

Infrastructure - Story #8363

indexer shutdown generates index tasks

2018-02-12 21:42 - Rob Nahf

Status:	New	Start date:	2018-02-12
Priority:	Normal	Due date:	
Assignee:	Rob Nahf	% Done:	0%
Category:	d1_indexer	Estimated time:	0.00 hour
Target version:	CCI-2.4.0		
Story Points:			

Description

Seen in STAGE, somehow the index processor generated about 15k tasks (after processing 215k tasks over the weekend) during a service stop. It also created about 12.5 failures. Before trying to stop services, this the status of postgres:

```
d1-index-queue=# select status, count(*) from index_task group by status;
```

status	count
NEW	5
FAILED	1659
IN PROCESS	367

(3 rows)

Execution of /etc/init.d/d1-index-task-processor stop timed out.
I performed /etc/init.d/d1-index-task-generator stop successfully, getting an [OK]
then I performed /etc/init.d/d1-processing stop on UCSB, also getting an [OK]

examination of the indexing log file a couple minuted later showed this:

```
[ INFO] 2018-02-12 20:36:08,975 (IndexTaskProcessor:logProcessorLoad:245) new tasks:0, tasks previously failed: 1661
[ INFO] 2018-02-12 20:36:09,361 (IndexTaskProcessor:processFailedIndexTaskQueue:226) IndexTaskProcessor.processFailedIndexTaskQueue with size 0
[ WARN] 2018-02-12 20:36:09,361 (IndexTaskProcessorJob:execute:58) processing job [org.dataone.cn.index.processor.IndexTaskProcessorJob@515de84e] finished execution of index task processor [org.dataone.cn.index.processor.IndexTaskProcessor@20621d44]
[ WARN] 2018-02-12 20:36:26,571 (IndexTaskProcessorScheduler:stop:99) stopping index task processor quartz scheduler [org.dataone.cn.index.processor.IndexTaskProcessorScheduler@103bbd22] ...
[ INFO] 2018-02-12 20:36:26,572 (QuartzScheduler:standby:572) Scheduler QuartzScheduler_$_NON_CLUSTERED paused.
[ INFO] 2018-02-12 20:36:26,572 (IndexTaskProcessorScheduler:stop:111) Scheuler.interrupt method can't succeed to interrupt the d1 index job and the static method IndexTaskProcessorJob.interruptCurrent() will be called.
[ WARN] 2018-02-12 20:36:26,572 (IndexTaskProcessorJob:interruptCurrent:92) IndexTaskProcessorJob class [1806183035] interruptCurrent called, shutting down processor [org.dataone.cn.index.processor.IndexTaskProcessor@20621d44]
[ WARN] 2018-02-12 20:36:26,573 (IndexTaskProcessor:shutdownExecutor:952) processor [org.dataone.cn.index.processor.IndexTaskProcessor@20621d44] Shutting down the ExecutorService. Will allow active tasks to finish; will cancel submitted tasks
and return them to NEW status, wait for active tasks to finish, then return any remaining task not yet submitted to NEW status....
[ WARN] 2018-02-12 20:36:26,573 (IndexTaskProcessor:shutdownExecutor:955) ...1.) closing Executorservice to new tasks...
[ WARN] 2018-02-12 20:36:26,574 (IndexTaskProcessor:shutdownExecutor:957) ...2.) cancelling cancelable futures...
[ WARN] 2018-02-12 20:36:26,575 (IndexTaskProcessor:shutdownExecutor:958) ...number of futures: 591344
[ WARN] 2018-02-12 20:36:26,575 (IndexTaskProcessor:shutdownExecutor:959) ... number of tasks in f
```

utures map: 591344

15 minutes or so later, the log showed this:

```
[ WARN] 2018-02-12 20:36:26,573 (IndexTaskProcessor:shutdownExecutor:955) ...1.) closing ExecutorsService to new tasks...
[ WARN] 2018-02-12 20:36:26,574 (IndexTaskProcessor:shutdownExecutor:957) ...2.) cancelling cancellable futures...
[ WARN] 2018-02-12 20:36:26,575 (IndexTaskProcessor:shutdownExecutor:958) ...number of futures: 591344
[ WARN] 2018-02-12 20:36:26,575 (IndexTaskProcessor:shutdownExecutor:959) ... number of tasks in futures map: 591344
[ WARN] 2018-02-12 20:52:30,811 (IndexTaskProcessor:shutdownExecutor:988) ...number of (cancellable) runnables/tasks reset to new: 0
[ WARN] 2018-02-12 20:52:30,811 (IndexTaskProcessor:shutdownExecutor:989) ...number of (cancellable) runnables not mapped to tasks: 0
[ WARN] 2018-02-12 20:52:30,811 (IndexTaskProcessor:shutdownExecutor:990) ...number of uncancellable runnables: 591344 (completed or in process)
[ WARN] 2018-02-12 20:52:30,812 (IndexTaskProcessor:shutdownExecutor:993) ...3.) waiting (with timeout) for active futures to finish...
[ WARN] 2018-02-12 20:52:30,812 (IndexTaskProcessor:shutdownExecutor:998) ...4.) Reviewing remaining uncancellables to check for completion, returning incomplete ones to NEW status...
[ WARN] 2018-02-12 20:52:30,835 (IndexTaskProcessor:shutdownExecutor:1026) ...5.) Calling shutdownNow on the executor service.
[ WARN] 2018-02-12 20:52:30,835 (IndexTaskProcessor:shutdownExecutor:1028) ... .... number of runnables still waiting: 0
[ WARN] 2018-02-12 20:52:30,835 (IndexTaskProcessor:shutdownExecutor:1030) ...6.) returning preSubmitted tasks to NEW status...
[ WARN] 2018-02-12 20:52:30,835 (IndexTaskProcessor:shutdownExecutor:1031) ... .... number of preSubmitted tasks: 34735
[ INFO] 2018-02-12 20:52:30,835 (IndexTask:markNew:454) Even though it was masked new, it is still considered failed for id testGetPackage_2017119234441164 since it was tried to many times.
[ERROR] 2018-02-12 20:52:30,891 (IndexTaskProcessor:shutdownExecutor:1038) ..... Exception thrown trying to return task to NEW status for pid: testGetPackage_2017119234441164
org.springframework.orm.hibernate3.HibernateOptimisticLockingFailureException: Object of class [org.dataone.cn.index.task.IndexTask] with identifier [13071797]: optimistic locking failed; nested exception is org.hibernate.StaleObjectStateException: Row was updated or deleted by another transaction (or unsaved-value mapping was incorrect): [org.dataone.cn.index.task.IndexTask#13071797]

...

[ INFO] 2018-02-12 20:54:19,618 (IndexTask:markNew:454) Even though it was masked new, it is still considered failed for id P3_201622214921901 since it was tried to many times.
[ WARN] 2018-02-12 20:54:19,621 (IndexTaskProcessor:shutdownExecutor:1036) ... preSubmittedTask for pid P3_201622214921901 returned to NEW status.
[ WARN] 2018-02-12 20:54:19,623 (IndexTaskProcessor:shutdownExecutor:1036) ... preSubmittedTask for pid resource_map_doi:10.5065/D6VD6WFP returned to NEW status.
[ INFO] 2018-02-12 20:54:19,623 (IndexTask:markNew:454) Even though it was masked new, it is still considered failed for id testGetPackage_NotAuthorized_201710605522454 since it was tried to many times.
[ WARN] 2018-02-12 20:54:19,626 (IndexTaskProcessor:shutdownExecutor:1036) ... preSubmittedTask for pid testGetPackage_NotAuthorized_201710605522454 returned to NEW status.
[ WARN] 2018-02-12 20:54:19,628 (IndexTaskProcessor:shutdownExecutor:1036) ... preSubmittedTask for pid resource_map_urn:uuid:d3606ccb-2d50-4723-ae45-c0d01b817e48 returned to NEW status.
[ WARN] 2018-02-12 20:54:19,631 (IndexTaskProcessor:shutdownExecutor:1036) ... preSubmittedTask for pid resource_map_doi:10.18739/A2165F returned to NEW status.
[ WARN] 2018-02-12 20:54:19,631 (IndexTaskProcessor:shutdownExecutor:1041) .....7.) DONE with shutting down IndexTaskProcessor.
[ INFO] 2018-02-12 20:54:19,631 (IndexTaskProcessorScheduler:stop:113) The scheduler.interrupt method seems not interrupt the dl index job and the static method IndexTaskProcessorJob.interruptCurrent() was called.
[ WARN] 2018-02-12 20:54:19,632 (IndexTaskProcessorScheduler:stop:128) Job scheduler [org.dataone.cn.index.processor.IndexTaskProcessorScheduler@103bbd22] finished executing all jobs. The dl-index-processor shut down successfully.=====
```

but postgres yielded this:

```
d1-index-queue=# select status, count(*) from index_task group by status;
 status | count
-----+-----
 NEW    | 15367
 FAILED | 14032
(2 rows)
```

indexer shutdowns are a stubborn problem...

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

History

#1 - 2018-09-24 15:22 - Rob Nahf

- Related to Story #8702: Indexing Refactor Strategy added

#2 - 2018-09-24 15:25 - Rob Nahf

the new design will introduce back-pressure so that many fewer tasks are spooled into IN_PROCESS status. It is likely issues in shutdown that cause new tasks to be created, instead of resetting the status of existing tasks, but the shutdown method should also be reviewed to see why this is occurring in the first place, even in small numbers.