

Arvados - Bug #14009

container request creation very slow when there are many potentially reusable containers

08/10/2018 12:24 PM - Joshua Randall

Status:	Resolved	Start date:	08/24/2018
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:	API	Estimated time:	0.00 hour
Target version:	2018-09-05 Sprint		
Description			
<p>Container request creation on our system was typically taking 90-120s (N.B. this is after disabling all audit logging so it has already been sped up dramatically from stock). We found this to be due to the queries that look up the reusable containers were the query conditions were pretty much all on columns that have no indices of any kind.</p> <p>Our interim solution has been to add a hash-based index (to keep index size down) on the `command` column:</p> <pre>arvados_api_production=# create index index_containers_on_command_hash on containers using hash (command);</pre> <p>After adding this index, container request creation was sped up by ~20x (to around 5s each).</p> <p>However, note that hash indices are not recommended for postgres versions < 10 as they are not included in WAL replication. An alternative would be to have explicit hash columns included in the table (which are btree indexed) that can be used to query the exact-match fields for container reuse.</p> <p>In addition to making for faster reuse queries, this should also pave the way to enable formulation of client-side container queries that don't need to transmit the large text values (such as mounts, command, etc) to the API server.</p>			
Subtasks:			
Task # 14046: Review 14009-container-reuse-index			Resolved

Associated revisions

Revision 9d413ff4 - 09/04/2018 03:36 PM - Tom Clegg

Merge branch '14009-container-reuse-index'

refs #14009

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <tclegg@veritasgenetics.com>

History

#1 - 08/15/2018 03:57 PM - Tom Morris

- Target version set to 2018-09-05 Sprint

#2 - 08/15/2018 04:23 PM - Tom Clegg

- Assigned To set to Tom Clegg

#3 - 08/24/2018 03:44 PM - Tom Clegg

- Status changed from New to In Progress

14009-container-reuse-index @ [6525b509825dbbf1cbe8b30b34080aafc4e5bde3](https://github.com/Arvados/arvados/commit/6525b509825dbbf1cbe8b30b34080aafc4e5bde3)

The test database is so small postgres declines to use the index, but I pasted the "create index" statement and a reuse query from the test suite into a dev cluster with lots of containers (9tee4) to confirm postgres uses it. The query took ~130ms before adding the index, ~1ms after.

#4 - 08/24/2018 08:30 PM - Peter Amstutz

Tom Clegg wrote:

14009-container-reuse-index @ [6525b509825dbbf1cbe8b30b34080aafc4e5bde3](https://github.com/Arvados/arvados/commit/6525b509825dbbf1cbe8b30b34080aafc4e5bde3)

The test database is so small postgres declines to use the index, but I pasted the "create index" statement and a reuse query from the test suite into a dev cluster with lots of containers (9tee4) to confirm postgres uses it. The query took ~130ms before adding the index, ~1ms after.

This seems like the sort of thing a hash index would be useful for, but Postgres 9.x hash indexes come with a recommendation not to use them, so this seems to be the next best thing. LGTM.

#5 - 09/04/2018 09:28 PM - Tom Clegg

- *Status changed from In Progress to Resolved*

#6 - 11/13/2018 09:02 PM - Tom Morris

- *Release set to 13*