

## Pulp - Issue #4289

Story # 3693 (CLOSED - CURRENTRELEASE): Lazy for Pulp3

### Content delivery breaks if remote is deleted

12/30/2018 06:47 PM - gmbnomis

<b>Status:</b>	CLOSED - CURRENTRELEASE	<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>Groomed:</b>	No
<b>Sprint/Milestone:</b>	3.0.0	<b>Sprint Candidate:</b>	No
<b>Severity:</b>	2. Medium	<b>Tags:</b>	
<b>Version:</b>		<b>Sprint:</b>	Sprint 48
<b>Platform Release:</b>		<b>Quarter:</b>	
<b>OS:</b>			
<b>Triaged:</b>	Yes		

#### Description

#### Problem

It is possible to delete a Remote which is referenced by RemoteArtifacts that are part of a Publication/Repo. As cascading delete is used (<https://github.com/pulp/pulp/blob/88a6ebf31de83a4877909fa93fd94bf122894838/pulpcore/app/models/content.py#L335>), this will delete the RemoteArtifacts as well.

If such an artifact is retrieved via the content app and needs to be streamed, this will fail with an internal server error:

```
Dec 30 17:41:45 pulp3.dev gunicorn[23755]: Traceback (most recent call last):
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:   File "/home/vagrant/.virtualenvs/pulp/lib64/python3.6
/site-packages/aiohttp/web_protocol.py", line 390, in start
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:     resp = await self._request_handler(request)
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:   File "/home/vagrant/.virtualenvs/pulp/lib64/python3.6
/site-packages/aiohttp/web_app.py", line 366, in _handle
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:     resp = await handler(request)
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:   File "/home/vagrant/devel/pulp/pulpcore/content/handl
er.py", line 65, in stream_content
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:     return await self._match_and_stream(path, request)
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:   File "/home/vagrant/devel/pulp/pulpcore/content/handl
er.py", line 179, in _match_and_stream
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:     return await self._stream_content_artifact(request,
web.StreamResponse(), ca)
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:   File "/home/vagrant/devel/pulp/pulpcore/content/handl
er.py", line 223, in _stream_content_artifact
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:     remote_artifact = content_artifact.remoteartifact_s
et.get()
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:   File "/home/vagrant/.virtualenvs/pulp/lib64/python3.6
/site-packages/django/db/models/manager.py", line 82, in manager_method
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:     return getattr(self.get_queryset(), name)(*args, **
kwargs)
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:   File "/home/vagrant/.virtualenvs/pulp/lib64/python3.6
/site-packages/django/db/models/query.py", line 399, in get
Dec 30 17:41:45 pulp3.dev gunicorn[23755]:     self.model._meta.object_name
Dec 30 17:41:45 pulp3.dev gunicorn[23755]: pulpcore.app.models.content.RemoteArtifact.DoesNotExist
: RemoteArtifact matching query does not exist.
```

#### Possible solutions:

1. Block the deletion of the remote as long as referencing RemoteArtifacts exist

2. Return a 404 not found error

**Related issues:**

Copied to Pulp - Test #4464: Content delivery breaks if remote is deleted

**CLOSED - COMPLETE**

**History**

**#1 - 01/08/2019 04:38 PM - jortel@redhat.com**

The cascading delete is the correct behavior. The streamer just needs to 404 when no matching RemoteArtifact records are found. Probably should log at DEBUG that it could not find any RemoteArtifact for the URL for troubleshooting.

**#2 - 01/11/2019 04:39 PM - CodeHeeler**

- *Tags Pulp 3 RC Blocker added*

**#3 - 01/15/2019 04:41 PM - CodeHeeler**

- *Triaged changed from No to Yes*

- *Sprint set to Sprint 47*

**#4 - 01/16/2019 07:40 PM - bmbouter**

I believe cascade deleting the RemoteArtifacts when the Remote is deleted makes sense because there is effectively no Remote to fetch it so all it can do is fail.

It should fail gracefully so +1 to option 2. It should raise an `aihttp.web_exceptions.HttpNotFound` exception in the case where the ContentArtifact exists, but has no Artifact **and** no RemoteArtifact. Is this right?

**#5 - 01/25/2019 11:38 PM - rchan**

- *Sprint changed from Sprint 47 to Sprint 48*

**#6 - 02/04/2019 08:44 PM - bmbouter**

- *Parent task set to #3693*

**#7 - 02/04/2019 08:47 PM - bmbouter**

- *Status changed from NEW to ASSIGNED*

- *Assignee set to bmbouter*

**#8 - 02/04/2019 09:00 PM - bmbouter**

- *Status changed from ASSIGNED to MODIFIED*

This was resolved accidentally by another PR. Specifically this change here:

<https://github.com/pulp/pulp/pull/3826/files#diff-8d791f887df7f2776bf3a0396432c40bR270>

**#9 - 02/25/2019 03:25 PM - bherring**

- *Copied to Test #4464: Content delivery breaks if remote is deleted added*

**#10 - 04/25/2019 06:44 PM - daviddavis**

- *Sprint/Milestone set to 3.0.0*

**#11 - 04/26/2019 10:32 PM - bmbouter**

- *Tags deleted (Pulp 3, Pulp 3 RC Blocker)*

**#12 - 12/13/2019 06:27 PM - bmbouter**

*- Status changed from MODIFIED to CLOSED - CURRENTRELEASE*