

## Pulp - Task #1181

Story # 1150 (CLOSED - NOTABUG): As a user, I can lazily fetch repositories

### Create a download\_unit platform task

08/17/2015 09:18 PM - bmbouter

<b>Status:</b> CLOSED - WONTFIX	<b>Start date:</b>
<b>Priority:</b> Normal	<b>Due date:</b>
<b>Assignee:</b>	<b>% Done:</b> 0%
<b>Category:</b>	<b>Estimated time:</b> 0:00 hour
<b>Sprint/Milestone:</b>	<b>Tags:</b> Pulp 2
<b>Platform Release:</b>	<b>Sprint:</b>
<b>Groomed:</b> No	<b>Quarter:</b>
<b>Sprint Candidate:</b> No	
<b>Description</b> A platform Celery task needs to exist which allows a single unit to be downloaded and saved. This task does not need a TaskStatus but it does need a reservation. The task needs to know which unit to download and will simply download it and put it on disk. Any existing symlinks will already be pointing to its final location so moving it in place is all that is necessary for all repos associated with the unit to begin using it.  The lock requirement ensures that if multiple users request a file, the created download_unit tasks will be processed serially ensuring that the copy into place will not have any corruption due to concurrent writes. If those writes are from multiple machines to the same NFS backend bad things could happen. The lock tag for this type of task should be enough to guarantee it is unique in the catalog. I expect this to be the "unit id + the unit path". Together these should handle both ContentUnits and SharedContentUnits uniquely.  Prior to downloading the unit the code should check if the file already exists. If the file does exist the task should exit immediately doing no work (a noop).	
<b>Related issues:</b> Blocks Pulp - Story #1180: As a user, I can identify a lazy-loader using serv... <b>CLOSED - CURRENTRELEASE</b>	

#### Associated revisions

Revision f5a83095 - 10/02/2015 09:47 PM - jortel@redhat.com

ref #1181 - add content download task.

Revision f5a83095 - 10/02/2015 09:47 PM - jortel@redhat.com

ref #1181 - add content download task.

#### History

#1 - 08/17/2015 09:18 PM - bmbouter

- Blocks Story #1180: As a user, I can identify a lazy-loader using server.conf and Pulp will redirect for content it does not have added

#2 - 08/17/2015 10:55 PM - bmbouter

- Description updated

#3 - 08/18/2015 05:13 PM - rbarlow

I have an implementation detail proposal to consider:

If multiple clients request the same file, we could queue one task per request with resource locking as you suggested in the description. However, each task could check to see if the file has been downloaded already before it does anything. If it finds that the file already exists, it could exit and simply become a NOOP. This is nice because it is simple, and Celery/qpid are known to be able to handle lots of tasks per second so I believe it would still perform well.

#4 - 08/18/2015 08:28 PM - bmbouter

- Description updated

rbarlow wrote:

I have an implementation detail proposal to consider:

If multiple clients request the same file, we could queue one task per request with resource locking as you suggested in the description. However, each task could check to see if the file has been downloaded already before it does anything. If it finds that the file already exists, it could exit and simply become a NOOP. This is nice because it is simple, and Celery/qpid are known to be able to handle lots of tasks per second so I believe it would still perform well.

This sounds like a good improvement, let's do it. I considered the same idea, but I've fixed two bugs before regarding a "file check" that incorrectly reported a file existed when it didn't. This was due to stale inode pointers cached on a client node in an NFS environment where the file was already deleted from the server. This is a different case because nothing is going to delete the file. I've updated the proposal to have it check the file and NOOP if it exists.

**#5 - 09/29/2015 04:30 PM - jortel@redhat.com**

- Status changed from NEW to ASSIGNED

- Assignee set to jortel@redhat.com

**#6 - 12/07/2015 04:22 PM - jcline@redhat.com**

- Status changed from ASSIGNED to CLOSED - WONTFIX

- Assignee deleted (jortel@redhat.com)

Since we no longer plan to have a task that is designed to download a unit, I'm going to close this.

**#7 - 04/15/2019 10:44 PM - bmbouter**

- Tags Pulp 2 added