

Arvados - Bug #13113

arvados-cwl-runner error while getting output object

02/21/2018 08:33 PM - Joshua Randall

Status:	Closed	Start date:	
Priority:	Normal	Due date:	
Assigned To:	Joshua Randall	% Done:	100%
Category:	Crunch	Estimated time:	0.00 hour
Target version:	2018-04-11 Sprint		
Description			
We are occasionally (when the API server is under heavy load) getting errors such as this from arvados-cwl-runner (a-c-r):			
<pre>2018-02-21T11:49:31.321102281Z arvados.cwl-runner INFO: [container haplotype_caller] reused container ncucu-dz642-f0nyoalc70zb00d 2018-02-21T11:49:37.106085474Z arvados.cwl-runner ERROR: [container haplotype_caller] while getting output object: <HttpError 502 when requesting https://arvados-api-ncucu.hgi.sanger.ac.uk/arvados/v1/collections/b32ff390527c3570482e73bbddf8fe3f%2B307?alt=json returned "Bad Gateway"> 2018-02-21T11:49:37.106085474Z Traceback (most recent call last): 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/arvados_cwl/arvcontainer.py", line 281, in done 2018-02-21T11:49:37.106085474Z outputs = done.done_outputs(self, container, "/tmp", self.outdir, "/keep") 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/arvados_cwl/done.py", line 53, in done_outputs 2018-02-21T11:49:37.106085474Z return self.collect_outputs("keep:" + record["output"]) 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/cwltool/draft2tool.py", line 519, in collect_output_ports 2018-02-21T11:49:37.106085474Z if fs_access.exists(custom_output): 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/arvados_cwl/fsaccess.py", line 123, in exists 2018-02-21T11:49:37.106085474Z collection, rest = self.get_collection(fn) 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/arvados_cwl/fsaccess.py", line 82, in get_collection 2018-02-21T11:49:37.106085474Z return (self.collection_cache.get(pdh), sp[1] if len(sp) == 2 else None) 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/arvados_cwl/fsaccess.py", line 57, in get 2018-02-21T11:49:37.106085474Z keep_client=self.keep_client) 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/arvados/collection.py", line 1662, in __init__ 2018-02-21T11:49:37.106085474Z super(CollectionReader, self).__init__(manifest_locator_or_text, *args, **kwargs) 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/arvados/collection.py", line 1259, in __init__ 2018-02-21T11:49:37.106085474Z self._populate() 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/arvados/collection.py", line 1352, in _populate 2018-02-21T11:49:37.106085474Z self._populate_from_api_server() 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/arvados/collection.py", line 1339, in _populate_from_api_server 2018-02-21T11:49:37.106085474Z num_retries=self.num_retries)) 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/oauth2client/util.py", line 140, in positional_wrapper 2018-02-21T11:49:37.106085474Z return wrapped(*args, **kwargs) 2018-02-21T11:49:37.106085474Z File "/usr/lib/python2.7/dist-packages/googleapiclient/http.py", line 840, in execute 2018-02-21T11:49:37.106085474Z raise HttpError(resp, content, uri=self.uri) 2018-02-21T11:49:37.106085474Z ApiError: <HttpError 502 when requesting https://arvados-api-ncucu.hgi.sanger.ac.uk/arvados/v1/collections/b32ff390527c3570482e73bbddf8fe3f%2B307?alt=json returned "Bad Gateway"></pre>			

It seems that these failures are not being retried by a-c-r -- it looks like a-c-r is not passing `num_retries` when it creates a `CollectionReader` in fsaccess: https://github.com/wtsi-hgi/arvados/blob/master/sdk/cwl/arvados_cwl/fsaccess.py#L56-L57

I will submit a pull request to address this.

Associated revisions

Revision 68cde7e2 - 02/23/2018 06:28 PM - Peter Amstutz

Merge branch 'wtsi/13113-acr-collectioncache-retries' refs #13113

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

History

#1 - 02/21/2018 08:57 PM - Joshua Randall

- % Done changed from 0 to 100

Pull request submitted (this time with a DCO and everything!): <https://github.com/curoverse/arvados/pull/64>

#2 - 04/11/2018 03:58 PM - Tom Morris

- Status changed from New to Closed

- Target version set to 2018-04-11 Sprint

Thanks for the fix!

#3 - 07/23/2018 06:52 PM - Tom Morris

- Release set to 13