

Arvados - Story #13146

[API] Endpoint to get projects shared with me

02/28/2018 09:07 PM - Peter Amstutz

Status: Resolved	Start date: 08/15/2018
Priority: Normal	Due date:
Assigned To: Peter Amstutz	% Done: 100%
Category:	Estimated time: 0.00 hour
Target version: 2018-09-05 Sprint	
Description	
<p>To get "projects shared with me" efficiently requires a new API endpoint. The way arv-mount currently determines what to list in "shared" currently requires looking at all projects and finding the ones where owner_uuid is not another project which is visible to us (meaning: users, non-project groups, or shared subprojects where the parent is not visible). This is expensive to compute on the client, but can probably be accomplished with a single query on the API server.</p> <p>Proposed endpoint: /arvados/v1/groups/shared</p> <p>Return in "items" a list of toplevel projects and collections which are shared with me. Toplevel projects are ones which are either owned by a user, owned by a non-project group, or owned by a group which the user cannot read.</p> <p>Support "includes" to return the set of owner objects as well (users, groups).</p> <p>Possible query (abbreviated):</p> <ul style="list-style-type: none">• select groups that are readable by current user AND<ul style="list-style-type: none">◦ the owner_uuid is a user (not a group) OR◦ the owner_uuid is a group is not a project or is not readable by the current user <pre>select uuid from groups where exists(select 1 from materialized_permission_view where user_uuid='2tlax-tpzed-51vcbmac4uv6bpb ' and target_uuid=groups.uuid) and (groups.owner_uuid in (select uuid from users) or not exists(select 1 from materialized_permission_view where user_uuid=' 2tlax-tpzed-51vcbmac4uv6bpb' and target_uuid=groups.owner_uuid)) and owner_uuid != '2tlax-tpzed-51vcbmac4uv6bpb';</pre>	
Subtasks:	
Task # 13950: Review 13146-shared-rails	Resolved
Task # 14060: shared endpoint	Resolved
Related issues:	
Related to Arvados - Bug #12990: [FUSE] Access shared/ is inefficient	In Progress
Related to Arvados - Story #13111: [WebDAV] Support browsing of project hiera...	Resolved 03/27/2018
Related to Arvados - Story #13218: Support browsing of projects shared with m...	New
Related to Arvados - Feature #14201: [API] Accept "exclude_home_project" flag...	Resolved 10/01/2018
Related to Arvados Workbench 2 - Support #15238: [Workbench] Change the conte...	New

Associated revisions

Revision 6d6a0b0f - 08/17/2018 05:13 PM - Peter Amstutz

Merge branch '13146-shared-rails' refs #13146

Arvados-DCO-1.1-Signed-off-by: Peter Amstutz <pamstutz@veritasgenetics.com>

History

#1 - 02/28/2018 09:07 PM - Peter Amstutz

- Status changed from New to In Progress

#2 - 02/28/2018 09:08 PM - Peter Amstutz

- Description updated

#3 - 02/28/2018 09:08 PM - Peter Amstutz

- Related to Bug #12990: [FUSE] Access shared/ is inefficient added

#4 - 03/07/2018 08:23 PM - Peter Amstutz

- Related to Story #13111: [WebDAV] Support browsing of project hierarchies added

#5 - 03/14/2018 07:24 PM - Tom Morris

- Tracker changed from Feature to Story
- Status changed from In Progress to New

#6 - 03/14/2018 07:28 PM - Tom Morris

- Related to Story #13218: Support browsing of projects shared with me in WebDAV added

#7 - 04/18/2018 03:06 PM - Peter Amstutz

- Description updated

#8 - 04/18/2018 03:16 PM - Peter Amstutz

- Description updated

The current algorithm used in the FUSE driver is:

- get all projects
- go through the list of projects and find projects which have an owner_uid that isn't in the list and isn't the current user. This has the effect of finding the projects for which the owner is either a user, a group which is not a project, or a group which is not readable by the current user.
- The list of "roots" consists of each owner (if the owner is readable but not a project) or the project itself (if not).

#9 - 04/18/2018 04:48 PM - Peter Amstutz

- Description updated
- Target version changed from To Be Groomed to Arvados Future Sprints
- Story points set to 3.0

#10 - 08/01/2018 04:01 PM - Tom Morris

- Target version changed from Arvados Future Sprints to 2018-08-15 Sprint

#11 - 08/01/2018 04:08 PM - Peter Amstutz

- Assigned To set to Peter Amstutz
- Target version changed from 2018-08-15 Sprint to Arvados Future Sprints

#12 - 08/01/2018 04:21 PM - Peter Amstutz

- Target version changed from Arvados Future Sprints to 2018-08-15 Sprint

#13 - 08/03/2018 01:41 PM - Peter Amstutz

- Status changed from New to In Progress

#14 - 08/15/2018 03:13 PM - Peter Amstutz

- Target version changed from 2018-08-15 Sprint to 2018-09-05 Sprint

#15 - 08/16/2018 03:36 PM - Peter Amstutz

13146-shared-rails @ [0ba8a53ba145475153b01dc498f85dbd2f03228b](#)

<https://ci.curoverse.com/view/Developer/job/developer-run-tests/850/>

I spent a while mulling over alternate APIs for "shared" but kept coming back the original definition. It is a bit ugly but it should still be much more efficient than listing every project and user and working it out on the client side.

#16 - 08/16/2018 04:18 PM - Tom Clegg

Suggestions

- rename the "include" response field to "included" like <http://jsonapi.org/format/#document-compound-documents>

- ...and rename "includes" to "included" in the comment in shared()
- skip the owners unless the client asks for them with ?include=owner like <http://jsonapi.org/format/#fetching-includes>
- apply @select to the related objects as well

Have you tried these queries on a production-size database to make sure performance is OK with the existing indices?

#18 - 08/16/2018 08:01 PM - Peter Amstutz

Tom Clegg wrote:

Suggestions

- rename the "include" response field to "included" like <http://jsonapi.org/format/#document-compound-documents>
- ...and rename "includes" to "included" in the comment in shared()
- skip the owners unless the client asks for them with ?include=owner like <http://jsonapi.org/format/#fetching-includes>>
- apply @select to the related objects as well

Done.

Have you tried these queries on a production-size database to make sure performance is OK with the existing indices?

See previous note.

Now 13146-shared-rails @ [a41c0f6aa41b658c8f2947c46cb90778894f5cf3](https://ci.curoverse.com/view/Developer/job/developer-run-tests/855/)

<https://ci.curoverse.com/view/Developer/job/developer-run-tests/855/>

#19 - 08/16/2018 08:30 PM - Tom Clegg

Experiments & include/included changes LGTM.

In the docs...

- The second paragraph should have some reference to the "include" flag, perhaps just the word "optionally" -- as written it implies "always".
- It feels like we're saying the same thing 4 different ways, including some pseudocode that has "a and b or c or d", which makes me think about operator precedence, and an "intended use" paragraph that seems to amount to "this API exists to support clients which wish to do what this API does". Could we pare this back to one or two ways of explaining it?

#20 - 08/17/2018 03:24 PM - Peter Amstutz

Tom Clegg wrote:

Experiments & include/included changes LGTM.

In the docs...

- The second paragraph should have some reference to the "include" flag, perhaps just the word "optionally" -- as written it implies "always".
- It feels like we're saying the same thing 4 different ways, including some pseudocode that has "a and b or c or d", which makes me think about operator precedence, and an "intended use" paragraph that seems to amount to "this API exists to support clients which wish to do what this API does". Could we pare this back to one or two ways of explaining it?

Sure, documentation updated @ [489aed58cb0d8bd816e07128cfc9f5a06224083](https://ci.curoverse.com/view/Developer/job/developer-run-tests/855/)

#21 - 08/17/2018 04:59 PM - Tom Clegg

LGTM, thanks.

(Nit: seems like "this is useful for clients which wish to browse the list of" could be replaced with "i.e.," since it's implicit that an API is useful for clients that wish to get the information it returns... but this is gold-plated enough and I'm ready to move on)

#22 - 08/21/2018 03:32 PM - Tom Morris

It looks like a branch got merged a few days ago, but this ticket is still marked as being "In Progress." Is there additional work to be done? We were looking at it during the Workbench2 review to see what its status was.

#23 - 08/21/2018 03:35 PM - Peter Amstutz

- Status changed from In Progress to Resolved

Tom Morris wrote:

It looks like a branch got merged a few days ago, but this ticket is still marked as being "In Progress." Is there additional work to be done? We

were looking at it during the Workbench2 review to see what it's status was.

This is done. The workbench2 folks should be able to use it (provided they are accessing one of the dev clusters with the latest code).

#24 - 08/21/2018 03:36 PM - Peter Amstutz

And it's documented at the bottom of the page here:

<http://doc.arvados.org/api/methods/groups.html>

#25 - 09/12/2018 05:04 PM - Tom Clegg

- Related to Feature #14201: [API] Accept "exclude_home_project" flag in groups#contents added

#26 - 09/17/2018 01:11 PM - Ward Vandewege

- Release set to 13

#27 - 06/26/2019 01:57 PM - Tom Morris

- Related to Support #15238: [Workbench] Change the content inside the owner column from uuid to Projects and user names added