

Arvados - Feature #14706

[Crunch2] Retain references + permissions to earlier containers when retrying a container request

01/08/2019 08:12 PM - Peter Amstutz

Status:	Resolved	Start date:	
Priority:	Normal	Due date:	
Assigned To:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
Description			
<p>The container request record lists the most recent container attempted to fulfill the request. This means when a cancelled container is retried, the earlier cancelled containers are not visible to the user: the container UUID is no longer mentioned in the container request record, which means that even if the client remembers the UUID, the user no longer has permission to retrieve the container record.</p> <p>(See #14870 for the related problem that the <i>logs</i> from previous attempts are not preserved in the container request's log collection.)</p> <p>Proposal:</p> <p>Need a column that has uuids of all containers. Can use array column, eg https://www.postgresql.org/docs/9.6/arrays.html, or JSONB column.</p> <p>Current data model has "container_uuid" as a singular value. It would be a backwards compatibility problem if that changed to be an array. API should report past attempts in a separate field, like "past_container_uuids".</p> <p>Unclear if it would be better in the underlying database to have a single array column (where first/last item is always the most recent attempt), or retain container_uuid column and add a past_container_uuids column.</p> <p>Need to be able to join array column to grant read permission to container records. Section 8.15.5 of postgres docs suggest this is something like:</p> <pre>container.uuid = ANY (container_request.past_container_uuids)</pre>			
Related issues:			
Related to Arvados - Feature #8018: [Crunch2] Identify container failure and ...		Resolved	09/23/2016
Related to Arvados - Story #14870: [API] Access logs from previous attempts a...		Resolved	03/01/2019

History

#1 - 01/08/2019 08:12 PM - Peter Amstutz

- Status changed from New to In Progress

#2 - 01/08/2019 08:30 PM - Peter Amstutz

- Description updated

- Status changed from In Progress to New

#3 - 01/08/2019 08:39 PM - Peter Amstutz

- Tracker changed from Bug to Feature

#4 - 02/15/2019 03:32 PM - Peter Amstutz

- Description updated

#5 - 02/19/2019 09:13 PM - Tom Clegg

- Related to Feature #8018: [Crunch2] Identify container failure and retry added

#6 - 02/19/2019 09:37 PM - Tom Clegg

I'm not sure adding an array of container UUIDs to the container_requests table would solve this problem. Often the most valuable troubleshooting information is in the log files, which would still be inaccessible.

It might be more useful to focus on preserving all relevant logs in the container request's log collection, even if they span multiple containers. Perhaps the API server could merge the logs: e.g., instead of replacing the CR's entire log collection when the container's log is updated, just copy the container's log files into a "container \${uuid}" subdir in the container request's log collection. This would disturb existing scripts/users who expect the log files to be at the top level, but it would be compatible with multiple concurrent containers (e.g., speculative retry without killing, and replication>1 service containers).

This also helps in the case where the container record itself is really what's wanted, since that is included in the container's log collection. (There are currently some exceptions -- e.g., a log collection isn't created at all when a container doesn't fit any instance type -- but those could be fixed.)

It's also worth addressing the permission issue, at least for admins (currently even the dispatcher isn't allowed to see that a container has state=Cancelled if all matching CRs have had different containers assigned!). If we need to do it for users, we should consider the performance implications of an array vs. a separate table to express the many-to-many relationship.

#7 - 02/20/2019 02:43 PM - Tom Clegg

One more refinement: Put a copy of the latest container's logs in the root dir of the container request's log collection, in addition to a subdir named after the container UUID. This way, existing scripts continue to work on new logs.

#8 - 02/20/2019 02:50 PM - Tom Clegg

- *Related to Story #14870: [API] Access logs from previous attempts after auto-retrying a container request added*

#9 - 02/21/2019 02:57 PM - Tom Clegg

- *Subject changed from [Crunch2] Retain record of container retries to [Crunch2] Retain references + permissions to earlier containers when retrying a container request*

- *Description updated*

#10 - 12/18/2019 08:47 PM - Peter Amstutz

- *Status changed from New to Resolved*

#11 - 12/18/2019 08:55 PM - Peter Amstutz

- *Target version deleted (To Be Groomed)*