Pulp - Issue #2331

The "unsafe_autoretry" feature does not retry with mongoengine.QuerySet codepaths

10/11/2016 08:13 PM - mhrivnak

Status:	CLOSED - WONTFIX	**Start date:**	
Priority:	High	**Due date:**	
Assignee:	daviddavis	**Estimated time:**	0:00 hour
Category:			
Sprint/Milestone:			
Severity:	2. Medium	**Groomed:**	No
Version:	2.8.7	**Sprint Candidate:**	No
Platform Release:		**Tags:**	Pulp 2
OS:	CentOS 7	**Sprint:**	Sprint 11
Triaged:	Yes	**Quarter:**	

Description

If on a running pulp deployment you stop mongod, pulp_celerybeat will eventually die and produce the following log output, including a traceback.

Instead, it should handle the connection failure gracefully and recover when mongod becomes available again.

This behavior has been seen "in the wild" when a full disk caused mongod to die.

```
Oct 11 14:04:09 localhost.localdomain pulp[12792]: pulp.server.async.worker_watcher:DEBUG: 'worker
-heartbeat' sent at time 2016-10-11 19:04:09.617800 from resource_manager@localhost.localdomain, r
eceived at time: 2016-10-11 18:04:09.635230
Oct 11 14:04:39 localhost.localdomain pulp[12792]: pulp.server.async.scheduler:ERROR: localhost:27
017: [Errno 111] Connection refused
Oct 11 14:04:47 localhost.localdomain pulp[12792]: pulp.server.async.scheduler:DEBUG: Checking if
pulp_workers, pulp_celerybeat, or pulp_resource_manager processes are missing for more than 300 se
conds
Oct 11 14:04:49 localhost.localdomain pulp[12792]: kombu.transport.qpid:DEBUG: Attempting to conne
ct to qpid with SASL mechanism ANONYMOUS
Oct 11 14:04:49 localhost.localdomain pulp[12792]: kombu.transport.qpid:INFO: Connected to qpid wi
th SASL mechanism ANONYMOUS
Oct 11 14:04:49 localhost.localdomain pulp[12792]: pulp.server.async.scheduler:INFO: Event Monitor
 Starting
Oct 11 14:04:49 localhost.localdomain pulp[12792]: kombu.transport.qpid:DEBUG: Attempting to conne
ct to qpid with SASL mechanism ANONYMOUS
Oct 11 14:04:49 localhost.localdomain pulp[12792]: kombu.transport.qpid:INFO: Connected to qpid wi
th SASL mechanism ANONYMOUS
Oct 11 14:04:49 localhost.localdomain pulp[12792]: kombu.mixins:INFO: Connected to qpid://localhos
t:5672//
Oct 11 14:04:49 localhost.localdomain pulp[12792]: pulp.server.async.worker_watcher:DEBUG: 'worker
-heartbeat' sent at time 2016-10-11 19:04:49.940497 from resource_manager@localhost.localdomain, r
eceived at time: 2016-10-11 18:04:49.961244
Oct 11 14:05:17 localhost.localdomain pulp[12792]: pulp.server.async.scheduler:ERROR: localhost:27
017: [Errno 111] Connection refused
Oct 11 14:05:18 localhost.localdomain pulp[12792]: pulp.server.async.worker_watcher:DEBUG: 'schedu
ler-event' sent at time 2016-10-11 18:05:18.236002 from scheduler@localhost.localdomain, received
at time: 2016-10-11 18:05:18.236004
Oct 11 14:05:20 localhost.localdomain pulp[12792]: pulp.server.async.scheduler:ERROR: localhost:27
017: [Errno 111] Connection refused
Oct 11 14:05:30 localhost.localdomain pulp[12792]: kombu.transport.qpid:DEBUG: Attempting to conne
ct to qpid with SASL mechanism ANONYMOUS
Oct 11 14:05:30 localhost.localdomain pulp[12792]: kombu.transport.qpid:INFO: Connected to qpid wi
th SASL mechanism ANONYMOUS
Oct 11 14:05:30 localhost.localdomain pulp[12792]: pulp.server.async.scheduler:INFO: Event Monitor
 Starting
Oct 11 14:05:30 localhost.localdomain pulp[12792]: kombu.transport.qpid:DEBUG: Attempting to conne
```

```
ct to qpid with SASL mechanism ANONYMOUS
Oct 11 14:05:30 localhost.localdomain pulp[12792]: kombu.transport.qpid:INFO: Connected to qpid wi
th SASL mechanism ANONYMOUS
Oct 11 14:05:30 localhost.localdomain pulp[12792]: kombu.mixins:INFO: Connected to qpid://localhos
t:5672//
Oct 11 14:05:30 localhost.localdomain pulp[12792]: pulp.server.async.worker_watcher:DEBUG: 'worker
-heartbeat' sent at time 2016-10-11 19:05:30.561665 from reserved_resource_worker-0@localhost.loca
ldomain, received at time: 2016-10-11 18:05:30.579888
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504) beat raised
 exception <class 'pymongo.errors.ServerSelectionTimeoutError'>: ServerSelectionTimeoutError('loca
lhost:27017: [Errno 111] Connection refused',)
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504) Traceback (
most recent call last):
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib/python2.7/site-packages/celery/apps/beat.py", line 112, in start_scheduler
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     beat.st
art()
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib/python2.7/site-packages/celery/beat.py", line 462, in start
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     interva
l = self.scheduler.tick()
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib/python2.7/site-packages/pulp/server/async/scheduler.py", line 252, in tick
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     worker_
watcher.handle_worker_heartbeat(scheduler_event)
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib/python2.7/site-packages/pulp/server/async/worker_watcher.py", line 75, in handle_worker_hear
tbeat
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     worker
= Worker.objects(name=event_info['worker_name']).first()
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib/python2.7/site-packages/mongoengine/queryset/base.py", line 264, in first
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     result
= queryset[0]
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib/python2.7/site-packages/mongoengine/queryset/base.py", line 161, in __getitem__
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     return
queryset._document._from_son(queryset._cursor[key],
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib64/python2.7/site-packages/pymongo/cursor.py", line 565, in __getitem__
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     for doc
 in clone:
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib64/python2.7/site-packages/pymongo/cursor.py", line 1097, in next
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     if len(
self.__data) or self._refresh():
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib64/python2.7/site-packages/pymongo/cursor.py", line 1019, in _refresh
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     self.__
read_concern))
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib64/python2.7/site-packages/pymongo/cursor.py", line 850, in __send_message
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     **kwarg
s)
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib64/python2.7/site-packages/pymongo/mongo_client.py", line 777, in _send_message_with_response
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     server
= topology.select_server(selector)
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib64/python2.7/site-packages/pymongo/topology.py", line 141, in select_server
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     address
))
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)   File "/us
r/lib64/python2.7/site-packages/pymongo/topology.py", line 117, in select_servers
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504)     self._e
rror_message(selector))
Oct 11 14:05:48 localhost.localdomain pulp[12792]: celery.beat:CRITICAL: (12792-69504) ServerSelec
```

```
tionTimeoutError: localhost:27017: [Errno 111] Connection refused
```

Related issues:

Related to Pulp - Issue #2423: As a user, I want pulp_celerybeat to survive a...	**CLOSED - WORKSFORME**
Related to Pulp - Task #2417: Ensure all processes have initial and reconnect...	**CLOSED - CURRENTRELEASE**

History

#2 - 10/11/2016 09:04 PM - mmccune@redhat.com

- Severity changed from 2. Medium to 3. High

I'd like to bump the severity to High as it isn't obvious from the user perspective of the interaction between the two. Generally we would always recommend a full restart of Mongo goes down but it would be much more resilient if services could handle outages better

#3 - 10/14/2016 04:34 PM - amacdona@redhat.com

- Priority changed from Normal to High

- Sprint/Milestone set to 27

- Severity changed from 3. High to 2. Medium

- Triaged changed from No to Yes

#4 - 10/24/2016 04:29 PM - jortel@redhat.com

- Sprint/Milestone changed from 27 to 28

#5 - 11/07/2016 08:15 PM - daviddavis

- Assignee set to daviddavis

#6 - 11/08/2016 01:06 PM - daviddavis

- Status changed from NEW to ASSIGNED

#7 - 11/08/2016 04:08 PM - bmbouter

I suspect it's actually all of the Pulp processes which are affected not just celerybeat. Confirming that and updating the issue would be good.

Quick background on the tasking system's usage of mongodb. There are really 3 entry points for Pulp. One is httpd itself, the second is celerybeat which uses scheduler.py and the third are the workers (both pulp_resource_manager and pulp_workers) share this one. All of them use a shared section of code to acquire the connection[0]. We manage the connection and the same code is used by all of our processes so that is why I think this issue affects all of our processes.

The idea of putting a few try/except in key places around the code would help, but it will be tough to get right with that approach for two reasons. (1) The exception could be raised at any point Pulp code calls into mongodb which is all over the place and (2) We want to retry-wait-and-continue at the moment the exception is raised so letting the exception propogate all the way up the stack would prevent the ability to resume at the point in the call stack.

We've had a similar problem before which caused us to "wrap" all the meaningful pyMongo methods to catch-and-retry. We do that here[1] with one of our features aptly named unsafe-retry. I think we want to do something similar here and put the try/except there and introduce a backoff and reconnect behavior with this exception type. Unlike the unsafe-retry feature which is turned on via a config in server.conf I think this backoff and retry would be enabled by default. See how the traceback flows through methods that are decorated by this section of code[1]. I think that is your opportunity.

[0]: https://github.com/pulp/pulp/blob/master/server/pulp/server/db/connection.py#L35-L168
[1]: https://github.com/pulp/pulp/blob/master/server/pulp/server/db/connection.py#L203-L260

#8 - 11/10/2016 12:39 AM - daviddavis

Digging into this bug, it looks like the places where we are failing with ServerSelectionTimeoutError, our retry code is not decorating the methods we need. There's two places we call decorate_instance, once for Documents[0] and once for pymongo Collections[1].

I tried adding 'next', 'first', etc to our list of methods to decorate but the object we're dealing with in this bug is a QuerySet. You can see that in the stacktrace and also you can see that first() calls pymongo's cursor to get the first item. This of course bypasses our retry code.

I debugged to confirm this suspicion and also noticed that there are other places that have the same problem (e.g. [2]). I was able to patch our QuerySet class to fix the problem[3] but that doesn't seem like an optimal solution in that we'd have to add/maintain yet another class from mongoengine/pymongo, and also, we'd have to go through QuerySet and add any methods that access the db to our list of decorated methods.

Not sure if there's perhaps some other solution to consider?

[0] https://git.io/vXuNj
[1] https://git.io/vX2FE
[2] https://git.io/vXu7c
[3] https://gist.github.com/daviddavis/58c1d81e3b2ee259667fb26d2eab631c

#9 - 11/10/2016 10:31 PM - bmbouter

- Subject changed from pulp_celerybeat dies if mongod stops to The "unsafe_autoretry" feature does not retry with mongoengine.QuerySet codepaths

daviddavis, great job on debugging all of this.

I retitled this bug to reflect what I've learned from your investigation. The unsafe_autoretry feature is supposed to catch-and-retry AutoReconnect exceptions. The ServerSelectionTimeoutError[0] exception is a subclass of AutoReconnect so that feature is designed to handle this case already. The unsafe_autoretry feature does not correctly catch-and-retry exceptions raised while calling into mongoengine.QuerySet objects.

This won't be an issue in Pulp3 because we won't be using mongodb at all. For the 2.y line, we have 2 options.

1. To fix this we would need to extend the unsafe_autoretry to cover QuerySet objects also.
2. The other option is to close it as WONTFIX since the issue will never go beyond the 2.y line

#10 - 11/14/2016 11:10 PM - mhrivnak

- Sprint/Milestone changed from 28 to 29

#11 - 11/14/2016 11:55 PM - daviddavis

- Related to Issue #2423: As a user, I want pulp_celerybeat to survive a brief database outage without having to use the unsafe_autoretry feature. added

#12 - 11/14/2016 11:56 PM - daviddavis

Opened up a related issue to look into having pulp_celerybeat not die when we can't connect to MongoDB:

#13 - 11/16/2016 05:15 PM - mhrivnak

The unsafe_autoretry feature is not widely used, and there does not seem to be broad demand for this particular behavior. Let's close this and let pulp 3 make it irrelevant.

#14 - 11/17/2016 04:14 PM - daviddavis

- Status changed from ASSIGNED to CLOSED - WONTFIX

#15 - 11/17/2016 04:14 PM - daviddavis

- Related to Task #2417: Ensure all processes have initial and reconnect support for PostgreSQL added

#16 - 03/08/2018 08:15 PM - bmbouter

- Sprint set to Sprint 11

#17 - 03/08/2018 08:17 PM - bmbouter

- Sprint/Milestone deleted (29)

#18 - 04/15/2019 10:24 PM - bmbouter

- Tags Pulp 2 added