

Infrastructure - Bug #8536

resource Map update when metadata SID is used is not indexed

2018-04-09 22:27 - Rob Nahf

Status:	In Progress	Start date:	2018-04-09
Priority:	Normal	Due date:	
Assignee:	Rob Nahf	% Done:	30%
Category:	d1_indexer	Estimated time:	0.00 hour
Target version:		Story Points:	
Milestone:	None		
Product Version:	*		

Description

the checkReadinessProcessResourceMap routine consistently keeps a resourceMap update from being indexed.

```
[ERROR] 2018-04-09 21:57:24,438 (IndexTaskProcessor:checkReadinessProcessResourceMap:353) We waited for another thread to finish indexing a pid with series id packageIndexingTest_P_S_P_T_T_T_20180400.14.25_md for a while. Now we quit and can't index id packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.2
```

```
[ERROR] 2018-04-09 21:57:24,438 (IndexTaskProcessor:processTask:297) Unable to process task for pid: packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.2
```

full log related to the resourceMap.1 and .2 shown below. Perhaps it's a race condition.

```
[ INFO] 2018-04-09 21:55:15,098 (HZEventFilter:filter:127) HZEventFilter.filter - the system metadata for the index event shows packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1 having a newer version than the SOLR server. So this event should be granted for indexing.
```

```
[ INFO] 2018-04-09 21:55:16,442 (IndexTaskProcessor:saveTask:865) IndexTaskProcess.saveTask save the index task packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1
```

```
[ INFO] 2018-04-09 21:55:16,442 (IndexTaskProcessor:getNextIndexTask:610) Start of indexing pid: packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1
```

```
[ INFO] 2018-04-09 21:55:16,467 (IndexTaskProcessor:getNextIndexTask:664) the original index task - IndexTask [id=8067008, pid=packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1, formatid=http://www.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.2018040914304929417.1, dateSysMetaModified=1523310751860, deleted=false, taskModifiedDate=1523310916438, priority=3, status=IN PROCESS]
```

```
[ INFO] 2018-04-09 21:55:16,467 (IndexTaskProcessor:getNextIndexTask:671) the new index task - IndexTask [id=8067008, pid=packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1, formatid=http://www.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.2018040914304929417.1, dateSysMetaModified=1523310751860, deleted=false, taskModifiedDate=1523310916438, priority=3, status=IN PROCESS]
```

```
[ INFO] 2018-04-09 21:55:21,278 (IndexTaskProcessor:processTask:284) *****start to process update index task for packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1 in thread 109
```

```
[ INFO] 2018-04-09 21:55:21,533 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/solr/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=1&indent=on
```

```
[ INFO] 2018-04-09 21:55:21,535 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/solr/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=1&indent=on
```

```
[ INFO] 2018-04-09 21:55:21,538 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/solr/search_core/select/?q=resourceMap%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=5000&indent=on
```

```
[ INFO] 2018-04-09 21:55:21,598 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/solr/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=1&indent=on
```

```
[ INFO] 2018-04-09 21:55:21,601 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
```

```
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:55:22,446 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1+OR+seri
esId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_data.1+OR+seriesId%3ApackageIndexingTest_P_S
_P_T_T_T_20180400.14.25_md&start=0&rows=3&indent=on
[ INFO] 2018-04-09 21:55:22,451 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1+OR+id%3Apacka
geIndexingTest_P_S_P_T_T_T_20180400.14.25_data.1+OR+id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.
14.25_md&start=0&rows=3&indent=on
[ INFO] 2018-04-09 21:55:25,522 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0
&rows=1&indent=on
[ INFO] 2018-04-09 21:55:25,524 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:55:25,561 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0
&rows=1&indent=on
[ INFO] 2018-04-09 21:55:25,564 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:55:25,567 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0
&rows=1&indent=on
[ INFO] 2018-04-09 21:55:25,575 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:55:25,578 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0
&rows=1&indent=on
[ INFO] 2018-04-09 21:55:25,581 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:55:26,096 (IndexTaskProcessor:processTask:288) *****end to p
rocess update index task for packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1 in thread 109
[ INFO] 2018-04-09 21:55:26,099 (IndexTaskProcessor:processTask:315) Indexing complete for pid: pa
ckageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1
[ INFO] 2018-04-09 21:57:12,851 (IndexTaskProcessor:saveTask:865) IndexTaskProcess.saveTask save t
he index task packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1
[ INFO] 2018-04-09 21:57:12,851 (IndexTaskProcessor:getNextIndexTask:610) Start of indexing pid: p
ackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1
[ INFO] 2018-04-09 21:57:12,867 (IndexTaskProcessor:getNextIndexTask:664) the original index task
- IndexTask [id=8067004, pid=packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1, formatid=htt
p://www.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.2018040914304929417.1, da
teSysMetaModified=1523310751860, deleted=false, taskModifiedDate=1523311032850, priority=3, status
=IN PROCESS]
[ INFO] 2018-04-09 21:57:12,867 (IndexTaskProcessor:getNextIndexTask:671) the new index task - Ind
exTask [id=8067004, pid=packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1, formatid=http://w
ww.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.2018040914304929417.1, dateSys
MetaModified=1523310751860, deleted=false, taskModifiedDate=1523311032850, priority=3, status=IN P
ROCESS]
[ INFO] 2018-04-09 21:57:15,910 (IndexTaskProcessor:saveTask:865) IndexTaskProcess.saveTask save t
he index task packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.2
[ INFO] 2018-04-09 21:57:15,910 (IndexTaskProcessor:getNextIndexTask:610) Start of indexing pid: p
ackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.2
[ INFO] 2018-04-09 21:57:15,956 (IndexTaskProcessor:getNextIndexTask:664) the original index task
- IndexTask [id=8067042, pid=packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.2, formatid=htt
p://www.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.2018040914550540142.1, da
teSysMetaModified=1523310752163, deleted=false, taskModifiedDate=1523311035906, priority=3, status
=IN PROCESS]
[ INFO] 2018-04-09 21:57:15,956 (IndexTaskProcessor:getNextIndexTask:671) the new index task - Ind
exTask [id=8067042, pid=packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.2, formatid=http://w
ww.openarchives.org/ore/terms, objectPath=/var/metacat/data/autogen.2018040914550540142.1, dateSys
MetaModified=1523310752163, deleted=false, taskModifiedDate=1523311035906, priority=3, status=IN P
ROCESS]
[ INFO] 2018-04-09 21:57:16,478 (IndexTaskProcessor:processTask:284) *****start to
```

```
process update index task for packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1 in thread 1
10
[ INFO] 2018-04-09 21:57:16,807 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0
&rows=1&indent=on
[ INFO] 2018-04-09 21:57:16,812 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:57:16,817 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=resourceMap%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&star
t=0&rows=5000&indent=on
[ INFO] 2018-04-09 21:57:16,821 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0
&rows=1&indent=on
[ INFO] 2018-04-09 21:57:16,824 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:57:18,827 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1+OR+seri
esId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_data.1+OR+seriesId%3ApackageIndexingTest_P_S
_P_T_T_T_20180400.14.25_md&start=0&rows=3&indent=on
[ INFO] 2018-04-09 21:57:18,831 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1+OR+id%3Apacka
geIndexingTest_P_S_P_T_T_T_20180400.14.25_data.1+OR+id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.
14.25_md&start=0&rows=3&indent=on
[ INFO] 2018-04-09 21:57:23,334 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0
&rows=1&indent=on
[ INFO] 2018-04-09 21:57:23,336 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:57:23,379 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0
&rows=1&indent=on
[ INFO] 2018-04-09 21:57:23,381 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:57:23,383 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0
&rows=1&indent=on
[ INFO] 2018-04-09 21:57:23,385 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:57:23,387 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=seriesId%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0
&rows=1&indent=on
[ INFO] 2018-04-09 21:57:23,389 (HTTPService:doRequest:381) REQUEST URI= http://localhost:8983/sol
r/search_core/select/?q=id%3ApackageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1&start=0&rows=
1&indent=on
[ INFO] 2018-04-09 21:57:24,056 (IndexTaskProcessor:processTask:288) *****end to p
rocess update index task for packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1 in thread 110
[ INFO] 2018-04-09 21:57:24,062 (IndexTaskProcessor:processTask:315) Indexing complete for pid: pa
ckageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.1
[ERROR] 2018-04-09 21:57:24,438 (IndexTaskProcessor:checkReadinessProcessResourceMap:353) We wait
ed for another thread to finish indexing a pid with series id packageIndexingTest_P_S_P_T_T_T_20180
400.14.25_md for a while. Now we quit and can't index id packageIndexingTest_P_S_P_T_T_T_20180400.
14.25_resMap.2
[ERROR] 2018-04-09 21:57:24,438 (IndexTaskProcessor:processTask:297) Unable to process task for pi
d: packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.2
java.lang.Exception: We waited for another thread to finish indexing a pid with series id packageI
ndexingTest_P_S_P_T_T_T_20180400.14.25_md for a while. Now we quit and can't index id packageIndex
ingTest_P_S_P_T_T_T_20180400.14.25_resMap.2
[ INFO] 2018-04-09 21:57:24,442 (IndexTaskProcessor:newOrFailedIndexTaskExists:890) IndexTaskProce
ss.newOrFailedIndexTaskExists for id packageIndexingTest_P_S_P_T_T_T_20180400.14.25_resMap.2
```

Related issues:

Related to Infrastructure - Story #8702: Indexing Refactor Strategy

New

2018-09-24

History

#1 - 2018-04-09 22:30 - Rob Nahf

- Assignee changed from Dave Vieglais to Jing Tao
- Description updated

#2 - 2018-04-12 20:55 - Rob Nahf

- % Done changed from 0 to 30
- Assignee changed from Jing Tao to Rob Nahf
- Status changed from New to In Progress

This mainly turned out to be a problem with the length one thread waited for the other to finish up. Creating resourceMap records tasks was observed to take between 2 and 8 seconds, while the wait timed out at 5 seconds.

I upped the number of half-second waits in SANDBOX (ORC) to 30 (15 seconds), and the misbehavior went away.

#3 - 2018-09-24 15:17 - Rob Nahf

- Related to Story #8702: Indexing Refactor Strategy added