

Arvados - Bug #6074

[API] collections are being returned too slowly + NoMemoryError

05/19/2015 04:55 PM - Ward Vandewege

| | | | |
|---|-------------------|------------------------|------------|
| Status: | Resolved | Start date: | 06/03/2015 |
| Priority: | Normal | Due date: | |
| Assigned To: | Tom Clegg | % Done: | 100% |
| Category: | API | Estimated time: | 0.00 hour |
| Target version: | 2015-06-10 sprint | | |
| Description | | | |
| <p>On Mon, May 18, 2015 at 06:51:02PM -0700, Misha Zatsman wrote:</p> <pre>> I've been working on the pull lists lately, using stale data because the > api server has been frustratingly slow to respond to the data manager. Here > you can see it taking about a minute plus to respond, before eventually > failing: > > Is there a fix in the works? If the API server takes one second per > collection to respond, the data manager will never be able to run in a > reasonable amount of time. > > This all happened around 6:30 PM san francisco time in case you want to > look at the qrlhi api server logs. > > 2015/05/19 01:23:41 15570 collections read, 49 new in last batch, > 2014-10-28T15:36:52Z latest modified date, 16113 8991033 250878093 > avg,max,total manifest size > 2015/05/19 01:24:36 15619 collections read, 49 new in last batch, > 2014-10-29T11:50:18Z latest modified date, 16064 8991033 250909919 > avg,max,total manifest size > 2015/05/19 01:25:36 15667 collections read, 48 new in last batch, > 2014-10-30T00:46:12Z latest modified date, 16018 8991033 250948820 > avg,max,total manifest size > 2015/05/19 01:26:36 15716 collections read, 49 new in last batch, > 2014-10-30T19:30:56Z latest modified date, 15976 8991033 251073129 > avg,max,total manifest size > 2015/05/19 01:27:43 15765 collections read, 49 new in last batch, > 2014-10-31T23:45:52Z latest modified date, 15927 8991033 251083331 > avg,max,total manifest size > 2015/05/19 01:28:58 15814 collections read, 49 new in last batch, > 2014-11-02T23:45:31Z latest modified date, 15879 8991033 251116752 > avg,max,total manifest size > 2015/05/19 01:30:19 15863 collections read, 49 new in last batch, > 2014-11-04T11:29:52Z latest modified date, 15840 8991033 251277070 > avg,max,total manifest size > 2015/05/19 01:31:33 15912 collections read, 49 new in last batch, > 2014-11-04T16:25:36Z latest modified date, 15798 8991033 251385001 > avg,max,total manifest size > 2015/05/19 01:32:47 15961 collections read, 49 new in last batch, > 2014-11-05T04:30:20Z latest modified date, 15751 8991033 251406912 > avg,max,total manifest size > 2015/05/19 01:33:58 16010 collections read, 49 new in last batch, > 2014-11-05T21:50:34Z latest modified date, 15705 8991033 251430464 > avg,max,total manifest size > 2015/05/19 01:35:21 16059 collections read, 49 new in last batch, > 2014-11-06T16:12:59Z latest modified date, 15658 8991033 251456615 > avg,max,total manifest size > 2015/05/19 01:37:35 16108 collections read, 49 new in last batch, > 2014-11-06T23:46:04Z latest modified date, 45871 54080537 738893606 > avg,max,total manifest size > 2015/05/19 01:38:46 16157 collections read, 49 new in last batch, > 2014-11-07T20:08:25Z latest modified date, 55791 54080537 901420414 > avg,max,total manifest size</pre> | | | |

```
> 2015/05/19 01:39:32 16206 collections read, 49 new in last batch,
> 2014-11-07T22:52:48Z latest modified date, 55637 54080537 901649169
> avg,max,total manifest size
> 2015/05/19 01:40:27 16255 collections read, 49 new in last batch,
> 2014-11-08T07:17:22Z latest modified date, 55473 54080537 901715163
> avg,max,total manifest size
> 2015/05/19 01:41:09 16303 collections read, 48 new in last batch,
> 2014-11-08T21:04:18Z latest modified date, 55314 54080537 901778597
> avg,max,total manifest size
> 2015/05/19 01:41:53 16352 collections read, 49 new in last batch,
> 2014-11-09T17:48:07Z latest modified date, 55149 54080537 901804169
> avg,max,total manifest size
> 2015/05/19 01:43:01 16401 collections read, 49 new in last batch,
> 2014-11-10T18:26:31Z latest modified date, 65035 54085104 1066634271
> avg,max,total manifest size
> 2015/05/19 01:44:01 Error querying collections: API server (
> qrlhi.arvadosapi.com) returned 422: 422 Unprocessable Entity. Additional
> details: #<NoMemoryError: failed to allocate memory>
> exit status 1
>
```

Also - qr1hi is currently running [31d8ece](#). That is after [#5834](#) was merged. And still there is a NoMemoryError. Here's the backtrace from the logs:

```
#<NoMemoryError: failed to allocate memory>
/data/www/qrlhi.arvadosapi.com/releases/20150508140613/app/controllers/application_controller.rb:1
89:in `dump'
/data/www/qrlhi.arvadosapi.com/releases/20150508140613/app/controllers/application_controller.rb:1
89:in `send_json'
/data/www/qrlhi.arvadosapi.com/releases/20150508140613/app/controllers/application_controller.rb:4
87:in `render_list'
/data/www/qrlhi.arvadosapi.com/releases/20150508140613/app/controllers/application_controller.rb:7
6:in `index'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_co
ntroller/metal/implicit_render.rb:4:in `send_action'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/abstract_
controller/base.rb:167:in `process_action'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_co
ntroller/metal/rendering.rb:10:in `process_action'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/abstract_
controller/callbacks.rb:18:in `block in process_action'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activerecord-3.2.17/lib/active_
support/callbacks.rb:579:in `__run__613498188452779378__process_action__1315541577497190102__callb
acks'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activerecord-3.2.17/lib/active_
support/callbacks.rb:405:in `__run_callback'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activerecord-3.2.17/lib/active_
support/callbacks.rb:385:in `__run_process_action_callbacks'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activerecord-3.2.17/lib/active_
support/callbacks.rb:81:in `run_callbacks'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/abstract_
controller/callbacks.rb:17:in `process_action'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_co
ntroller/metal/rescue.rb:29:in `process_action'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_co
ntroller/metal/instrumentation.rb:30:in `block in process_action'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activerecord-3.2.17/lib/active_
support/notifications.rb:123:in `block in instrument'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activerecord-3.2.17/lib/active_
support/notifications/instrumenter.rb:20:in `instrument'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activerecord-3.2.17/lib/active_
support/notifications.rb:123:in `instrument'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_co
ntroller/metal/instrumentation.rb:29:in `process_action'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_co
ntroller/metal/params_wrapper.rb:207:in `process_action'
/var/www/qrlhi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activerecord-3.2.17/lib/active_
```

```
record/railties/controller_runtime.rb:18:in `process_action'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/abstract_  
controller/base.rb:121:in `process'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/abstract_  
controller/rendering.rb:45:in `process'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_co  
ntroller/metal.rb:203:in `dispatch'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_co  
ntroller/metal/rack_delegation.rb:14:in `dispatch'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_co  
ntroller/metal.rb:246:in `block in action'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/routing/route_set.rb:73:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/routing/route_set.rb:73:in `dispatch'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/routing/route_set.rb:36:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/journey-1.0.4/lib/journey/route  
r.rb:68:in `block in call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/journey-1.0.4/lib/journey/route  
r.rb:56:in `each'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/journey-1.0.4/lib/journey/route  
r.rb:56:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/routing/route_set.rb:608:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/omniauth-1.1.1/lib/omniauth/str  
ategy.rb:177:in `call!'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/omniauth-1.1.1/lib/omniauth/str  
ategy.rb:157:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/omniauth-1.1.1/lib/omniauth/bui  
lder.rb:48:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/middleware/best_standards_support.rb:17:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-1.4.5/lib/rack/etag.rb:23:  
in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-1.4.5/lib/rack/conditional  
get.rb:25:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/middleware/head.rb:14:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/middleware/params_parser.rb:21:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/middleware/flash.rb:242:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-1.4.5/lib/rack/session/abs  
tract/id.rb:210:in `context'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-1.4.5/lib/rack/session/abs  
tract/id.rb:205:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/middleware/cookies.rb:341:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activerecord-3.2.17/lib/active_  
record/query_cache.rb:64:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activerecord-3.2.17/lib/active_  
record/connection_adapters/abstract/connection_pool.rb:479:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/middleware/callbacks.rb:28:in `block in call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activesupport-3.2.17/lib/active  
_support/callbacks.rb:405:in `__run__2522496659413941100__call__2134072409839546070__callbacks'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activesupport-3.2.17/lib/active  
_support/callbacks.rb:405:in `__run_callback'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activesupport-3.2.17/lib/active  
_support/callbacks.rb:385:in `_run_call_callbacks'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activesupport-3.2.17/lib/active  
_support/callbacks.rb:81:in `run_callbacks'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/middleware/callbacks.rb:27:in `call'  
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_di  
spatch/middleware/debug_exceptions.rb:16:in `call'
```

```

/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_dispatch/middleware/show_exceptions.rb:56:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/railties-3.2.17/lib/rails/rack/logger.rb:32:in `call_app'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/railties-3.2.17/lib/rails/rack/logger.rb:16:in `block in call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activesupport-3.2.17/lib/active_support/tagged_logging.rb:22:in `tagged'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/railties-3.2.17/lib/rails/rack/logger.rb:16:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_dispatch/middleware/request_id.rb:22:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-1.4.5/lib/rack/methodoverride.rb:21:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-1.4.5/lib/rack/runtime.rb:17:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/activesupport-3.2.17/lib/active_support/cache/strategy/local_cache.rb:72:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-1.4.5/lib/rack/lock.rb:15:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-cache-1.2/lib/rack/cache/context.rb:136:in `forward'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-cache-1.2/lib/rack/cache/context.rb:245:in `fetch'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-cache-1.2/lib/rack/cache/context.rb:185:in `lookup'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-cache-1.2/lib/rack/cache/context.rb:66:in `call!'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/rack-cache-1.2/lib/rack/cache/context.rb:51:in `call'
/data/www/qr1hi.arvadosapi.com/releases/20150508140613/app/middlewares/arvados_api_token.rb:59:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/actionpack-3.2.17/lib/action_dispatch/middleware/remote_ip.rb:31:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/railties-3.2.17/lib/rails/engine.rb:484:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/railties-3.2.17/lib/rails/application.rb:231:in `call'
/var/www/qr1hi.arvadosapi.com/shared/vendor_bundle/ruby/2.1.0/gems/railties-3.2.17/lib/rails/railtie/configurable.rb:30:in `method_missing'
/usr/lib/ruby/vendor_ruby/phusion_passenger/rack/thread_handler_extension.rb:74:in `process_request'
/usr/lib/ruby/vendor_ruby/phusion_passenger/request_handler/thread_handler.rb:141:in `accept_and_process_next_request'
/usr/lib/ruby/vendor_ruby/phusion_passenger/request_handler/thread_handler.rb:109:in `main_loop'
/usr/lib/ruby/vendor_ruby/phusion_passenger/request_handler.rb:455:in `block (3 levels) in start_threads'

```

```

Error 1431999841+3bb31288: 422
  Rendered text template (0.0ms)
Completed 422 Unprocessable Entity in 60139.7ms (Views: 1.1ms | ActiveRecord: 48283.5ms)

```

Subtasks:

| | |
|--|-----------------|
| Task # 6123: Determine why still NoMemoryError even with API memory-limiting feature | Resolved |
| Task # 6110: Review 6074-collections-index @ 22c8b63 | Resolved |

Related issues:

| | | |
|--|-----------------|-------------------|
| Related to Arvados - Bug #5834: [API] Be careful about memory usage when resp... | Resolved | 04/30/2015 |
| Blocks Arvados - Story #3408: [Keep] Implement Production Data Manager | Resolved | 07/29/2014 |

Associated revisions

Revision a12cf91d - 06/04/2015 07:29 PM - Tom Clegg

Merge branch '6074-collections-index' closes #6074

History

#1 - 05/19/2015 04:56 PM - Ward Vandewege

- Description updated

#2 - 05/19/2015 05:59 PM - Ward Vandewege

- Description updated

#3 - 05/19/2015 06:03 PM - Ward Vandewege

- Subject changed from [API] collections are being returned too slowly to [API] collections are being returned too slowly + NoMemoryError

- Description updated

#4 - 05/20/2015 06:20 PM - Brett Smith

- Category set to API

- Target version changed from Bug Triage to Arvados Future Sprints

#5 - 05/20/2015 07:02 PM - Brett Smith

- Target version changed from Arvados Future Sprints to 2015-06-10 sprint

#6 - 05/20/2015 07:27 PM - Tom Clegg

- Assigned To set to Tom Clegg

#7 - 05/20/2015 07:32 PM - Tom Clegg

- Story points set to 1.0

#8 - 06/03/2015 07:54 PM - Tom Clegg

At least part of the slowness is fixed by using `octet_length()` instead of `length()` in `ApplicationController#limit_database_read`, [69f592e](#)

Before

```
arvados_production=> explain analyze SELECT collections."manifest_text", collections."owner_uuid", collection
s."uuid", collections."modified_at", (length(manifest_text)) as read_length FROM "collections" WHERE (expires
_at IS NULL or expires_at > CURRENT_TIMESTAMP) AND ((collections.modified_at >= '2014-11-08 07:17:22.000000'))
ORDER BY "collections"."id" ASC LIMIT 1000 OFFSET 0;
```

QUERY PLAN

```
-----
Limit (cost=562.06..564.56 rows=1000 width=212) (actual time=4721.771..4723.577 rows=1000 loops=1)
-> Sort (cost=562.06..568.73 rows=2667 width=212) (actual time=4721.764..4722.385 rows=1000 loops=1)
    Sort Key: id
    Sort Method: top-N heapsort Memory: 352kB
    -> Seq Scan on collections (cost=0.00..415.84 rows=2667 width=212) (actual time=0.026..4717.554 row
s=3633 loops=1)
    Filter: ((modified_at >= '2014-11-08 07:17:22'::timestamp without time zone) AND ((expires_at I
S NULL) OR (expires_at > now())))
Total runtime: 4724.374 ms
(7 rows)
```

After

```
arvados_production=> explain analyze SELECT collections."manifest_text", collections."owner_uuid", collection
s."uuid", collections."modified_at", (octet_length(manifest_text)) as read_length FROM "collections" WHERE (e
xpries_at IS NULL or expires_at > CURRENT_TIMESTAMP) AND ((collections.modified_at >= '2014-11-08 07:17:22.000
000')) ORDER BY "collections"."id" ASC LIMIT 1000 OFFSET 0;
```

QUERY PLAN

```
-----
Limit (cost=562.06..564.56 rows=1000 width=212) (actual time=11.248..13.639 rows=1000 loops=1)
-> Sort (cost=562.06..568.73 rows=2667 width=212) (actual time=11.247..12.075 rows=1000 loops=1)
    Sort Key: id
    Sort Method: top-N heapsort Memory: 352kB
    -> Seq Scan on collections (cost=0.00..415.84 rows=2667 width=212) (actual time=0.006..6.007 rows=3
633 loops=1)
    Filter: ((modified_at >= '2014-11-08 07:17:22'::timestamp without time zone) AND ((expires_at I
S NULL) OR (expires_at > now())))
```

Total runtime: 14.570 ms
(7 rows)

#9 - 06/03/2015 09:45 PM - Brett Smith

- Category deleted (API)

Reviewing [22c8b63](#). To be clear upfront: this is a worthwhile change, and the code is good to merge. My points are more about understanding context than anything.

This change should help mitigate Misha's concern about slowness, but it doesn't seem to do anything to mitigate the memory pressure that Ward observed. Do I understand correctly? Is this branch expected to close the ticket, or just one improvement of several?

Related to the change to make `max_index_database_read` a hard limit:

- It'd be good to update or remove the comment in `application.default.yml`, "Note this is a threshold, not a limit. Record collection stops **after** reading this much data."
- Off-by-one concerns: Should the `if read_total` condition use `>` instead of `>=`? Or at least not decrement `new_limit` when `read_total == max_index_database_read`?

Thanks.

#10 - 06/03/2015 09:45 PM - Brett Smith

- Category set to API

#11 - 06/04/2015 05:37 AM - Tom Clegg

Brett Smith wrote:

This change should help mitigate Misha's concern about slowness, but it doesn't seem to do anything to mitigate the memory pressure that Ward observed. Do I understand correctly? Is this branch expected to close the ticket, or just one improvement of several?

Sorry for missing explanation. Here's what I should have offered.

Two memory-impacting problems are fixed here:

1. Our "scouting" query was unintentionally retrieving `manifest_text`, because
 - the client-provided `@select` was applied to the `@objects` query before `limit_database_read`, and
 - `ActiveRecord's #select` method adds to the list of selected fields, rather than replacing it.
2. The effective maximum size was the configured maximum (default 128 MiB) plus the maximum `manifest_text` size.

The first problem *might* not have contributed to the observed OOM problem, if Ruby's garbage collection was doing well. But even if it didn't, it was wasteful, so I fixed it.

The second problem, again, *might* not have caused the observed OOM problem, but it seems clear that it makes the effective limit much larger than anticipated. The new behavior still provides the same non-starvation guarantee (i.e., you will always make progress with each page) but limits the result to approximately `max(configured,max(actual))` instead of `configured+max(actual)`. IOW, if you have enough memory to return your biggest record, then you have enough memory to do every "index" query.

Related to the change to make `max_index_database_read` a hard limit:

- It'd be good to update or remove the comment in `application.default.yml`, "Note this is a threshold, not a limit. Record collection stops **after** reading this much data."

Indeed, thanks. Updated.

- Off-by-one concerns: Should the `if read_total` condition use `>` instead of `>=`? Or at least not decrement `new_limit` when `read_total == max_index_database_read`?

I considered that, and yes that would be more closely adhere to the config knob documentation, but

- The number is already at a much higher precision than is sensible. Will anyone really care whether we stop at 134217727 or 134217728?
- The row will have a bunch of other stuff besides `manifest_text`, so this means we'll be using more than `max_index_database_read` bytes when we fetch it.

But if you still think it's worth adding the extra test "return the last one too if it fills the max right down to the byte", saying "yes please" will be enough to twist my arm.

...and now for the good stuff

Meanwhile, I decided I'd made the use of `#find_each` superfluous, so I changed it to `#each`. And then I read the docs. And then things really started getting good. I ended up at [602706e](#).

I still think the above changes are worth keeping, but they pale in comparison to the `find_each` bug. I confirmed that `find_each` was checking the manifest sizes on the wrong records, which obviously would have made the "prevent OOM" mechanism ineffective in many cases like the one at hand which uses "order":"modified_at". (To see this, you can change it back to `find_each` and run the new test on it. Looking at the `length()` queries in your older log files doesn't demonstrate it very well because of the testing bugs.)

Underneath it all was another oddity/bug, "order params are silently ignored if you use "select" and don't include the columns you're ordering on" but I think that can be dealt with separately so I added [#6240](#).

#12 - 06/04/2015 07:10 PM - Brett Smith

Thanks for the detailed explanation, and for fixing my bugs. [602706ef](#) is good to merge.

#13 - 06/04/2015 07:35 PM - Tom Clegg

- *Status changed from New to Resolved*

Applied in changeset `arvados|commit:a12cf91db3eeced782d9718af30553f8866ad81`.