

## Pulp - Issue #5906

### tasks can be assigned to dead workers

12/19/2019 09:41 PM - jsherril@redhat.com

|                          |                     |                          |           |
|--------------------------|---------------------|--------------------------|-----------|
| <b>Status:</b>           | CLOSED - WORKSFORME | <b>Start date:</b>       |           |
| <b>Priority:</b>         | Normal              | <b>Due date:</b>         |           |
| <b>Assignee:</b>         | bmbouter            | <b>Estimated time:</b>   | 0:00 hour |
| <b>Category:</b>         |                     |                          |           |
| <b>Sprint/Milestone:</b> |                     |                          |           |
| <b>Severity:</b>         | 2. Medium           | <b>Groomed:</b>          | No        |
| <b>Version:</b>          |                     | <b>Sprint Candidate:</b> | No        |
| <b>Platform Release:</b> |                     | <b>Tags:</b>             |           |
| <b>OS:</b>               |                     | <b>Sprint:</b>           | Sprint 65 |
| <b>Triaged:</b>          | Yes                 | <b>Quarter:</b>          |           |

#### Description

I hit a situation with pulp3 where i synced a repository and the task never got picked up. After restarting ALL pulp3 services, i was still seeing this behavior. Other repositories synced without issue.

Turns out, all these stuck syncs are being assigned to the same worker that no longer existed:

```
- pulp_href: "/pulp/api/v3/tasks/04b2f793-ef3e-49fa-8698-b3f86654d48e/"
pulp_created: '2019-12-19T20:28:02.060+00:00'
state: waiting
name: pulp_container.app.tasks.synchronize.synchronize
worker: "/pulp/api/v3/workers/3880936c-1a64-46ff-9651-7fefef088c69/"
progress_reports: []
created_resources: []
reserved_resources_record:
- "/pulp/api/v3/remotes/container/container/366ec86d-5910-499f-ab51-07bc73cafb80/"
- "/pulp/api/v3/repositories/container/container/1026248f-1e62-497d-9dee-d8c959dea8d4/"
```

#### status output:

```
{
  "versions": [
    {
      "component": "pulpcore",
      "version": "3.0.0rc9"
    },
    {
      "component": "pulp_2to3_migration",
      "version": "0.0.1a1.dev0"
    },
    {
      "component": "pulp_file",
      "version": "0.1.0.dev0"
    },
    {
      "component": "pulp_container",
      "version": "1.0.0rc2.dev0"
    },
    {
      "component": "pulp_ansible",
      "version": "0.2.0b7.dev0"
    },
    {
      "component": "pulp_rpm",
      "version": "3.0.0rc2.dev0"
    }
  ]
}
```

```

],
"online_workers": [
  {
    "pulp_href": "/pulp/api/v3/workers/7f91ec97-98e3-47f2-a1d8-10e0f2bf13ec/",
    "pulp_created": "2019-12-19T20:27:44.849421Z",
    "name": "23430@devel.balmora.example.com",
    "last_heartbeat": "2019-12-19T20:38:33.138047Z"
  },
  {
    "pulp_href": "/pulp/api/v3/workers/19bf0a63-d8c9-4f9b-b7e2-4597c903c31c/",
    "pulp_created": "2019-12-19T20:27:45.006246Z",
    "name": "23429@devel.balmora.example.com",
    "last_heartbeat": "2019-12-19T20:38:33.138315Z"
  },
  {
    "pulp_href": "/pulp/api/v3/workers/68acaa8e-5b2e-4fb5-b17a-f903171691db/",
    "pulp_created": "2019-12-11T17:10:07.192264Z",
    "name": "resource-manager",
    "last_heartbeat": "2019-12-19T20:38:35.546686Z"
  }
],
"online_content_apps": [
  {
    "name": "23444@devel.balmora.example.com",
    "last_heartbeat": "2019-12-19T20:38:31.081544Z"
  },
  {
    "name": "23442@devel.balmora.example.com",
    "last_heartbeat": "2019-12-19T20:38:31.088967Z"
  }
],
"database_connection": {
  "connected": true
},
"redis_connection": {
  "connected": true
},
"storage": {
  "total": 42927656960,
  "used": 23027937280,
  "free": 19899719680
}
}

```

#### Dead worker details:

```

$ curl https://`hostname`/pulp/api/v3/workers/3880936c-1a64-46ff-9651-7fefef088c69/ -u admin:password | jq
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload  Total   Spent    Left  Speed
100  206  100    206    0     0    655      0  --:--:--  --:--:--  --:--:--   658
{
  "pulp_href": "/pulp/api/v3/workers/3880936c-1a64-46ff-9651-7fefef088c69/",
  "pulp_created": "2019-12-16T15:04:54.639299Z",
  "name": "1989@devel.balmora.example.com",
  "last_heartbeat": "2019-12-19T19:53:36.786369Z"
}

```

## History

### #1 - 12/19/2019 10:06 PM - jsherril@redhat.com

I fetched all the tasks associated with those reserved resources:

```

$ curl https://`hostname`/pulp/api/v3/tasks/?reserved_resources_record=/pulp/api/v3/remotes/container/containers/366ec86d-5910-499f-ab51-07bc73cafb80/ -u admin:password |jq
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload  Total   Spent    Left  Speed

```

```

{
  "count": 4,
  "next": null,
  "previous": null,
  "results": [
    {
      "pulp_href": "/pulp/api/v3/tasks/04b2f793-ef3e-49fa-8698-b3f86654d48e/",
      "pulp_created": "2019-12-19T20:28:02.060726Z",
      "state": "waiting",
      "name": "pulp_container.app.tasks.synchronize.synchronize",
      "started_at": null,
      "finished_at": null,
      "error": null,
      "worker": "/pulp/api/v3/workers/3880936c-1a64-46ff-9651-7fefef088c69/",
      "progress_reports": [],
      "created_resources": [],
      "reserved_resources_record": [
        "/pulp/api/v3/remotes/container/container/366ec86d-5910-499f-ab51-07bc73cafb80/",
        "/pulp/api/v3/repositories/container/container/1026248f-1e62-497d-9dee-d8c959dea8d4/"
      ]
    },
    {
      "pulp_href": "/pulp/api/v3/tasks/92bbd7a5-8714-424d-9eef-2245e5cd4ad6/",
      "pulp_created": "2019-12-19T20:18:57.083880Z",
      "state": "canceled",
      "name": "pulp_container.app.tasks.synchronize.synchronize",
      "started_at": null,
      "finished_at": null,
      "error": null,
      "worker": "/pulp/api/v3/workers/3880936c-1a64-46ff-9651-7fefef088c69/",
      "progress_reports": [],
      "created_resources": [],
      "reserved_resources_record": [
        "/pulp/api/v3/repositories/container/container/1026248f-1e62-497d-9dee-d8c959dea8d4/",
        "/pulp/api/v3/remotes/container/container/366ec86d-5910-499f-ab51-07bc73cafb80/"
      ]
    },
    {
      "pulp_href": "/pulp/api/v3/tasks/8b2a3659-396e-4f25-8bb3-bc5f1b02ef18/",
      "pulp_created": "2019-12-19T20:05:18.591109Z",
      "state": "waiting",
      "name": "pulp_container.app.tasks.synchronize.synchronize",
      "started_at": null,
      "finished_at": null,
      "error": null,
      "worker": "/pulp/api/v3/workers/3880936c-1a64-46ff-9651-7fefef088c69/",
      "progress_reports": [],
      "created_resources": [],
      "reserved_resources_record": [
        "/pulp/api/v3/repositories/container/container/1026248f-1e62-497d-9dee-d8c959dea8d4/",
        "/pulp/api/v3/remotes/container/container/366ec86d-5910-499f-ab51-07bc73cafb80/"
      ]
    },
    {
      "pulp_href": "/pulp/api/v3/tasks/ea008903-fac7-46d1-ae76-7cac037b87b6/",
      "pulp_created": "2019-12-19T19:50:00.190200Z",
      "state": "failed",
      "name": "pulp_container.app.tasks.synchronize.synchronize",
      "started_at": "2019-12-19T19:50:00.357552Z",
      "finished_at": "2019-12-19T19:53:36.773635Z",
      "error": {
        "traceback": "",
        "description": "None"
      },
      "worker": "/pulp/api/v3/workers/3880936c-1a64-46ff-9651-7fefef088c69/",
      "progress_reports": [
        {
          "message": "Downloading tag list",
          "code": "downloading.tag_list",
          "state": "running",
          "total": 1,
          "done": 0,
          "suffix": null
        }
      ]
    }
  ]
}

```

```

    {
      "message": "Downloading Artifacts",
      "code": "downloading.artifacts",
      "state": "running",
      "total": null,
      "done": 0,
      "suffix": null
    },
    {
      "message": "Associating Content",
      "code": "associating.content",
      "state": "running",
      "total": null,
      "done": 0,
      "suffix": null
    }
  ],
  "created_resources": [
    null
  ],
  "reserved_resources_record": [
    "/pulp/api/v3/repositories/container/container/1026248f-1e62-497d-9dee-d8c959dea8d4/",
    "/pulp/api/v3/remotes/container/container/366ec86d-5910-499f-ab51-07bc73cafb80/"
  ]
}
]
}

```

## #2 - 12/19/2019 10:12 PM - jsherril@redhat.com

It looks like i might have shut down the workers around the time that initial sync failed:

```

Dec 19 19:53:36 devel.balmora.example.com systemd[1]: Stopping Pulp RQ Worker...
Dec 19 19:53:36 devel.balmora.example.com systemd[1]: Stopping Pulp RQ Worker...
Dec 19 19:53:37 devel.balmora.example.com systemd[1]: pulpcore-worker@2.service: main process exited, code=exi
ted, status=120/n/a
Dec 19 19:53:37 devel.balmora.example.com systemd[1]: Stopped Pulp RQ Worker.
Dec 19 19:53:37 devel.balmora.example.com systemd[1]: Unit pulpcore-worker@2.service entered failed state.
Dec 19 19:53:37 devel.balmora.example.com systemd[1]: pulpcore-worker@2.service failed.
Dec 19 19:53:37 devel.balmora.example.com systemd[1]: Started Pulp RQ Worker.
Dec 19 19:53:37 devel.balmora.example.com systemd[1]: pulpcore-worker@1.service: main process exited, code=exi
ted, status=120/n/a
Dec 19 19:53:37 devel.balmora.example.com systemd[1]: Stopped Pulp RQ Worker.
Dec 19 19:53:37 devel.balmora.example.com systemd[1]: Unit pulpcore-worker@1.service entered failed state.
Dec 19 19:53:37 devel.balmora.example.com systemd[1]: pulpcore-worker@1.service failed.
Dec 19 19:53:37 devel.balmora.example.com systemd[1]: Started Pulp RQ Worker.

```

## #3 - 12/20/2019 12:10 PM - meldaresearch

- File 13b.jpg added

## #4 - 12/20/2019 04:34 PM - fao89

- Triaged changed from No to Yes

- Sprint set to Sprint 63

## #5 - 12/20/2019 04:35 PM - fao89

- File deleted (13b.jpg)

## #6 - 01/15/2020 10:02 PM - rchan

- Sprint changed from Sprint 63 to Sprint 64

## #7 - 01/24/2020 02:58 PM - rchan

- Sprint changed from Sprint 64 to Sprint 65

## #8 - 01/27/2020 08:41 PM - bmbouter

@jsherrill what do you think I can do to reproduce this?

## #9 - 01/27/2020 08:42 PM - bmbouter

- Status changed from NEW to ASSIGNED

- Assignee set to bmbouter

#### #10 - 01/27/2020 09:19 PM - jsherril@redhat.com

You can try:

1. starting a sync
2. systemctl restart 'pulpcore-workers@\*'
3. Start a new sync

but i have no idea how reproducible this is. I only saw it on the one installation (which is why i immediately came on #pulp-dev to ask for what info was needed). As soon as i reset my environment i haven't seen it again.

Here's some info from the chat:

```
jsherrill      i searched for tasks with the same reservation, and there was 2 waiting, 1 cancelled (that i
cancelled), and 1 'failed'
jsherrill      looks like the failed one was maybe running at some point
jsherrill      it has 'running' progress reports
jsherrill      i even see this in my logs:
jsherrill      Dec 16 19:11:25 devel rq: pulp: rq.worker:INFO: Cleaning registries for queue: 1989@devel.bal
mora.example.com
jsherrill      which seems to be the worker in question
jsherrill      but i guess thats more an rq thing
dkliban jsherrill: i think you are onto something
jsherrill      let me put the task list in the issue too
jsherrill      https://pulp.plan.io/issues/5906#note-1
jsherrill      that last task is the one that 'died' or 'failed'
jsherrill      and i'm not sure what put it in that state
jsherrill      let me doublecheck journalctl
jsherrill      it looks like maybe i had shut down the workers at that 19:53
```

#### #11 - 01/28/2020 04:02 PM - bmbouter

- Status changed from ASSIGNED to CLOSED - WORKSFORME

I added a 10-second sleep statement to the sync task to slow it down. Then I reset my pulp environment and ran a script that:

makes a repo  
makes a remote  
sync's the remote from the remote to successfully bring in a few content units (from fedorapeople)

Then I took my Pulp system down to 1 worker. Then I start the sync and first I used systemctl to restart the worker. Later I would do this again, only I would pkill -9 -f pulpcore-worker which did successfully kill the worker. In both cases the task was shown as 'canceled' as follows:

```
{
  "count": 1,
  "next": null,
  "previous": null,
  "results": [
    {
      "pulp_href": "/pulp/api/v3/tasks/7718277d-33f4-41da-9933-6a2e5bfebe05/",
      "pulp_created": "2020-01-28T14:27:20.001474Z",
      "state": "canceled",
      "name": "pulp_file.app.tasks.synchronizing.synchronize",
      "started_at": "2020-01-28T14:27:20.122284Z",
      "finished_at": null,
      "error": null,
      "worker": "/pulp/api/v3/workers/770e12fb-6a8f-4b03-9de2-24ae24f000b7/",
      "progress_reports": [],
      "created_resources": [],
      "reserved_resources_record": [
        "/pulp/api/v3/repositories/file/file/23027b38-2770-4144-bce3-b2edf189d71c/",
        "/pulp/api/v3/remotes/file/file/a6f3e605-79a5-4b10-9a74-0c1a78578a6c/"
      ]
    }
  ]
}
```

In the case of prestart the worker gracefully canceled it's own work upon shutdown (as expected). In the case of pkill, the resource manager cleaned up and logged that the worker went missing (as expected).

Then after restarting the worker I performed more sync's and those were always scheduled on the current worker.