

Arvados - Bug #13681

[CWL] Cannot set submit-runner-ram when running from Workbench

06/25/2018 04:56 PM - Peter Amstutz

Status: Resolved	Start date: 07/05/2018
Priority: Normal	Due date:
Assigned To: Peter Amstutz	% Done: 100%
Category:	Estimated time: 0.00 hour
Target version: 2018-07-18 Sprint	
Description Workbench has the RAM allocation for arvados-cwl-runner jobs hardcoded at 256MB. This is way too stingy (when you run from the command line, the default is 1GB) and users with large collection inputs may need as much as 8GB. When running from workbench, there is no way to set --submit-runner-ram. This should be settable, at minimum when using a-c-r --create-workflow. Proposed fix: <ul style="list-style-type: none">• Add a "workflow_runner_ram" column to the Workflows table• When using "a-c-r --create-workflow" or "a-c-r --update-workflow", looks for "--submit-runner-ram" and set "workflow_runner_ram" in the request• When workbench constructs a container request to run a-c-r, use the value of "workflow_runner_ram" for "--submit-runner-ram" on the a-c-r command line.	
Subtasks:	
Task # 13721: a-c-r understands runner ram hint	Resolved
Task # 13722: Workbench understands runner ram hint	Resolved
Task # 13723: Review 13681-wb-acr-submit-ram	Resolved
Task # 13724: Document runner ram hint	Resolved
Related issues:	
Related to Arvados - Bug #13766: [API] bug in libyaml / Pysch used by API ser...	In Progress 07/09/2018

Associated revisions

Revision 77cfd1ee - 07/09/2018 12:48 PM - Peter Amstutz

Merge branch '13681-wb-acr-submit-ram' closes #13681

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

History

#1 - 06/25/2018 04:56 PM - Peter Amstutz

- Status changed from New to In Progress

#2 - 06/25/2018 04:58 PM - Peter Amstutz

- Description updated

#3 - 06/25/2018 05:10 PM - Tom Morris

- Release set to 13

#4 - 06/25/2018 05:11 PM - Peter Amstutz

- Description updated

- Status changed from In Progress to New

- Release deleted (13)

#5 - 06/25/2018 05:12 PM - Peter Amstutz

- Release set to 14

#6 - 06/25/2018 05:12 PM - Peter Amstutz

- Release changed from 14 to 13

#7 - 06/27/2018 01:55 PM - Tom Clegg

Is there an appropriate place to put this information in the workflow definition itself, rather than a separate column?

#8 - 07/02/2018 08:25 PM - Peter Amstutz

Tom Clegg wrote:

Is there an appropriate place to put this information in the workflow definition itself, rather than a separate column?

Yes, I suppose we could add a new hint for it. Workbench would still need to know how to extract it in order to construct the container request.

#9 - 07/03/2018 03:40 PM - Tom Morris

- Target version changed from *To Be Groomed* to *2018-07-18 Sprint*

- Story points set to 1.0

#10 - 07/03/2018 03:44 PM - Peter Amstutz

- Assigned To set to Peter Amstutz

#11 - 07/05/2018 07:19 PM - Peter Amstutz

13681-wb-acr-submit-ram @ [4e714090c2c81cddcef6da393b5680f3fcfe40b1](https://ci.curoverse.com/job/developer-run-tests/784/)

<https://ci.curoverse.com/job/developer-run-tests/784/>

- New hint arv:WorkflowRunnerResources
- Recognized by a-c-r Runner class and Workbench to set runtime_constraints for "ram" and "vcpus"
- Providing --submit-runner-ram with --create/update-workflow will implicitly set arv:WorkflowRunnerResources in hints
- Updated documentation

#12 - 07/06/2018 01:42 PM - Lucas Di Pentima

Reviewing [334c069f4da11278844f4276b5c254e8cbf06412](https://ci.curoverse.com/job/developer-run-tests/784/)

- There're some sdk/cwl test failures on jenkins (already talked about this, I'm just recording the observation)
- Just some code style observations on file sdk/cwl/arvados_cwl/arvworkflow.py
 - Line 55: Could that be replaced with this?: hints = main.get('hints', [])
 - The found var on line 57 could be replaced with an else clause on the for loop?

#13 - 07/06/2018 05:52 PM - Peter Amstutz

Lucas Di Pentima wrote:

Reviewing [334c069f4da11278844f4276b5c254e8cbf06412](https://ci.curoverse.com/job/developer-run-tests/784/)

- There're some sdk/cwl test failures on jenkins (already talked about this, I'm just recording the observation)

Yea, I'm going to revert the change which is causing those failures, I filed [#13766](#) to dig deeper.

- Just some code style observations on file sdk/cwl/arvados_cwl/arvworkflow.py
 - Line 55: Could that be replaced with this?: hints = main.get('hints', [])

Yes, I don't know why I wrote it that way.

- The found var on line 57 could be replaced with an else clause on the for loop?

No, because (according to my research) the "else" clause doesn't fire for empty arrays.

#14 - 07/06/2018 05:53 PM - Peter Amstutz

13681-wb-acr-submit-ram @ [7cebf3bba60a98897ebddf5e613db6337e5f1513](https://ci.curoverse.com/job/developer-run-tests/786/)

<https://ci.curoverse.com/job/developer-run-tests/786/>

#15 - 07/06/2018 05:53 PM - Peter Amstutz

- Related to Bug [#13766](#): [API] bug in libyaml / Pysch used by API server to parse yaml added

#16 - 07/06/2018 06:07 PM - Lucas Di Pentima

Providing that tests pass OK, LGTM.

Regarding the else clause, it seems that they work even with empty lists:

```
Python 2.7.13 (default, Jun 27 2017, 10:00:24)
[GCC 4.2.1 Compatible Apple LLVM 8.1.0 (clang-802.0.42)] on darwin
Type "help", "copyright", "credits" or "license" for more information.
>>> for a in []:
...     print "loop"
... else:
...     print "else"
...
else
```

#17 - 07/09/2018 12:55 PM - Peter Amstutz

- *Status changed from New to Resolved*

- *% Done changed from 75 to 100*

Applied in changeset [arvados|77cfd1ee241b4a2d408bd3ec5ebc69f34c7496f5](#).