

Nectar - Issue #7911

Unable to sync large openshift docker repos

12/02/2020 12:24 PM - ipanova@redhat.com

Status: MODIFIED	Start date:
Priority: Normal	Due date:
Assignee:	Estimated time: 0:00 hour
Category:	
Sprint/Milestone:	
Severity: 2. Medium	Groomed: No
Version - Nectar:	Sprint Candidate: No
Platform Release: 2.21.4	Tags: Pulp 2
Target Release - Nectar:	Sprint: Sprint 86
OS:	Quarter:
Triaged: Yes	

Description

This is a regression of <https://pulp.plan.io/issues/6471>

We are trying to sync the following docker images from satellite server and we got the error while doing sync 1.openshift3/mysql-apb 2.openshift3/mediawiki-apb 3.openshift3/mariadb-apb 4.openshift3/metrics-hawkular-metrics

Getting the below error, when syncing the repository from <http://registry.redhat.io/https://registry.redhat.io>

Error:

```
DKR1020: Image download(s) from http://registry.redhat.io/v2/openshift3/mysql-apb/blobs/sha256:1e3f598e03f841ddc4933fafb143d44393d4ce0517604c5cb2367024975e871c failed. Sync task has failed to prevent a corrupted repository.
```

```
DKR1020: Image download(s) from https://registry.redhat.io/v2/openshift3/mysql-apb/blobs/sha256:2a81c83bf7ad14e424acbfc7eecaf84aec2b273c25124807d820c4da607644e1 failed. Sync task has failed to prevent a corrupted repository.
```

Stacktrace:

```
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) (104, 'Connection reset by peer')
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) Traceback (most recent call last):
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) File "/usr/lib/python2.7/site-packages/nectar/downloaders/threaded.py", line 292, in _fetch
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) for chunk in chunks:
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) File "/usr/lib/python2.7/site-packages/requests/models.py", line 656, in generate
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) for chunk in self.raw.stream(chunk_size, decode_content=True):
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) File "/usr/lib/python2.7/site-packages/requests/packages/urllib3/response.py", line 273, in stream
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) data = self.read(amt=amt, decode_content=decode_content)
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) File "/usr/lib/python2.7/site-packages/requests/packages/urllib3/response.py", line 203, in read
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) data = self._fp.read(amt)
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) File "/usr/lib64/python2.7/httplib.py", line 611, in read
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) s = self.fp.read(amt)
```

```
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) File "/usr/lib64/python2.7/socket.py", line 380, in read
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) data = self._sock.recv(left)
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) File "/usr/lib/python2.7/site-packages/requests/packages/urllib3/contrib/pyopenssl.py", line 188, in recv
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) data = self.connection.recv(*args, **kwargs)
Apr 1 09:45:10 satellite pulp: nectar.downloaders.threaded:ERROR: (22861-60960) SysCallError: (104, 'Connection reset by peer')
```

The workaround is to limit the tags to sync but it is not feasible for some users because they will have to manually update the tags to sync when the upstream repo is updated

Related issues:

Copied from Docker Support - Issue #6471: Unable to sync large openshift dock...

CLOSED - CURRENTRELEASE

Associated revisions

Revision f19d8a30 - 12/02/2020 01:01 PM - ipanova@redhat.com

Fixed RST handling.

Check correct item in tuple.

closes #7911 <https://pulp.plan.io/issues/7911>

Revision 9afa0e9b - 12/03/2020 10:17 PM - ipanova@redhat.com

Check on the exception args.

[noissue]

History

#1 - 12/02/2020 12:24 PM - ipanova@redhat.com

- Copied from Issue #6471: Unable to sync large openshift docker repos added

#2 - 12/02/2020 12:25 PM - ipanova@redhat.com

- Sprint set to Sprint 86

#4 - 12/03/2020 09:33 PM - rchan

@ipanova - did this get merged as per your email?

#5 - 12/03/2020 09:42 PM - ipanova@redhat.com

- Status changed from NEW to MODIFIED

Applied in changeset [f19d8a30f8a5f1b48aee2549136e6000f3fbd112](https://github.com/pulp/nectar/pull/75).

<https://github.com/pulp/nectar/pull/75><https://github.com/pulp/nectar/pull/76>