

## Pulp - Issue #1669

### Memory leak in Pulp celery processes

02/11/2016 09:49 PM - bmbouter

<b>Status:</b>	CLOSED - WORKSFORME	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>		<b>Estimated time:</b>	0:00 hour
<b>Category:</b>		<b>Groomed:</b>	No
<b>Sprint/Milestone:</b>		<b>Sprint Candidate:</b>	No
<b>Severity:</b>	2. Medium	<b>Tags:</b>	Pulp 2
<b>Version:</b>		<b>Sprint:</b>	
<b>Platform Release:</b>	2.8.1	<b>Quarter:</b>	
<b>OS:</b>			
<b>Triaged:</b>	Yes		

#### Description

Via IRC I received a report that on at least two occasions, Pulp workers of a large Pulp installation showed signs of memory leak. The leak consumed so much memory that at a certain point, all subsequent tasks fail with "Cannot allocate memory".

When they are in this situation there are no tasks waiting or running in Pulp as is evident by their output provided for:

No running or waiting tasks in pulp

```
pulp:PRIMARY> db.task_status.find({"state":{"$in":["running", "waiting"]}})
pulp:PRIMARY>
```

Also you can clearly see the offending processes are Pulp celery processes which are idle but consuming large amounts of memory.

```
bash-4.1# ps u 16012
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
apache   16012  0.2 15.1 3016808 2475812 ?        S1   Feb10    4:17 /usr/bin/python -m celery.__main_
_ worker -c 1 -n reserved_resource_worker-1@pulp04.example.com --events --a
bash-4.1# ps u 16121
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
apache   16121  0.3 13.6 2767072 2224064 ?        S1   Feb10    5:24 /usr/bin/python -m celery.__main_
_ worker -c 1 -n reserved_resource_worker-5@pulp04.example.com --events --a
bash-4.1# ps u 16206
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
apache   16206  0.2 12.8 3092120 2105596 ?        S1   Feb10    3:34 /usr/bin/python -m celery.__main_
_ worker -c 1 -n reserved_resource_worker-7@pulp04.example.com --events --a
bash-4.1# ps u 16141
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
apache   16141  0.3 12.5 2621912 2056276 ?        S1   Feb10    4:31 /usr/bin/python -m celery.__main_
_ worker -c 1 -n reserved_resource_worker-6@pulp04.example.com --events --a
```

It's not yet clear if this is happening over time versus instantaneously, or if it is related to a specific task type. It has been suggested that this may be related[0]. We should also consider that this could be an upstream Celery issue also which has several memory leak related bug reports[1].

[0]: <https://dzone.com/articles/diagnosing-memory-leaks-python>

[1]: <https://github.com/celery/celery/issues?utf8=%E2%9C%93&q=is%3Aissue+is%3Aopen+leak>

#### History

#1 - 02/12/2016 05:13 PM - bmbouter

- Platform Release set to 2.8.1

- Triaged changed from No to Yes

#2 - 02/25/2016 04:10 PM - bmbouter

- Status changed from NEW to CLOSED - WORKSFORME

I was contacted by the original reporter and they identified the issue as not being a root cause in Pulp code. They had made some modifications to the code which introduced a memory issue. They are not exactly sure of the root cause, but by modifying their implementation they confirmed it was resolved.

If other users observe this, please re-open.

**#3 - 04/15/2019 10:35 PM - bmbouter**

- Tags Pulp 2 added