

## Pulp - Issue #4161

### Content Unit Unassociation progress reporting is highly non-granular

11/14/2018 09:20 PM - dalley

<b>Status:</b>	NEW	<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>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>	
<b>Triaged:</b>	Yes	<b>Quarter:</b>	
<b>Description</b>			
<p>If you sync a large number of units, then do another mirror sync against a different remote, the task will "apparently" hang for several minutes with no progress being made while it processes content unit unassociations. The progress value is only updated at the end of the process (hence, the end of the task), with no intermittent updates to inform the user of progress.</p> <p>In my case, with 20,000 units being unassociated, this took nearly half an hour and resulted in believing I was looking at a problem with the tasking system.</p> <p>Along these lines, content unassociation seemed much slower than association, exacerbating the problem. Perhaps we should investigate unassociation performance.</p>			

#### History

##### #1 - 11/16/2018 04:48 PM - daviddavis

- Triaged changed from No to Yes

##### #2 - 04/26/2019 10:33 PM - bmbouter

- Tags deleted (Pulp 3)

##### #3 - 08/31/2020 06:23 PM - dalley

```
(pulp) [vagrant@pulp3-source-fedora32 devel]$ python rpm_perf_test.py
The task was successful.
Sync time: 378 seconds
The task was successful.
Sync time: 504 seconds
```

Second sync was mirror-mode. It's unclear that the reason it took longer was unassociation vs the actual sync itself (it's probably QueryExistingContent), but it's at least much better than half an hour difference.

I didn't really check progress reporting though, this was a 50,000k unit repo so even if there was no progress report for a minute that's not the worst thing in the world.