

Arvados - Bug #18316

arv-mount tries to save collection on file close even in read-only mode

11/01/2021 05:44 PM - Peter Amstutz

Status:	Resolved	Start date:	11/12/2021
Priority:	Normal	Due date:	
Assigned To:	Peter Amstutz	% Done:	100%
Category:	FUSE	Estimated time:	0.00 hour
Target version:	2021-11-24 sprint		
Description			
De-identified customer backtrace:			
2021-11-01 15:22:50 arvados.arvados_fuse ¹²⁷³⁰³ ERROR: Error during finalize of inode 13			
Traceback (most recent call last):			
File "/usr/share/python3/dist/python3-arvados-fuse/lib/python3.8/site-packages/arvados_fuse/__init__.py", line 314, in clear_v.finalize()			
File "/usr/share/python3/dist/python3-arvados-fuse/lib/python3.8/site-packages/arvados_fuse/fusedir.py", line 581, in finalize self.collection.save()			
File "/usr/share/python3/dist/python3-arvados-fuse/lib/python3.8/site-packages/arvados/arvfile.py", line 469, in must_be_writable_wrapper			
return orig_func(self, *args, **kwargs)			
File "/usr/share/python3/dist/python3-arvados-fuse/lib/python3.8/site-packages/arvados/arvfile.py", line 270, in synchronized_wrapper			
return orig_func(self, *args, **kwargs)			
File "/usr/share/python3/dist/python3-arvados-fuse/lib/python3.8/site-packages/arvados/retry.py", line 177, in num_retries_setter			
return orig_func(self, *args, **kwargs)			
File "/usr/share/python3/dist/python3-arvados-fuse/lib/python3.8/site-packages/arvados/collection.py", line 1626, in save self._remember_api_response(self._my_api().collections().update(
File "/usr/share/python3/dist/python3-arvados-fuse/lib/python3.8/site-packages/googleapiclient/_helpers.py", line 134, in positional_wrapper			
return wrapped(*args, **kwargs)			
File "/usr/share/python3/dist/python3-arvados-fuse/lib/python3.8/site-packages/googleapiclient/http.py", line 915, in execute raise HttpError(resp, content, uri=self.uri)			
arvados.errors.ApiError: <HttpError 403 when requesting https://.....com/arvados/v1/collections/.....-4zz18-.....?alt=json returned "request failed: http://localhost:8000/arvados/v1/collections/.....-4zz18-..... : 403 Forbidden: Uuid-4zz18-..... is not writable by-tpzed-..... (req-ixqstv61tuf3vnx21lwf)">			
The token doesn't have write access, but it is trying to save the collection anyway. It shouldn't be trying to save the collection when it closes the file.			
There seem to be 2 main bugs here:			
<ul style="list-style-type: none">• Not tracking change operations, tries to saves the collection record on close every time even though nothing changed, which is unnecessary and incorrect• Arv-mount was started in read-only mode but flush() still tries to save the collection.			
Subtasks:			
Task # 18358: Review 18316-fuse-read-only			Resolved

Associated revisions

Revision a4d47d01 - 11/15/2021 05:50 PM - Peter Amstutz

Merge branch '18316-fuse-read-only' refs #18316

Arvados-DCO-1.1-Signed-off-by: Peter Amstutz <peter.amstutz@curii.com>

Revision bd8ee613 - 11/17/2021 07:44 PM - Peter Amstutz

Merge branch '18316-fuse-read-only' refs #18316

Arvados-DCO-1.1-Signed-off-by: Peter Amstutz <peter.amstutz@curii.com>

History

#1 - 11/01/2021 05:56 PM - Peter Amstutz

- Description updated

#2 - 11/10/2021 04:20 PM - Peter Amstutz

- Assigned To set to Peter Amstutz

#3 - 11/12/2021 10:43 PM - Peter Amstutz

18316-fuse-read-only @ [d3ffe252f9d7cbbad9a7bf61ccf5d26129720f43](https://ci.arvados.org/view/Developer/job/developer-run-tests/2795/)

<https://ci.arvados.org/view/Developer/job/developer-run-tests/2795/>

This propagates the enable_write flag to the FUSE directory/file python objects to ensure that when in read-only mode, it won't try to save the collection.

Does *not* fix the related bug that in read-write mode, it still doesn't track if the user has write permission on a collection or not, or if you tried to change the collection or not. So you'll still get error-on-close if you access a file in a collection that was loaded by uuid (via MagicDirectory or ProjectDirectory). This requires a bit more thought on how best to implement properly.

#4 - 11/15/2021 02:53 PM - Peter Amstutz

- Status changed from New to In Progress

#5 - 11/15/2021 04:13 PM - Lucas Di Pentima

Although the updates and test look good, I'm wondering why the original report said that this didn't happen on versions previous to 2.3.0

I've tried to do a manual test against a local arvbox instance, and I'm getting I/O Errors when doing a simple cat <file>. Am I doing something wrong here?

```
(venv-arvados) lucas@buster:~$ arv-mount --version
/home/lucas/venv-arvados/bin/arv-mount 2.4.0.dev20211112222152
(venv-arvados) lucas@buster:~$ arv-mount ./tmp
(venv-arvados) lucas@buster:~$ cd ./tmp/by_id/x240b-4zz18-els01lgw4y50icd
(venv-arvados) lucas@buster:~/tmp/by_id/x240b-4zz18-els01lgw4y50icd$ cat test.js
cat: test.js: Input/output error
```

#6 - 11/15/2021 04:33 PM - Lucas Di Pentima

Same thing with --debug:

```
(venv-arvados) lucas@buster:~$ arv-mount --debug ./tmp
2021-11-15 13:27:53 arvados.arv-mount[24252] DEBUG: arv-mount debugging enabled
2021-11-15 13:27:53 arvados.arv-mount[24252] INFO: /home/lucas/venv-arvados/bin/arv-mount 2.4.0.dev20211112222
152 started
2021-11-15 13:27:53 arvados.arv-mount[24252] INFO: enable write is False
2021-11-15 13:27:54 arvados.arvados_fuse[24252] DEBUG: InodeCache touched inode 5 (size 0) (uuid
x240b-tpzed-k1liuqi1ds4s827) total now 0 (1 entries)
FUSE library version: 2.9.9
(venv-arvados) lucas@buster:~$ cd ./tmp/by_id/x240b-4zz18-els01lgw4y50icd
(venv-arvados) lucas@buster:~/tmp/by_id/x240b-4zz18-els01lgw4y50icd$ cat test.js
cat: test.js: Input/output error
```

#7 - 11/15/2021 04:44 PM - Lucas Di Pentima

Let me re-do my manual test with arvbox in dev mode instead of publicdev because part of the log says:

```
2021-11-15 13:40:10 arvados.keep[14412] DEBUG: Request fail: GET http://localhost:25107/136b64e537cfd8920f0364
6ab852fb61+46+A1975c2037e57581f288f8820b26e005b86e5ac92@61a5022b => <class 'arvados.errors.HttpError': (0, "(
7, 'Failed to connect to localhost port 25107: Connection refused')")
2021-11-15 13:40:10 arvados.arvados_fuse[14412] ERROR: Unhandled exception during FUSE operation
Traceback (most recent call last):
  File "/home/lucas/venv-arvados/lib/python3.7/site-packages/arvados_fuse-2.4.0.dev20211112222152-py3.7.egg/ar
vados_fuse/__init__.py", line 327, in catch_exceptions_wrapper
    return orig_func(self, *args, **kwargs)
  File "/home/lucas/venv-arvados/lib/python3.7/site-packages/arvados_fuse-2.4.0.dev20211112222152-py3.7.egg/ar
vados_fuse/__init__.py", line 655, in read
    r = handle.obj.readfrom(off, size, self.num_retries)
  File "/home/lucas/venv-arvados/lib/python3.7/site-packages/arvados_fuse-2.4.0.dev20211112222152-py3.7.egg/ar
vados_fuse/fusefile.py", line 66, in readfrom
    return self.arvfile.readfrom(off, size, num_retries, exact=True)
  File "/home/lucas/venv-arvados/lib/python3.7/site-packages/arvados/arvfile.py", line 1106, in readfrom
    block = self.parent._my_block_manager().get_block_contents(lr.locator, num_retries=num_retries, cache_only
=(bool(data) and not exact))
```

```
File "/home/lucas/venv-arvados/lib/python3.7/site-packages/arvados/arvfile.py", line 784, in get_block_contents
    return self._keep.get(locator, num_retries=num_retries)
File "/home/lucas/venv-arvados/lib/python3.7/site-packages/arvados/retry.py", line 177, in num_retries_setter
    return orig_func(self, *args, **kwargs)
File "/home/lucas/venv-arvados/lib/python3.7/site-packages/arvados/keep.py", line 1056, in get
    return self._get_or_head(loc_s, method="GET", **kwargs)
File "/home/lucas/venv-arvados/lib/python3.7/site-packages/arvados/keep.py", line 1182, in _get_or_head
    "[{}]] failed to read {} after {}".format(request_id, loc_s, loop.attempts_str()), service_errors, label="service")
arvados.errors.KeepReadError: [req-qgwoe2t8nm779xdla27] failed to read 136b64e537cfd8920f03646ab852fb61+46+A1975c2037e57581f288f8820b26e005b86e5ac92@61a5022b after 5 attempts: service http://localhost:25108/ responded with 0 (7, 'Failed to connect to localhost port 25108: Connection refused'); service http://localhost:25107/ responded with 0 (7, 'Failed to connect to localhost port 25107: Connection refused')
```

Maybe I'm having a keepproxy issue.

#8 - 11/15/2021 04:52 PM - Lucas Di Pentima

It worked on dev mode, LGTM.

#9 - 11/15/2021 05:58 PM - Peter Amstutz

- Status changed from *In Progress* to *Resolved*

#10 - 11/17/2021 07:22 PM - Peter Amstutz

- Release set to 45