

## Pulp - Issue #8978

### Users are reporting timeout issues during sync

06/28/2021 09:38 AM - quba42

<b>Status:</b>	CLOSED - DUPLICATE	<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>	
<b>Version:</b>		<b>Sprint:</b>	
<b>Platform Release:</b>		<b>Quarter:</b>	
<b>OS:</b>			
<b>Triaged:</b>	No		

#### Description

I am uncertain if this is a bug in pulpcore or pulp\_deb, or if it can be remedied by mere configuration.

But users are frequently reporting that large syncs (think initial syncs of Debian main) run into timeouts when creating the new DebDeclarativeVersion. Apparently this can then be worked around by re-syncing several times until the sync eventually completes. Presumably the later syncs don't need to re-download many packages and therefore avoid the timeout.

See the following foreman community discussion: <https://community.theforeman.org/t/add-debian-10-repository/24157/13>

See the following backtrace from a Katello user:

```
{ "traceback"=>
  " File \"/usr/lib/python3.6/site-packages/rq/worker.py", line 975, in perform_job\n" +
  "   rv = job.perform()\n" +
  " File \"/usr/lib/python3.6/site-packages/rq/job.py", line 696, in perform\n" +
  "   self._result = self._execute()\n" +
  " File \"/usr/lib/python3.6/site-packages/rq/job.py", line 719, in _execute\n" +
  "   return self.func(*self.args, **self.kwargs)\n" +
  " File \"/usr/lib/python3.6/site-packages/pulp_deb/app/tasks/synchronizing.py", line 122,
in synchronize\n" +
  "   DebDeclarativeVersion(first_stage, repository, mirror=mirror).create()\n" +
  " File \"/usr/lib/python3.6/site-packages/pulpcore/plugin/stages/declarative_version.py",
line 147, in create\n" +
  "   loop.run_until_complete(pipeline)\n" +
  " File \"/usr/lib64/python3.6/asyncio/base_events.py", line 484, in run_until_complete\n"
+
  "   return future.result()\n" +
  " File \"/usr/lib/python3.6/site-packages/pulpcore/plugin/stages/api.py", line 225, in cr
eate_pipeline\n" +
  "   await asyncio.gather(*futures)\n" +
  " File \"/usr/lib/python3.6/site-packages/pulpcore/plugin/stages/api.py", line 43, in __c
all__\n" +
  "   await self.run()\n" +
  " File \"/usr/lib/python3.6/site-packages/pulpcore/plugin/stages/artifact_stages.py", lin
e 171, in run\n" +
  "   pb.done += task.result() # download_count\n" +
  " File \"/usr/lib/python3.6/site-packages/pulpcore/plugin/stages/artifact_stages.py", lin
e 197, in _handle_content_unit\n" +
  "   await asyncio.gather(*downloaders_for_content)\n" +
  " File \"/usr/lib/python3.6/site-packages/pulpcore/plugin/stages/models.py", line 88, in
download\n" +
  "   download_result = await downloader.run(extra_data=self.extra_data)\n" +
  " File \"/usr/lib/python3.6/site-packages/pulpcore/download/base.py", line 227, in run\n"
+
  "   return await self._run(extra_data=extra_data)\n" +
```

```
" File \"/usr/lib/python3.6/site-packages/backoff/_async.py\", line 133, in retry\n"   ret = await target(*args, **kwargs)\n" File \"/usr/lib/python3.6/site-packages/pulpcore/download/http.py\", line 218, in _run\n"+\n"   to_return = await self._handle_response(response)\n" File \"/usr/lib/python3.6/site-packages/pulpcore/download/http.py\", line 185, in _handle_response\n"+\n"   chunk = await response.content.read(1048576) # 1 megabyte\n" File \"/usr/lib64/python3.6/site-packages/aiohttp/streams.py\", line 380, in read\n"   await self._wait(\"read\")\n" File \"/usr/lib64/python3.6/site-packages/aiohttp/streams.py\", line 306, in _wait\n"   await waiter\n" File \"/usr/lib64/python3.6/site-packages/aiohttp/helpers.py\", line 656, in __exit__\n"+\n"   raise asyncio.TimeoutError from None\n",\n"description"=>""},
```

## History

### #1 - 06/28/2021 11:37 AM - quba42

Reported in python3-pulpcore-3.11.2

### #2 - 06/29/2021 05:01 PM - bmbouter

- Status changed from NEW to CLOSED - NOTABUG

Pulp 3.14 is getting retry-by-default logic, and we expect that to retry on timeout errors, so please retest with that. Also you can configure the Timeout settings on a Remote, so maybe try those also. See the settings on Remotes with \*\_timeout in the name [https://docs.pulpproject.org/pulpcore/restapi.html#operation/remotes\\_file\\_file\\_create](https://docs.pulpproject.org/pulpcore/restapi.html#operation/remotes_file_file_create)

I'm going to close as NOTABUG since I'm not sure what we can do to improve. We are very interested in this working great for users though, so please comment or reopen with any new info on how we can help you.

### #3 - 06/29/2021 05:03 PM - dalley

- Status changed from CLOSED - NOTABUG to CLOSED - DUPLICATE

I'm moving this to DUPLICATE - it's definitely a bug in practice, just one we have already addressed in the upcoming release

### #4 - 10/18/2021 03:39 PM - quba42

A user running Katello 4.2 is still reporting this issue:

<https://community.theforeman.org/t/correct-instructions-and-process-for-creating-a-katello-debian-repository-and-importing-gpg-keys/25736>

I have suggested increasing the timeout, but it feels like things should not take this long to download...

Anyway, I am just adding the link for future reference.