

## Pulp - Issue #3931

### Pagination of content units returns less than requested limit

08/23/2018 05:42 PM - lucarval

<b>Status:</b>	CLOSED - WONTFIX	<b>Start date:</b>	
<b>Priority:</b>	High	<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>	Master	<b>Sprint Candidate:</b>	No
<b>Platform Release:</b>		<b>Tags:</b>	Pulp 2
<b>OS:</b>	RHEL 7	<b>Sprint:</b>	
<b>Triaged:</b>	Yes	<b>Quarter:</b>	

#### Description

For a pulp repository with the following amount of content units:

```
26 "docker_blob"
16 "docker_image"
39 "docker_manifest"
10 "docker_manifest_list"
52 "docker_tag"
```

When paginating via all content units in repository, pulp groups content units by their type, which is probably ok. However, if the requested page size is less than the size of one of the groups, pulp returns partial results.

With the example repository above, notice the skip and limit value and the amount of items that are returned:

1. POST /pulp/api/v2/repositories/redhat-e2e-container-e2e-container-test-product/search/units/ {"criteria": {"sort": {"unit": "["unit\\_id"](#), "["ascending"](#)}, "skip": 0, "type\_ids": ["docker\_image", "docker\_manifest", "docker\_blob", "docker\_tag", "docker\_manifest\_list", "iso"], "limit": 25, "filters": {"unit": {"\_last\_updated": {"\$lte": 1535038100.200567}}}} - **Returns 25 items**
2. POST /pulp/api/v2/repositories/redhat-e2e-container-e2e-container-test-product/search/units/ {"criteria": {"sort": {"unit": "["unit\\_id"](#), "["ascending"](#)}, "skip": 25, "type\_ids": ["docker\_image", "docker\_manifest", "docker\_blob", "docker\_tag", "docker\_manifest\_list", "iso"], "limit": 25, "filters": {"unit": {"\_last\_updated": {"\$lte": 1535038100.200567}}}} - **Returns 1 item**
3. POST /pulp/api/v2/repositories/redhat-e2e-container-e2e-container-test-product/search/units/ {"criteria": {"sort": {"unit": "["unit\\_id"](#), "["ascending"](#)}, "skip": 26, "type\_ids": ["docker\_image", "docker\_manifest", "docker\_blob", "docker\_tag", "docker\_manifest\_list", "iso"], "limit": 25, "filters": {"unit": {"\_last\_updated": {"\$lte": 1535038100.200567}}}} - **Returns 25 items**

If limit is set to, at least, the amount of content units in the largest group (52 in example above), then this behavior doesn't happen.

Why is this important?

I'd like to implement a "safe" pagination process in dockpulp[0]. This would be done by paginating through content units with one item overlap. Then I can check for the expected overlap and have a bit more confidence that no content units were skipped. When pulp returns less items than expected, it's not clear whether there are more items to be fetched or pulp is just done processing a particular content unit type.

[0] <https://github.com/release-engineering/dockpulp/>

```
$ rpm -qa '*pulp*'
python-pulp-client-lib-2.15.1-1.el7.noarch
pulp-docker-plugins-3.1.1-2.el7eng.noarch
python-pulp-common-2.15.1-1.el7.noarch
python-pulp-agent-lib-2.15.1-1.el7.noarch
pulp-selinux-2.15.1-1.el7.noarch
python-pulp-repoauth-2.15.1-1.el7.noarch
pulp-admin-client-2.15.1-1.el7.noarch
python-pulp-docker-common-3.1.1-2.el7eng.noarch
pulp-server-2.15.1-1.el7.noarch
python-pulp-bindings-2.15.1-1.el7.noarch
```

## History

---

### #1 - 08/24/2018 04:46 PM - CodeHeeler

- Priority changed from Normal to High
- Triaged changed from No to Yes

### #2 - 08/24/2018 05:29 PM - ipanova@redhat.com

I took a quick glance and it looks like not a bug.

~~your first query has skip 0, other query has skip 25. skip is a number that gives the index of the first resource to select. The limit field is a number that gives the maximum amount of resources to select so if skip is 25, the first element selected is 25th and having limit of 25 you will have just 1 result.~~

~~when skip > limit, most likely skip is just ignored, that's why when skip is 26 and limit is 25, you get all 25 results.~~

No my understanding was wrong, We need to look how we handle multiple type\_ids. Was the pagination written for single type only or multiple as well?

### #3 - 08/24/2018 05:35 PM - lucarval

skip > limit seems like a valid use case for me. How else would I get 25 items that are not at the beginning of resource set?

The issue is that when skip is 25 and limit is 25 a single item is returned. But if skip is 26 and limit is 25, then I get 25 additional items.

I'd expect if skip is 25 and limit is 25, that 25 items are returned, because there are enough items to fill up that request.

### #4 - 08/24/2018 05:44 PM - lucarval

Also, when skip is 26 and limit is 25, the items returned are unique and not a repetition of the first 25 items. Skip is definitely **not** getting ignored.

### #5 - 04/12/2019 09:56 PM - bmbouter

- Status changed from NEW to CLOSED - WONTFIX

### #6 - 04/15/2019 10:09 PM - bmbouter

- Tags Pulp 2 added

### #7 - 06/16/2020 10:54 PM - bmbouter

- Category deleted (14)

We are removing the 'API' category per open floor discussion June 16, 2020.