclipse.jetty.util.thread.QueuedThreadPool$3.run(QueuedThreadPool.java:555)
at java.lang.Thread.run(Thread.java:745)
Caused by: java.lang.OutOfMemoryError: Java heap space
{"code":500}
```

To temporarily alleviate the issue, I restarted Solr on the 3 CNs. That had no effect, so I modified the Solr configuration in /var/solr/solr.in.sh to request 2048m of Java heap space (up from 512m that was in the file), and restarted Solr. That, too, had no effect. I then bumped it to 4096m, restarted, and all queries started to work again.

Related issues:

Related to Infrastructure - Bug #8027: Heap memory errors from solr	Closed	2017-02-27
---	---------------	-------------------

History

#1 - 2016-01-05 15:05 - Chris Jones

- Description updated

#2 - 2016-02-25 14:14 - Dave Vieglais

- Assignee changed from Skye Roseboom to Andrei Buium

reassigning.

#3 - 2016-04-01 15:42 - Andrei Buium

- Status changed from New to In Progress

- % Done changed from 0 to 30

Lauren has split the UI query most likely to be the culprit (the last one containing many facets) into multiple queries.

Other options to consider:

Using docValues="true" in the schema for fields often sorted/faceted on (would speed up query time and lower memory use at query time but increase total index size due to a forward map being created).

Checking if store="true" is absolutely necessary on all the current schema fields (NOTE: displaying formats without an existing xslt relies on solr stored fields).

Delivering a cached result (preferably also created from multiple queries).

Using SolrCloud to shard the index and split memory use more evenly across different nodes.

Further increasing memory allowed. (May be necessary regardless. Is it currently possible for a user to kill solr with a complicated query?)

#4 - 2017-03-28 16:15 - Dave Vieglais

- Project changed from CN Index to Infrastructure

- Tracker changed from Task to Bug

- Target version set to CCI-2.4.0

- Milestone set to None

- Category set to d1_indexer

- Assignee changed from Andrei Buium to Jing Tao

#5 - 2017-03-28 16:23 - Dave Vieglais

- Related to Bug #8027: Heap memory errors from solr added

#6 - 2017-04-26 21:01 - Rob Nahf

fwiw, I ran across this while investigating solr heap errors reported in a different bug ([#8027](#))

<http://solr.pl/en/2010/08/11/6-deadly-sins-in-the-context-of-query/>