

Arvados - Bug #17171

[sdk] unreasonable python SDK timeout when ARVADOS_API_HOST is reachable

11/30/2020 04:24 PM - Ward Vandewege

Status:	Resolved	Start date:	04/15/2021
Priority:	Normal	Due date:	
Assigned To:	Lucas Di Pentima	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	2021-04-28 bughunt sprint		
Description			
The Python SDK seems to take a long time to time out when ARVADOS_API_HOST is unreachable:			
<pre>\$ time arv-user-activity --days 1 Traceback (most recent call last): File "/usr/bin/arv-user-activity", line 8, in <module> sys.exit(main()) File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/arvados_ user_activity/main.py", line 48, in main arv = arvados.api() File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/arvados/ api.py", line 214, in api version=version, cache=cache, request_id=request_id, **kwargs) File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/arvados/ api.py", line 282, in api_from_config return api(version=version, host=host, token=token, insecure=insecure, **kwargs) File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/arvados/ api.py", line 240, in api svc = apiclient_discovery.build('arvados', version, cache_discovery=False, **kwargs) File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/googleap iclient/_helpers.py", line 130, in positional_wrapper return wrapped(*args, **kwargs) File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/googleap iclient/discovery.py", line 222, in build requested_url, discovery_http, cache_discovery, cache) File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/googleap iclient/discovery.py", line 269, in _retrieve_discovery_doc resp, content = http.request(actual_url) File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/arvados/ api.py", line 111, in _intercept_http_request return self.orig_http_request(uri, method, headers=headers, **kwargs) File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/httplib2 /__init__.py", line 1994, in request cachekey, File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/httplib2 /__init__.py", line 1651, in _request conn, request_uri, method, body, headers File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/httplib2 /__init__.py", line 1557, in _conn_request conn.connect() File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/httplib2 /__init__.py", line 1391, in connect raise socket_err File "/usr/share/python3/dist/python3-arvados-user-activity/lib/python3.7/site-packages/httplib2 /__init__.py", line 1324, in connect sock.connect((self.host, self.port)) TimeoutError: [Errno 110] Connection timed out real 6m38.024s user 0m0.366s sys 0m0.024s</pre>			
Subtasks:			

Associated revisions

Revision b82f118f - 04/15/2021 08:03 PM - Lucas Di Pentima

Merge branch '17171-pysdk-timeout-fix'
Closes #17171

Arvados-DCO-1.1-Signed-off-by: Lucas Di Pentima <lucas@di-pentima.com.ar>

History

#1 - 11/30/2020 04:24 PM - Ward Vandewege

- Description updated

#3 - 04/13/2021 05:09 PM - Peter Amstutz

- Target version changed from To Be Groomed to 2021-04-28 bughunt sprint

#4 - 04/14/2021 03:57 PM - Lucas Di Pentima

- Assigned To set to Lucas Di Pentima

#5 - 04/14/2021 06:17 PM - Lucas Di Pentima

- Status changed from New to In Progress

#6 - 04/15/2021 05:13 PM - Lucas Di Pentima

Updates at [80f26d7](#) - branch 17171-pysdk-timeout-fix
Test run: <https://ci.arvados.org/job/developer-run-tests/2416/>

- Fixes the bug by removing the default timeout value (5 minutes) of `arvados.api()`, it's left to the callers the option to set it depending on the different needs.
- Adds test.

#7 - 04/15/2021 06:23 PM - Nico César

review @ [80f26d7ffad85dc15a89bbf6f33a66e206832704](#)

```
time arv-user-activity --days 1
```

takes a lot if None is there.

#8 - 04/15/2021 07:59 PM - Lucas Di Pentima

Thanks for spotting my mistake, I was doing manual tests with low values and assumed setting it to None would be equivalent.

After further investigation I realized that this wasn't a regression from [71e0f70a3095f669e4b136c9d7624c26634db0d3](#) but it was a real bug. AFAICT, when passing None as a timeout to `httplib2`, it uses the socket module's default, that seems to be "no timeout": <https://docs.python.org/3/library/socket.html#socket.getdefaulttimeout>

Updates at [aa8cdf27e](#) sets `arvados.api()` default timeout to 10 seconds. Test updated to check for that.
Test run: <https://ci.arvados.org/job/developer-run-tests/2421/>

#9 - 04/15/2021 08:01 PM - Nico César

review @ [aa8cdf27e282f48981fca4c68b010cd04dc61213](#)

This looks good to me

#10 - 04/15/2021 08:38 PM - Lucas Di Pentima

- Status changed from In Progress to Resolved