

## Infrastructure - Bug #8696

### double indexing of a resource map and another not processed because of resource contention (lock) on member

2018-09-12 00:18 - Rob Nahf

<b>Status:</b>	New	<b>Start date:</b>	2018-09-12
<b>Priority:</b>	Urgent	<b>Due date:</b>	
<b>Assignee:</b>	Rob Nahf	<b>% Done:</b>	0%
<b>Category:</b>	d1_indexer	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	CCI-2.3.10	<b>Story Points:</b>	
<b>Milestone:</b>	None		
<b>Product Version:</b>	*		

#### Description

In production, the ORE 'a1a0e96a-3cde-4f3c-829c-29650b09f22b' was not processed because a member was also referenced by the ORE it obsoleted, 'dc39515e-440b-4673-9f63-962c7374bf48'. The task failed without being requeued. Below is the log output.

```
rnahf@cn-orc-1:/var/log/dataone/index$ grep ala0e96a-3cde-4f3c-829c-29650b09f22b cn-index-processor-daemon.log.*
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:27,384 (IndexTaskProcessor:saveTask:865)
IndexTaskProcess.saveTask save the index task ala0e96a-3cde-4f3c-829c-29650b09f22b
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:27,384 (IndexTaskProcessor:getNextIndexTask:610) Start of indexing pid: ala0e96a-3cde-4f3c-829c-29650b09f22b
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:34,832 (IndexTaskProcessor:getNextIndexTask:664) the original index task - IndexTask [id=18085996, pid=ala0e96a-3cde-4f3c-829c-29650b09f22b, formatid=http://www.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.2018091015425874434.1, dateSysMetaModified=1536087134490, deleted=false, taskModifiedDate=1536619467383, priority=3, status=IN PROCESS]
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:34,832 (IndexTaskProcessor:getNextIndexTask:671) the new index task - IndexTask [id=18085996, pid=ala0e96a-3cde-4f3c-829c-29650b09f22b, formatid=http://www.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.2018091015425874434.1, dateSysMetaModified=1536087134490, deleted=false, taskModifiedDate=1536619467383, priority=3, status=IN PROCESS]
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:34,901 (IndexTaskProcessor:checkReadinessProcessResourceMap:369) #####Another resource map is process the referenced id ee73cf7f-1005-4b89-bab9-3a7fa01d27c6 as well. So the thread to process id ala0e96a-3cde-4f3c-829c-29650b09f22b has to wait 0.5 seconds.
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:35,402 (IndexTaskProcessor:checkReadinessProcessResourceMap:369) #####Another resource map is process the referenced id ee73cf7f-1005-4b89-bab9-3a7fa01d27c6 as well. So the thread to process id ala0e96a-3cde-4f3c-829c-29650b09f22b has to wait 0.5 seconds.
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:35,902 (IndexTaskProcessor:checkReadinessProcessResourceMap:369) #####Another resource map is process the referenced id ee73cf7f-1005-4b89-bab9-3a7fa01d27c6 as well. So the thread to process id ala0e96a-3cde-4f3c-829c-29650b09f22b has to wait 0.5 seconds.
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:36,402 (IndexTaskProcessor:checkReadinessProcessResourceMap:369) #####Another resource map is process the referenced id ee73cf7f-1005-4b89-bab9-3a7fa01d27c6 as well. So the thread to process id ala0e96a-3cde-4f3c-829c-29650b09f22b has to wait 0.5 seconds.
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:36,903 (IndexTaskProcessor:checkReadinessProcessResourceMap:369) #####Another resource map is process the referenced id ee73cf7f-1005-4b89-bab9-3a7fa01d27c6 as well. So the thread to process id ala0e96a-3cde-4f3c-829c-29650b09f22b has to wait 0.5 seconds.
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:37,403 (IndexTaskProcessor:checkReadinessProcessResourceMap:369) #####Another resource map is process the referenced id ee73cf7f-1005-4b89-bab9-3a7fa01d27c6 as well. So the thread to process id ala0e96a-3cde-4f3c-829c-29650b09f22b has to wait 0.5 seconds.
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:37,903 (IndexTaskProcessor:checkReadinessProcessResourceMap:369) #####Another resource map is process the referenced id ee73cf7f-1005-4b89-bab9-3a7fa01d27c6 as well. So the thread to process id ala0e96a-3cde-4f3c-829c-29650b09f22b has to wait 0.5 seconds.
```

```
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:38,403 (IndexTaskProcessor:checkReadiness
ProcessResourceMap:369) #####Another resource map is process the referenced id ee73c
f7f-1005-4b89-bab9-3a7fa01d27c6 as well. So the thread to process id ala0e96a-3cde-4f3c-829c-29650
b09f22b has to wait 0.5 seconds.
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:38,904 (IndexTaskProcessor:checkReadiness
ProcessResourceMap:369) #####Another resource map is process the referenced id ee73c
f7f-1005-4b89-bab9-3a7fa01d27c6 as well. So the thread to process id ala0e96a-3cde-4f3c-829c-29650
b09f22b has to wait 0.5 seconds.
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:39,404 (IndexTaskProcessor:checkReadiness
ProcessResourceMap:369) #####Another resource map is process the referenced id ee73c
f7f-1005-4b89-bab9-3a7fa01d27c6 as well. So the thread to process id ala0e96a-3cde-4f3c-829c-29650
b09f22b has to wait 0.5 seconds.
cn-index-processor-daemon.log.6:[ERROR] 2018-09-10 22:44:39,904 (IndexTaskProcessor:checkReadiness
ProcessResourceMap:384) We waited for another thread to finish indexing a resource map which has t
he referenced id ee73cf7f-1005-4b89-bab9-3a7fa01d27c6 for a while. Now we quited and can't index i
d ala0e96a-3cde-4f3c-829c-29650b09f22b
cn-index-processor-daemon.log.6:[ERROR] 2018-09-10 22:44:39,904 (IndexTaskProcessor:processTask:29
7) Unable to process task for pid: ala0e96a-3cde-4f3c-829c-29650b09f22b
cn-index-processor-daemon.log.6:java.lang.Exception: We waited for another thread to finish indexi
ng a resource map which has the referenced id ee73cf7f-1005-4b89-bab9-3a7fa01d27c6 for a while. No
w we quited and can't index id ala0e96a-3cde-4f3c-829c-29650b09f22b
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:39,906 (IndexTaskProcessor:newOrFailedInd
exTaskExists:890) IndexTaskProcess.newOrFailedIndexTaskExists for id ala0e96a-3cde-4f3c-829c-29650
b09f22b
rnahf@cn-orc-1:/var/log/dataone/index$ date
Tue Sep 11 23:46:56 UTC 2018
rnahf@cn-orc-1:/var/log/dataone/index$ grep dc39515e-440b-4673-9f63-962c7374bf48 cn-index-processor-
daemon.log.*
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:12,133 (HZEEventFilter:filter:127) HZEEvent
Filter.filter - the system metadata for the index event shows shows dc39515e-440b-4673-9f63-962c73
74bf48 having a newer version than the SOLR server. So this event should be granted for indexing.
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:13,347 (HZEEventFilter:filter:127) HZEEvent
Filter.filter - the system metadata for the index event shows shows dc39515e-440b-4673-9f63-962c73
74bf48 having a newer version than the SOLR server. So this event should be granted for indexing.
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:18,677 (IndexTaskProcessor:saveTask:865)
IndexTaskProcess.saveTask save the index task dc39515e-440b-4673-9f63-962c7374bf48
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:18,677 (IndexTaskProcessor:getNextIndexTa
sk:610) Start of indexing pid: dc39515e-440b-4673-9f63-962c7374bf48
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:25,783 (IndexTaskProcessor:getNextIndexTa
sk:664) the original index task - IndexTask [id=18086020, pid=dc39515e-440b-4673-9f63-962c7374bf48
, formatid=http://www.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.20170725141
44216470.1, dateSysMetaModified=1536087137440, deleted=false, taskModifiedDate=1536619458675, prio
rity=3, status=IN PROCESS]
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:25,783 (IndexTaskProcessor:getNextIndexTa
sk:671) the new index task - IndexTask [id=18086020, pid=dc39515e-440b-4673-9f63-962c7374bf48, for
matid=http://www.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.2017072514144216
470.1, dateSysMetaModified=1536087137440, deleted=false, taskModifiedDate=1536619458675, priority=
3, status=IN PROCESS]
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:27,221 (IndexTaskProcessor:processTask:28
4) *****start to process update index task for dc39515e-440b-4673-9f63-962c7374bf4
8 in thread 20
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:43,513 (IndexTaskProcessor:processTask:28
8) *****end to process update index task for dc39515e-440b-4673-9f63-962c7374bf48
in thread 20
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:44:43,519 (IndexTaskProcessor:processTask:31
5) Indexing complete for pid: dc39515e-440b-4673-9f63-962c7374bf48
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:46:11,604 (IndexTaskProcessor:saveTask:865)
IndexTaskProcess.saveTask save the index task dc39515e-440b-4673-9f63-962c7374bf48
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:46:11,604 (IndexTaskProcessor:getNextIndexTa
sk:610) Start of indexing pid: dc39515e-440b-4673-9f63-962c7374bf48
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:46:18,731 (IndexTaskProcessor:getNextIndexTa
sk:664) the original index task - IndexTask [id=18086015, pid=dc39515e-440b-4673-9f63-962c7374bf48
, formatid=http://www.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.20170725141
44216470.1, dateSysMetaModified=1536087137440, deleted=false, taskModifiedDate=1536619571603, prio
rity=3, status=IN PROCESS]
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:46:18,732 (IndexTaskProcessor:getNextIndexTa
sk:671) the new index task - IndexTask [id=18086015, pid=dc39515e-440b-4673-9f63-962c7374bf48, for
```

```
matid=http://www.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.2017072514144216
470.1, dateSysMetaModified=1536087137440, deleted=false, taskModifiedDate=1536619571603, priority=
3, status=IN PROCESS]
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:46:20,164 (IndexTaskProcessor:processTask:28
4) *****start to process update index task for dc39515e-440b-4673-9f63-962c7374bf4
8 in thread 20
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:46:36,252 (IndexTaskProcessor:processTask:28
8) *****end to process update index task for dc39515e-440b-4673-9f63-962c7374bf48
in thread 20
cn-index-processor-daemon.log.6:[ INFO] 2018-09-10 22:46:36,255 (IndexTaskProcessor:processTask:31
5) Indexing complete for pid: dc39515e-440b-4673-9f63-962c7374bf48
cn-index-processor-daemon.log.7:[ INFO] 2018-09-10 21:44:09,798 (HZEventFilter:compareReplicaList:
256) HZEventFilter.compareReplicaList - the system metadata for the index event shows dc39515e-440
b-4673-9f63-962c7374bf48 having the same replica list as the solr doc.
cn-index-processor-daemon.log.7:[ INFO] 2018-09-10 21:44:09,798 (HZEventFilter:filter:164) HZEvent
Filter.filter - the system metadata for the index event shows dc39515e-440b-4673-9f63-962c7374bf48
having the same modification date as the SOLR server. Also both have the same replica list. So th
is event has been filtered out for indexing (no indexing).
rnahf@cn-orc-1:/var/log/dataone/index$
```

**Related issues:**

Related to Infrastructure - Story #8702: Indexing Refactor Strategy New 2018-09-24

**History**

**#1 - 2018-09-12 00:20 - Rob Nahf**

- Description updated

**#2 - 2018-09-12 00:22 - Rob Nahf**

- Description updated

**#3 - 2018-09-21 22:51 - Rob Nahf**

- Related to Story #8702: Indexing Refactor Strategy added