

Arvados - Bug #18004

Cached federated token creation race condition on RailsAPI

08/10/2021 08:04 PM - Lucas Di Pentima

Status:	Resolved	Start date:	08/17/2021
Priority:	Normal	Due date:	
Assigned To:	Lucas Di Pentima	% Done:	100%
Category:	API	Estimated time:	0.00 hour
Target version:	2021-09-01 sprint		
Description			
We got a report that some users were sometimes seeing 502 errors on wb2 when accessing collections, nginx error log show a primary key violation on insert:			
<pre>App 18334 stderr: [2021-08-06 09:54:55.4629 18396/0x00000000022f0380(Worker 1) utils.rb:87]: *** Exception ActiveRecord::RecordNotUnique in Rack application object (PG::UniqueViolation: ERROR: duplicate key value violates unique constraint "index_api_client_authorizations_on_api_token" App 18334 stderr: DETAIL: Key (api_token)=(TOKEN) already exists. App 18334 stderr: : INSERT INTO "api_client_authorizations" ("api_token", "api_client_id", "user_i d", "expires_at", "created_at", "updated_at", "uuid") VALUES (\$1, \$2, \$3, \$4, \$5, \$6, \$7) RETURNIN G "id") (process 18396, thread 0x00000000022f0380(Worker 1)): App 18334 stderr: from /var/www/arvados-api/shared/vendor_bundle/ruby/2.7.0/gems/activerecord-5.2. 6/lib/active_record/connection_adapters/postgresql_adapter.rb:611:in `async_exec_params' [...] App 18334 stderr: from /var/www/arvados-api/current/app/models/arvados_model.rb:145:in `create' App 18334 stderr: from /var/www/arvados-api/shared/vendor_bundle/ruby/2.7.0/gems/activerecord-5.2. 6/lib/active_record/relation.rb:85:in `block in create' App 18334 stderr: from /var/www/arvados-api/shared/vendor_bundle/ruby/2.7.0/gems/activerecord-5.2. 6/lib/active_record/relation.rb:281:in `scoping' App 18334 stderr: from /var/www/arvados-api/shared/vendor_bundle/ruby/2.7.0/gems/activerecord-5.2. 6/lib/active_record/relation.rb:85:in `create' App 18334 stderr: from /var/www/arvados-api/shared/vendor_bundle/ruby/2.7.0/gems/activerecord-5.2. 6/lib/active_record/relation.rb:164:in `find_or_create_by' App 18334 stderr: from /var/www/arvados-api/shared/vendor_bundle/ruby/2.7.0/gems/activerecord-5.2. 6/lib/active_record/querying.rb:8:in `find_or_create_by' App 18334 stderr: from /var/www/arvados-api/current/app/models/api_client_authorization.rb:367:in `block in validate' App 18334 stderr: from /var/www/arvados-api/current/lib/current_api_client.rb:138:in `block in act _as_system_user' App 18334 stderr: from /var/www/arvados-api/current/lib/current_api_client.rb:149:in `act_as_user' App 18334 stderr: from /var/www/arvados-api/current/lib/current_api_client.rb:137:in `act_as_syste m_user' App 18334 stderr: from /var/www/arvados-api/current/app/models/api_client_authorization.rb:312:in `validate' App 18334 stderr: from /var/www/arvados-api/current/app/middlewares/arvados_api_token.rb:52:in `bl ock in call' App 18334 stderr: from /var/www/arvados-api/current/app/middlewares/arvados_api_token.rb:49:in `ea ch' App 18334 stderr: from /var/www/arvados-api/current/app/middlewares/arvados_api_token.rb:49:in `ca ll' App 18334 stderr: from /var/www/arvados-api/shared/vendor_bundle/ruby/2.7.0/gems/railties-5.2.6/li b/rails/engine.rb:524:in `call' App 18334 stderr: from /var/www/arvados-api/shared/vendor_bundle/ruby/2.7.0/gems/railties-5.2.6/li b/rails/railtie.rb:190:in `public_send' App 18334 stderr: from /var/www/arvados-api/shared/vendor_bundle/ruby/2.7.0/gems/railties-5.2.6/li b/rails/railtie.rb:190:in `method_missing' App 18334 stderr: from /usr/share/ruby/vendor_ruby/phusion_passenger/rack/thread_handler_extensio n.rb:97:in `process_request' App 18334 stderr: from /usr/share/ruby/vendor_ruby/phusion_passenger/request_handler/thread_handle r.rb:152:in `accept_and_process_next_request' App 18334 stderr: from /usr/share/ruby/vendor_ruby/phusion_passenger/request_handler/thread_handle r.rb:113:in `main_loop' App 18334 stderr: from /usr/share/ruby/vendor_ruby/phusion_passenger/request_handler.rb:416:in `bl</pre>			

```
ock (3 levels) in start_threads'
App 18334 stderr: from /usr/share/ruby/vendor_ruby/phusion_passenger/utils.rb:113:in `block in create_thread_and_abort_on_exception'
[ 2021-08-06 09:54:55.4633 18288/7f15cbfff700 age/Cor/Con/InternalUtils.cpp:112 ]: [Client 2-787]
Sending 502 response: application did not send a complete response
```

It seems the problem is related to a race when the remote token isn't cached on the cluster and many request for the same token are in flight, some of them will try to create the same cached token producing a collision.

Subtasks:

Task # 18025: Review 18004-cached-token-race-condition **Resolved**

Associated revisions

Revision 9214382f - 08/11/2021 07:40 PM - Lucas Di Pentima

Merge branch '18005-collection-versioning-regression' into main. Closes #18004

Arvados-DCO-1.1-Signed-off-by: Lucas Di Pentima <lucas.dipentima@curii.com>

Revision 22fcbed2 - 08/24/2021 03:31 PM - Lucas Di Pentima

Merge branch '18004-cached-token-race-condition' into main. Closes #18004

Arvados-DCO-1.1-Signed-off-by: Lucas Di Pentima <lucas.dipentima@curii.com>

History

#1 - 08/10/2021 08:25 PM - Lucas Di Pentima

Test exposing the bug at [0be7e79](#) - branch 18004-cached-token-race-condition
Test run with failure: <https://ci.arvados.org/job/developer-run-tests-remainder/2730/>

#2 - 08/11/2021 04:08 PM - Tom Clegg

Maybe we just need something like this (from users_controller.rb):

```
begin
  u = User.find_or_create_by(uuid: uuid)
rescue ActiveRecord::RecordNotUnique
  retry
end
```

...but we should make sure to only retry once, not infinity times, in case RecordNotUnique is caused by a constraint other than uuid.

#3 - 08/11/2021 07:38 PM - Lucas Di Pentima

- Status changed from New to In Progress

#4 - 08/11/2021 07:43 PM - Lucas Di Pentima

- % Done changed from 0 to 100
- Status changed from In Progress to Resolved

Applied in changeset [arvados!9214382fc7083d0310675ce3760144ad7b6cb070](#).

#5 - 08/11/2021 08:41 PM - Lucas Di Pentima

- Status changed from Resolved to In Progress

#6 - 08/17/2021 06:27 PM - Lucas Di Pentima

Updates at [516669845](#)
Test run: <https://ci.arvados.org/job/developer-run-tests/2643/>

- Found (and fixed) an additional race condition bug (both detected by the same test): the user record creation code suffered from the same issue as the token code.
- Locally ran the test 10 times in a row without failure (takes ~70s per run), every run creates 100 concurrent requests so I'm pretty confident the problem is fixed.

#7 - 08/18/2021 01:54 PM - Lucas Di Pentima

- Target version changed from 2021-08-18 sprint to 2021-09-01 sprint

#8 - 08/24/2021 03:28 PM - Tom Clegg

LGTM, thanks!

#9 - 08/24/2021 05:53 PM - Lucas Di Pentima

- % Done changed from 0 to 100

- Status changed from In Progress to Resolved

Applied in changeset [arvados|22fcbed20b2691bcdfa5854004dcb95aa1f2e40d](#).