

RPM Support - Issue #9107

filelists and changelog metadata is not parsed properly - Pulp saves incorrect filelists and changelog metadata and generates incorrect metadata

07/18/2021 12:47 AM - optiz0r

Status: CLOSED - CURRENTRELEASE	Start date:
Priority: Urgent	Due date:
Assignee: dalley	Estimated time: 0:00 hour
Category:	
Sprint/Milestone: 3.14.0	
Severity: 4. Urgent	Groomed: No
Version:	Sprint Candidate: No
Platform Release:	Tags: Katello
OS:	Sprint: Sprint 101
Triaged: Yes	Quarter:

Description

Versions:

- Foreman 2.5.2
- Katello 4.1.1
- pulp-rpm-3.13.3

Upstream: http://lon.mirror.rackspace.com/almalinux/8/BaseOS/x86_64/os/repoata/.

I am seeing filelists not being parsed correctly for Alma Linux repositories. Within Katello UI, I have checked a random sample of packages from all of the Alma repositories I have sync'd (18 in total, x86_64, sources and debug-x86_64 for each of baseos, appstream, highavailability, powertools, extras and devel), and none of them list any files. For comparison, CentOS 8 Stream RPMS do list files.

The generated filelist served up by katello contains each of the packages in the repo but with no package contents. For example, the upstream baseos filelist is 2MB whereas the generated filelist for the katello-hosted repo is just 105KB and looks like the below:

```
<?xml version="1.0" encoding="UTF-8"?>
<filelists xmlns="http://linux.duke.edu/metadata/filelists" packages="1985">
<package pkgid="59c2172d6e423d8adc3248c1983146492471678afccdfae8ab0f66e18a1aaaa5" name="
kernel-debug" arch="x86_64">
  <version epoch="0" ver="4.18.0" rel="305.e18"/>
</package>
<package pkgid="c8120b541261d0ad425f369b1da5eef0aaad5883760ad0e149137027d102207c" name="avahi"
arch="i686">
  <version epoch="0" ver="0.7" rel="20.e18"/>
</package>
```

This is breaking system installs which use the local mirrored content, because packages included in the kickstart base package set have dependencies specified using files, and these are all unsatisfied due to the empty filelists. I have triggered resyncs from upstream mirrors without any change in behaviour.

May be related to similar-looking issue [#8955](#) ?

Related issues:

Related to RPM Support - Test #8972: Add test for properly reading and writin... **CLOSED - COMPLETE**

Associated revisions

Revision 552b0b23 - 07/23/2021 07:46 PM - dalley

Fix unreliable filelist parsing

.findall() is not reliable in the face of namespaces. Don't use it anymore.

closes: #9107 <https://pulp.plan.io/issues/9107> closes: #8972 <https://pulp.plan.io/issues/8972>

History

#1 - 07/19/2021 05:41 AM - dalley

- Assignee set to dalley
- Priority changed from Normal to High
- Severity changed from 2. Medium to 3. High
- Triaged changed from No to Yes

Reproduced

#2 - 07/19/2021 05:41 AM - dalley

- Sprint set to Sprint 101

#3 - 07/19/2021 05:42 AM - dalley

- Sprint/Milestone set to 3.14.0

#4 - 07/19/2021 05:44 AM - dalley

- Related to Test #8972: Add test for properly reading and writing the metadata added

#5 - 07/19/2021 07:16 AM - pulpbot

- Status changed from NEW to POST

PR: https://github.com/pulp/pulp_rpm/pull/2045

#6 - 07/20/2021 08:13 PM - dalley

- Subject changed from AlmaLinux filelists not parsed, generates empty filelists to filelists metadata is not parsed properly, Pulp generates empty filelists
- Priority changed from High to Urgent
- Severity changed from 3. High to 4. Urgent

#7 - 07/20/2021 08:13 PM - dalley

- Subject changed from filelists metadata is not parsed properly, Pulp generates empty filelists to filelists and changelog metadata is not parsed properly - Pulp saves incorrect filelists and changelog metadata and generates incorrect metadata

#9 - 07/22/2021 06:52 AM - gvde

What must be done with the packages which are already synced into the database with the incorrect/missing metadata? If I understand the PR correctly, there are new functions allowing to validate the data in the database with the data in source repository. So the Katello people would have to implement calls to these new functions to allow the broken information in the Katello/pulp3 installation to be fixed?

#10 - 07/22/2021 06:59 AM - dalley

@gvde Most likely we will just write a script that will do an in-place repair - it hasn't been decided yet whether Katello would do anything to run it automatically. My assumption is - probably not.

We'll try to get 3.14.0 released by the end of the day (Thursday), which should let Katello package it by the end of the week. And then once that's out we'll evaluate the best way to address any existing issues.

#11 - 07/22/2021 07:12 AM - dalley

Just to provide a bit of context on what happened, this is the same bug from 3.13.0, except the "fix" apparently didn't work on Python 3.6, which is what most Katello users are using.

Basically the fix relied on a feature that wasn't clearly documented as being exclusive to Python 3.8+, and the development was done on Python 3.8 so it didn't show up, and it doesn't return any kind of error, it just silently did the wrong thing.

The CI does run on Python 3.6, so it would have caught it... if our tests covered it, which they did not unfortunately. They do now.

So this is a really unfortunate sequence of events, but we're going to make sure that A) there's guidance for addressing it, hopefully without too much trouble, and B) making sure it doesn't happen again.

#12 - 07/22/2021 07:29 AM - gvde

Thanks for the information.

But for the future: I guess the metadata won't be fully checked with the metadata in the database during each sync. So in general, it would be probably a good idea to have a check and/or repair function for this. The check could run in a katello cron job every week (like the remove orphans job) and find inconsistencies and send out a notification, if there are any. And the advanced sync repository options could include a metadata repair option to make the corrections.

#13 - 07/22/2021 02:35 PM - dalley

At least as far as having a common place to put repair operations like this, I agree. Pulpcore has repair functionality which will correct missing or corrupted files in storage, or notify if they can't be redownloaded - and we could probably add plugin hooks for additional checks.

I don't know that parsing the metadata a second time during the repair is realistic because it would need to be a completely different parser in order to come up with a different result than the sync originally had, and that is not really an option. But we do now have a few tests to verify the metadata against the original metadata and we will try to keep expanding that coverage.

#14 - 07/23/2021 07:46 PM - dalley

- Status changed from POST to MODIFIED

Applied in changeset [552b0b232da14b06ccc1486e861d7533352dc420](https://pulp.plan.io/changesets/details/552b0b232da14b06ccc1486e861d7533352dc420).

#15 - 07/24/2021 06:11 AM - pulpbot

- Status changed from MODIFIED to CLOSED - CURRENTRELEASE

#16 - 07/25/2021 10:19 AM - gvde

How is this in-place repair going to work? That is going to be essential to get those repositories working again.

Basically, all RPMs synced since the installation of the broken version are missing the information. And as I understand it, it won't be automatically fixed during the next sync with the patched version nor will a repair help...

Obviously, it's urgent as I haven't been able to update a lot of clients because of this bug...

#17 - 07/26/2021 04:58 AM - dalley

I already have a working repair script for the case where all of the RPMs are already downloaded, but the case where there are no RPMs on disk will take a little longer. It still needs some polish and review and a bit of work to make it run in non-development environments.

It should probably be ready within a day or two.

#18 - 07/27/2021 01:55 AM - dalley

The repair script is ready, just waiting on Katello to get the fresh RPMs published (since it is not very useful otherwise).

#19 - 07/27/2021 11:24 AM - dalley

- File repair.py added

@gvde I expect you will be the first "real" user of this script so let me know if the instructions are inadequate. Still waiting on the updated Katello RPMs to be released so be aware that if you do any syncs prior to updating you will need to re-run the script.

Repair Instructions

Attached to this comment is a Python script. Please download it and execute with the following command:

```
DJANGO_SETTINGS_MODULE=pulpcore.app.settings PULP_SETTINGS=/etc/pulp/settings.py python3 repair.py
```

When you run it, the output will look like this:

```
#####  
# Repair tool for Pulp RPM issue #9107 #  
# https://pulp.plan.io/issues/9107 #  
#####
```

Analysis

=====

```
Total packages with broken metadata: 543/1074  
Total packages with broken metadata, with local RPMs available: 508/543  
Total packages with broken metadata, without local RPMs available: 35/543
```

Repair

=====

Package metadata repaired: 543/543 ...FINISHED

Please save the following information for reference, as it cannot be reproduced.
The following publications (potentially) contain incorrect metadata:

```
/pulp/api/v3/publications/rpm/rpm/97dd1c1e-c243-4e4a-8fa4-7498a2790913/  
Created from repository 'on_demand_broken' version 1  
/pulp/api/v3/repositories/rpm/rpm/cab586d7-5f2e-45a5-a16c-6084a6e69417/versions/1/
```

```
Currently being (directly) distributed by:  
/pulp/api/v3/distributions/rpm/rpm/623ba5bd-20b6-455c-bf95-0ecf11641533/
```

One thing to note is that if you're frequently updating your repos, you don't necessarily need to worry too much about the existing publications since the metadata will be refreshed as a matter of course. But if your use case demands ensuring that they be removed immediately, we've provided enough information to do so using the pulp CLI.

I would probably recommend, at minimum, deleting the publications like so:

```
pulp rpm publication destroy --publication /pulp/api/v3/publications/rpm/rpm/97dd1c1e-c243-4e4a-8fa4-7498a2790913/
```

In that case though, please heed the warning that re-runs of the script will not be able to re-print this information, so copy it down somewhere if you need to reference it in the future.

One last note, for "on_demand" synced packages which have not been downloaded by Pulp, listed as "Total packages with broken metadata, without local RPMs available" in the analysis section above, this repair script will download the RPMs (to parse their metadata) without making any changes to Pulp. This may consume a significant amount of time and bandwidth, so if this is a limitation for you, please comment below.

#20 - 07/27/2021 12:08 PM - optiz0r

Am I right in thinking the steps to get back to a working state with Katello would be:

- update the katello rpms, once available, and then one of:
 - Run the repair script, which will download all broken rpms, including ones which have never been needed before, and may not be needed in future
 - Remove and re-add the products/repositories/content-views/versions so that metadata is sync'd correctly on the first sync, which means losing all current content-view versions.

Neither is a particularly big deal in my specific case, just want to understand the options so I can choose the best way to move forward.

#21 - 07/27/2021 12:28 PM - gvde

dalley wrote:

@gvde I expect you will be the first "real" user of this script so let me know if the instructions are inadequate. Still waiting on the updated Katello RPMs to be released so be aware that if you do any syncs prior to updating you will need to re-run the script.

I'll see when I can try it. One thing I have noticed from the source code:

If I understand correctly you are checking for packages with empty filelist and changelogs which translate to this in the database:

```
pulpcore=# select name, epoch, version, release, arch, to_timestamp(time_build) from rpm_package where files =  
'[]' and changelogs = '[]';
```

Does every rpm package must have a non empty filelist and changelog? Or is it allowed to have an rpm which actually has both lists empty?

If it is allowed, I think the script should count the rpms which actually have empty filelist and changelogs and write that number in the end, i.e. packages that in deed have an empty filelist and changelog.

Otherwise it can be confusing for users if they run the script multiple times and there is always a number of rpms to be fixed which never goes away because those rpms "look" broken but in fact are not...

#22 - 07/27/2021 04:05 PM - dalley

optiz0r wrote:

Am I right in thinking the steps to get back to a working state with Katello would be:

- update the katello rpms, once available, and then one of:
 - Run the repair script, which will download all broken rpms, including ones which have never been needed before, and may not be needed in future

It doesn't store them in Pulp, though, it just deletes them afterwards. Unless Pulp already has them in which case it's not needed to download them.

- Remove and re-add the products/repositories/content-views/versions so that metadata is sync'd correctly on the first sync, which means losing all current content-view versions.

Neither is a particularly big deal in my specific case, just want to understand the options so I can choose the best way to move forward.

You *can* do that, and it would work, but I don't think it's necessarily the easiest way to go. You just need to do *\$something* that would trigger a re-publish. A normal sync could do that, for example. I'm going to ask @jsherrill to chime in here.

#23 - 07/27/2021 04:22 PM - dalley

gvde wrote:

dalley wrote:

@gvde I expect you will be the first "real" user of this script so let me know if the instructions are inadequate. Still waiting on the updated Katello RPMs to be released so be aware that if you do any syncs prior to updating you will need to re-run the script.

I'll see when I can try it. One thing I have noticed from the source code:

If I understand correctly you are checking for packages with empty filelist and changelogs which translate to this in the database:

```
pulpcore=# select name, epoch, version, release, arch, to_timestamp(time_build) from rpm_package where file  
s = '[]' and changelogs = '[]';
```

Does every rpm package must have a non empty filelist and changelog? Or is it allowed to have an rpm which actually has both lists empty?

If it is allowed, I think the script should count the rpms which actually have empty filelist and changelogs and write that number in the end, i.e. packages that in deed have an empty filelist and changelog.

Otherwise it can be confusing for users if they run the script multiple times and there is always a number of rpms to be fixed which never goes away because those rpms "look" broken but in fact are not..

It covers only packages that have *both* no files *and* no changelogs. So in practice I think it is exceedingly unlikely (though not strictly speaking impossible) to mis-categorize any RPMs, since nearly all RPMs have at least one file, and all upstream RPMs (provided by RHEL, CentOS, SUSE, EPEL, Fedora) should have at least one changelog due to packaging guidelines.

I can add some text explaining that if it does happen, this could be a possibility, though.

#24 - 07/27/2021 04:46 PM - jsherril@redhat.com

If you need to republish the metadata with katello its fairly easy. For a library repository you can navigate to:

Content > Products > click a product > click a repository > click the actions drop down > republish repository metadata

For a content view version:

Content > Content Views > click a content view > look for the actions drop down in each version row > click 'regenerate repository metadata'

#25 - 07/27/2021 04:55 PM - gvde

It covers only packages that have *both* no files *and* no changelogs. So in practice I think it is exceedingly unlikely (though not strictly speaking impossible) to mis-categorize any RPMs, since nearly all RPMs have at least one file, and all upstream RPMs (provided by RHEL, CentOS, SUSE, EPEL, Fedora) should have at least one changelog due to packaging guidelines.

I can add some text explaining that if it does happen, this could be a possibility, though.

Yes. That would be good. Better than have people ask later, why each run of the repair script (after the first run) is reporting the same, constant number of broken RPMs. It would give the impression that there are some rpms which the repair script could not repair..

#26 - 07/27/2021 04:59 PM - gvde

dalley wrote:

```
DJANGO_SETTINGS_MODULE=pulpcore.app.settings PULP_SETTINGS=/etc/pulp/settings.py python3 repair.py
```

Sorry: one more question: which user should run this? Is it safe to run this as root on a foreman/katello server? Or should it be the pulp user? I guess it's only using the pulp API therefore root should be O.K., but I better confirm before the repair creates the next problems...

#27 - 07/27/2021 05:08 PM - dalley

gvde wrote:

dalley wrote:

```
DJANGO_SETTINGS_MODULE=pulpcore.app.settings PULP_SETTINGS=/etc/pulp/settings.py python3 repair.py
```

Sorry: one more question: which user should run this? Is it safe to run this as root on a foreman/katello server? Or should it be the pulp user? I guess it's only using the pulp API therefore root should be O.K., but I better confirm before the repair creates the next problems...

No problem!

On my system, it worked fine when run as just the normal user. However, my system is not built from Katello RPMs, it uses the upstream installer, so it is possible there are some differences. Katello will be testing this themselves today.

I don't think that the repair wouldn't be able to mess anything up if it failed, the only permanent changes it makes are atomic so if they fail, no harm done.

#28 - 07/27/2021 09:26 PM - iballou

I just tested the repair script. On pulp-rpm 3.13.3, I synced AlmaLinux and verified that the filelist repodata was wrong. I then ran the script and republished the repository's metadata via Katello's UI. The filelist repodata was correct after the fact.

#29 - 07/29/2021 08:29 AM - gvde

iballou wrote:

I just tested the repair script. On pulp-rpm 3.13.3, I synced AlmaLinux and verified that the filelist repodata was wrong. I then ran the script and republished the repository's metadata via Katello's UI. The filelist repodata was correct after the fact.

How exactly did you do that? Run the following as root on the Katello server?

```
# DJANGO_SETTINGS_MODULE=pulpcore.app.settings PULP_SETTINGS=/etc/pulp/settings.py python3 repair.py
```

#30 - 07/29/2021 04:33 PM - dalley

When I was testing this, I never even tried running it as root - I don't think it's necessary. Try running without root, first.

#31 - 07/30/2021 06:50 AM - gvde

O.K. On my katello server I tested the repair script as user pulp. The settings.py is readable for root and group pulp (users apache and pulp).

It repaired ~4000 packages and in the end printed out a list of 307 publications with incorrect metadata. I guess for the standard Katello users that information would be overwhelming...

I have 94 packages left in the database with empty filelist and changelog. I checked a few there and they have in fact at least an empty filelist, e.g. http://software.internet2.edu/rpms/el7/x86_64/main/repodata/ packages pscheduler-bundle-*

Thus, as expected the repair script keeps "repairing" those 94 packages during every following run. Thus the user should know that this happens or the script should be modified to check for packages which really have empty filelists and changelog and count them out as "no repair necessary".

#32 - 07/30/2021 06:58 AM - dalley

@gvde, Thanks for the feedback, I'll update the script and instructions tomorrow. I'm not sure if it's possible to eliminate the false positives but we can do as you suggested and add a note.

Your installation should be fixed now, let us know if you have any other issues.

#33 - 07/30/2021 11:29 AM - voorman

I'm new here, so first of all, thanks for all the good work!

I try to repair my repo's, packages and content views with the repair script provided here. There seemed to be something wrong with a downloaded package from the Rocky repo's. Since I'm not using the product yet, I removed everything Rocky related. First output:

```
Package metadata repaired: 54777/54828Traceback (most recent call last):
  File "repair.py", line 114, in <module>
    packages_to_update = loop.run_until_complete(asyncio.gather(*packages_to_update))
  File "/usr/lib64/python3.6/asyncio/base_events.py", line 484, in run_until_complete
    return future.result()
  File "repair.py", line 100, in repair_on_demand_content
    download_result = await downloader.run()
```

```

File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 258, in run
    return await download_wrapper()
File "/usr/lib/python3.6/site-packages/backoff/_async.py", line 133, in retry
    ret = await target(*args, **kwargs)
File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 256, in download_wrapper
    return await self._run(extra_data=extra_data)
File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 280, in _run
    to_return = await self._handle_response(response)
File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 212, in _handle_response
    await self.finalize()
File "/usr/lib/python3.6/site-packages/pulpcore/download/base.py", line 162, in finalize
    self.validate_digests()
File "/usr/lib/python3.6/site-packages/pulpcore/download/base.py", line 212, in validate_digests
    raise DigestValidationError(self.url)
pulpcore.exceptions.validation.DigestValidationError: A file located at the url https://download.rockylinux.org/pub/rocky/8/BaseOS/x86_64/os/Packages/fwupdate-libs-11-3.el8.x86_64.rpm failed validation due to checksum.
-bash-4.2$

```

But after I removed Rocky and ran the repair script again, I got this error:

```

Repair
=====
Package metadata repaired: 0/51Traceback (most recent call last):
  File "repair.py", line 115, in <module>
    Package.objects.bulk_update(packages_to_update, fields=["files", "changelogs"])
  File "/usr/lib/python3.6/site-packages/django/db/models/manager.py", line 82, in manager_method
    return getattr(self.get_queryset(), name)(*args, **kwargs)
  File "/usr/lib/python3.6/site-packages/django/db/models/query.py", line 493, in bulk_update
    if any(obj.pk is None for obj in objs):
  File "/usr/lib/python3.6/site-packages/django/db/models/query.py", line 493, in <genexpr>
    if any(obj.pk is None for obj in objs):
AttributeError: 'NoneType' object has no attribute 'pk'
-bash-4.2$

```

I guess the script has to finish completely without errors for the fix to work right?

#34 - 07/30/2021 11:55 AM - voorman

Never mind, I had some orphaned content. After removing that, the script finishes without any errors.

#35 - 08/01/2021 08:03 PM - optiz0r

The repair script does not appear to handle errors from the upstream yum repositories very well, it aborts on the first failure, leaving all other packages that could potentially be fixed not fixed.

Examples:

- yum.theforeman.org was returning 503 for much of today
- One of the repos is serving up 404 for a package I've downloaded in the past, since the version has been updated upstream, and the old version was not retained in the repo.

In both cases, the repair script terminated leaving >2000 packages unfixed.

```

Repair
=====
Package metadata repaired: 0/2467Backing off download_wrapper(...) for 0.8s (aiohttp.client_exceptions.ClientResponseError: 503, message='Backend is unhealthy', url=URL('https://yum.theforeman.org/pulpcore/3.14/el8/source/python-djangorestframework-3.12.4-1.el8.src.rpm'))
pulp [None]: backoff:INFO: Backing off download_wrapper(...) for 0.8s (aiohttp.client_exceptions.ClientResponseError: 503, message='Backend is unhealthy', url=URL('https://yum.theforeman.org/pulpcore/3.14/el8/source/python-djangorestframework-3.12.4-1.el8.src.rpm'))
Backing off download_wrapper(...) for 0.2s (aiohttp.client_exceptions.ClientResponseError: 503, message='Backend is unhealthy', url=URL('https://yum.theforeman.org/pulpcore/3.14/el8/source/python-djangorestframework-3.12.4-1.el8.src.rpm'))
pulp [None]: backoff:INFO: Backing off download_wrapper(...) for 0.2s (aiohttp.client_exceptions.ClientResponseError: 503, message='Backend is unhealthy', url=URL('https://yum.theforeman.org/pulpcore/3.14/el8/source/python-djangorestframework-3.12.4-1.el8.src.rpm'))
Backing off download_wrapper(...) for 2.6s (aiohttp.client_exceptions.ClientResponseError: 503, message='Backend is unhealthy', url=URL('https://yum.theforeman.org/pulpcore/3.14/el8/source/python-djangorestframework-3.12.4-1.el8.src.rpm'))
pulp [None]: backoff:INFO: Backing off download_wrapper(...) for 2.6s (aiohttp.client_exceptions.ClientResponseError: 503, message='Backend is unhealthy', url=URL('https://yum.theforeman.org/pulpcore/3.14/el8/source/python-djangorestframework-3.12.4-1.el8.src.rpm'))
Giving up download_wrapper(...) after 4 tries (aiohttp.client_exceptions.ClientResponseError: 503, message='Ba

```

```
ckend is unhealthy', url=URL('https://yum.theforeman.org/pulpcore/3.14/el8/source/python-djangorestframework-3.12.4-1.el8.src.rpm'))
pulp [None]: backoff:ERROR: Giving up download_wrapper(...) after 4 tries (aiohttp.client_exceptions.ClientResponseError: 503, message='Backend is unhealthy', url=URL('https://yum.theforeman.org/pulpcore/3.14/el8/source/python-djangorestframework-3.12.4-1.el8.src.rpm'))
Traceback (most recent call last):
  File "repair.py", line 147, in <module>
    packages_to_update = loop.run_until_complete(asyncio.gather(*packages_to_update))
  File "/usr/lib64/python3.6/asyncio/base_events.py", line 484, in run_until_complete
    return future.result()
  File "repair.py", line 128, in repair_on_demand_content
    download_result = await downloader.run()
  File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 258, in run
    return await download_wrapper()
  File "/usr/lib/python3.6/site-packages/backoff/_async.py", line 133, in retry
    ret = await target(*args, **kwargs)
  File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 256, in download_wrapper
    return await self._run(extra_data=extra_data)
  File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 279, in _run
    self.raise_for_status(response)
  File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 194, in raise_for_status
    response.raise_for_status()
  File "/usr/lib64/python3.6/site-packages/aiohttp/client_reqrep.py", line 1005, in raise_for_status
    headers=self.headers,
aiohttp.client_exceptions.ClientResponseError: 503, message='Backend is unhealthy', url=URL('https://yum.theforeman.org/pulpcore/3.14/el8/source/python-djangorestframework-3.12.4-1.el8.src.rpm')
```

Repair
=====

```
Package metadata repaired: 0/2417Giving up download_wrapper(...) after 1 tries (aiohttp.client_exceptions.ClientResponseError: 404, message='Not Found', url=URL('https://download.opensuse.org/repositories/shells:/fish:/release:/3/CentOS_8/src/fish-3.3.1-1.1.src.rpm'))
pulp [None]: backoff:ERROR: Giving up download_wrapper(...) after 1 tries (aiohttp.client_exceptions.ClientResponseError: 404, message='Not Found', url=URL('https://download.opensuse.org/repositories/shells:/fish:/release:/3/CentOS_8/src/fish-3.3.1-1.1.src.rpm'))
Giving up download_wrapper(...) after 1 tries (aiohttp.client_exceptions.ClientResponseError: 404, message='Not Found', url=URL('https://download.opensuse.org/repositories/shells:/fish:/release:/3/CentOS_8/x86_64/fish-debugsource-3.3.1-1.1.x86_64.rpm'))
pulp [None]: backoff:ERROR: Giving up download_wrapper(...) after 1 tries (aiohttp.client_exceptions.ClientResponseError: 404, message='Not Found', url=URL('https://download.opensuse.org/repositories/shells:/fish:/release:/3/CentOS_8/x86_64/fish-debugsource-3.3.1-1.1.x86_64.rpm'))
Traceback (most recent call last):
  File "repair.py", line 147, in <module>
    packages_to_update = loop.run_until_complete(asyncio.gather(*packages_to_update))
  File "/usr/lib64/python3.6/asyncio/base_events.py", line 484, in run_until_complete
    return future.result()
  File "repair.py", line 128, in repair_on_demand_content
    download_result = await downloader.run()
  File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 258, in run
    return await download_wrapper()
  File "/usr/lib/python3.6/site-packages/backoff/_async.py", line 133, in retry
    ret = await target(*args, **kwargs)
  File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 256, in download_wrapper
    return await self._run(extra_data=extra_data)
  File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 279, in _run
    self.raise_for_status(response)
  File "/usr/lib/python3.6/site-packages/pulpcore/download/http.py", line 194, in raise_for_status
    response.raise_for_status()
  File "/usr/lib64/python3.6/site-packages/aiohttp/client_reqrep.py", line 1005, in raise_for_status
    headers=self.headers,
aiohttp.client_exceptions.ClientResponseError: 404, message='Not Found', url=URL('https://download.opensuse.org/repositories/shells:/fish:/release:/3/CentOS_8/src/fish-3.3.1-1.1.src.rpm')
Backing off download_wrapper(...) for 0.7s (aiohttp.client_exceptions.ServerDisconnectedError: Server disconnected)
pulp [None]: backoff:INFO: Backing off download_wrapper(...) for 0.7s (aiohttp.client_exceptions.ServerDisconnectedError: Server disconnected)
Backing off download_wrapper(...) for 0.1s (aiohttp.client_exceptions.ServerDisconnectedError: Server disconnected)
pulp [None]: backoff:INFO: Backing off download_wrapper(...) for 0.1s (aiohttp.client_exceptions.ServerDisconnectedError: Server disconnected)
Backing off download_wrapper(...) for 0.5s (aiohttp.client_exceptions.ServerDisconnectedError: Server disconnected)
pulp [None]: backoff:INFO: Backing off download_wrapper(...) for 0.5s (aiohttp.client_exceptions.ServerDisconnectedError: Server disconnected)
```

ectedError: Server disconnected

#36 - 08/02/2021 07:16 AM - dalley

- File deleted (repair.py)

#37 - 08/02/2021 07:26 AM - dalley

- File repair.py added

Repair script "version 2" attached. I added a note about false positives as well as some logic to skip repairing on-demand RPMs when they are no longer available.

Hopefully that resolves your problems @optiz0r, but let me know if it doesn't, and I'll keep tweaking.

#38 - 08/02/2021 12:09 PM - optiz0r

Thanks. I worked around it last night by removing the 404ing RPMs from my system entirely. Fortunately it was just a vendor repo with 3 packages in it, so easy to delete previous content-view versions with the older package versions, and then purge orphaned packages from the system.

I'm still having some troubles upgrading systems, but I've not quite worked out if it's fallout from this issue or something separate.

```
[root@test-centos3 ~]# dnf upgrade
Updating Subscription Management repositories.
subscription-manager plugin disabled 6 system repositories with respect of configuration in /etc/dnf/plugins/s
ubscription-manager.conf
CentOS 8 Stream PowerTools (x86_64)
  102 kB/s | 3.0 kB    00:00
CentOS 8 Stream Extras (x86_64)
  52 kB/s | 2.0 kB    00:00
CentOS 8 Stream AppStream (x86_64)
  73 kB/s | 3.0 kB    00:00
CentOS 8 Stream Base OS (x86_64)
  80 kB/s | 2.7 kB    00:00
centos-8-fish
  8.3 kB/s | 2.0 kB    00:00
Choria (Source)
  6.9 kB/s | 1.7 kB    00:00
Choria (x86_64)
  6.7 kB/s | 1.7 kB    00:00
CentOS 8 Stream SIG Ansible-2.9
  8.3 kB/s | 2.0 kB    00:00
Puppet Tools (EL8)
  36 kB/s | 1.7 kB    00:00
Puppet 6 (EL8)
  7.1 kB/s | 1.7 kB    00:00
Error:
  Problem 1: package crypto-policies-scripts-20210617-1.gitc776d3e.el8.noarch requires /usr/bin/sh, but none of
  the providers can be installed
  - cannot install both bash-4.4.20-2.el8.x86_64 and bash-4.4.20-1.el8_4.x86_64
  - cannot install both bash-4.4.20-2.el8.x86_64 and bash-4.4.19-14.el8.x86_64
  - cannot install both bash-4.4.20-1.el8_4.x86_64 and bash-4.4.20-2.el8.x86_64
  - cannot install the best update candidate for package crypto-policies-scripts-20210617-1.gitc776d3e.el8.noa
  rch
  - cannot install the best update candidate for package bash-4.4.20-1.el8_4.x86_64
  Problem 2: package cups-client-1:2.2.6-40.el8.x86_64 requires /usr/sbin/alternatives, but none of the provide
  rs can be installed
  - cannot install both chkconfig-1.19.1-1.el8.x86_64 and chkconfig-1.13-2.el8.x86_64
  - cannot install the best update candidate for package cups-client-1:2.2.6-40.el8.x86_64
  - cannot install the best update candidate for package chkconfig-1.13-2.el8.x86_64
  Problem 3: package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/ar, but none of the providers can be i
  nstalled
  - package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/strings, but none of the providers can be inst
  alled
  - package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/strip, but none of the providers can be instal
  led
  - cannot install both binutils-2.30-108.el8.x86_64 and binutils-2.30-104.el8.x86_64
  - cannot install both binutils-2.30-108.el8.x86_64 and binutils-2.30-102.el8.x86_64
  - cannot install both binutils-2.30-108.el8.x86_64 and binutils-2.30-101.el8.x86_64
  - cannot install both binutils-2.30-104.el8.x86_64 and binutils-2.30-108.el8.x86_64
  - cannot install the best update candidate for package redhat-lsb-core-4.1-47.el8.x86_64
  - cannot install the best update candidate for package binutils-2.30-104.el8.x86_64
  Problem 4: package systemd-239-49.el8.x86_64 requires /usr/bin/getent, but none of the providers can be insta
  lled
  - cannot install both glibc-common-2.28-162.el8.x86_64 and glibc-common-2.28-161.el8.x86_64
```

- cannot install both glibc-common-2.28-162.el8.x86_64 and glibc-common-2.28-160.el8.x86_64
- cannot install both glibc-common-2.28-162.el8.x86_64 and glibc-common-2.28-155.el8.x86_64
- cannot install both glibc-common-2.28-158.el8.x86_64 and glibc-common-2.28-162.el8.x86_64
- cannot install the best update candidate for package systemd-239-48.el8.x86_64
- cannot install the best update candidate for package glibc-common-2.28-161.el8.x86_64

Problem 5: package systemd-239-49.el8.x86_64 requires /usr/bin/getent, but none of the providers can be installed

- cannot install both glibc-common-2.28-162.el8.x86_64 and glibc-common-2.28-161.el8.x86_64
- cannot install both glibc-common-2.28-162.el8.x86_64 and glibc-common-2.28-160.el8.x86_64
- cannot install both glibc-common-2.28-162.el8.x86_64 and glibc-common-2.28-155.el8.x86_64
- cannot install both glibc-common-2.28-158.el8.x86_64 and glibc-common-2.28-162.el8.x86_64
- package systemd-udev-239-49.el8.x86_64 requires systemd(x86-64) = 239-49.el8, but none of the providers can be installed
- package glibc-2.28-162.el8.x86_64 requires glibc-common = 2.28-162.el8, but none of the providers can be installed
- cannot install the best update candidate for package systemd-udev-239-48.el8.x86_64
- cannot install the best update candidate for package glibc-2.28-161.el8.x86_64

Problem 6: problem with installed package redhat-lsb-core-4.1-47.el8.x86_64

- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/[, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/basename, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/cat, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/chgrp, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/chmod, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/chown, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/cksum, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/comm, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/cp, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/csplit, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/cut, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/date, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/dd, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/df, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/dirname, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/du, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/echo, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/env, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/expand, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/expr, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/false, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/fold, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/groups, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/head, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/id, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/install, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/join, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/ln, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/logname, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/ls, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/md5sum, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/mkdir, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/mkfifo, but none of the providers can be installed

```

lled
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/mknod, but none of the providers can be instal
led
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/mktemp, but none of the providers can be insta
lled
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/mv, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/nice, but none of the providers can be install
ed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/nl, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/nohup, but none of the providers can be instal
led
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/od, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/paste, but none of the providers can be instal
led
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/pathchk, but none of the providers can be inst
alled
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/pr, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/printf, but none of the providers can be insta
lled
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/pwd, but none of the providers can be installe
d
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/rm, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/rmdir, but none of the providers can be instal
led
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/seq, but none of the providers can be installe
d
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/sleep, but none of the providers can be instal
led
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/sort, but none of the providers can be install
ed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/split, but none of the providers can be instal
led
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/stty, but none of the providers can be install
ed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/sync, but none of the providers can be install
ed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/tail, but none of the providers can be install
ed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/tee, but none of the providers can be installe
d
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/test, but none of the providers can be install
ed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/touch, but none of the providers can be instal
led
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/tr, but none of the providers can be installed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/true, but none of the providers can be install
ed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/tsort, but none of the providers can be instal
led
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/tty, but none of the providers can be installe
d
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/uname, but none of the providers can be instal
led
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/unexpand, but none of the providers can be ins
talled
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/uniq, but none of the providers can be install
ed
- package redhat-lsb-core-4.1-47.el8.x86_64 requires /usr/bin/wc, but none of the providers can be installed
- cannot install both coreutils-8.30-12.el8.x86_64 and coreutils-8.30-10.el8.x86_64
- package coreutils-8.30-12.el8.x86_64 conflicts with coreutils-single provided by coreutils-single-8.30-9.e
l8.x86_64
- cannot install both coreutils-8.30-10.el8.x86_64 and coreutils-8.30-12.el8.x86_64
- cannot install both coreutils-8.30-9.el8.x86_64 and coreutils-8.30-12.el8.x86_64
- package coreutils-8.30-12.el8.x86_64 conflicts with coreutils-single provided by coreutils-single-8.30-8.e
l8.x86_64
- package coreutils-8.30-12.el8.x86_64 conflicts with coreutils-single provided by coreutils-single-8.30-10.
el8.x86_64
- cannot install both coreutils-8.30-8.el8.x86_64 and coreutils-8.30-12.el8.x86_64
- cannot install the best update candidate for package coreutils-8.30-10.el8.x86_64
Problem 7: package systemd-libs-239-49.el8.x86_64 requires /usr/bin/getent, but none of the providers can be
installed
- cannot install both glibc-common-2.28-162.el8.x86_64 and glibc-common-2.28-161.el8.x86_64
- cannot install both glibc-common-2.28-162.el8.x86_64 and glibc-common-2.28-160.el8.x86_64
- cannot install both glibc-common-2.28-162.el8.x86_64 and glibc-common-2.28-155.el8.x86_64
- cannot install both glibc-common-2.28-158.el8.x86_64 and glibc-common-2.28-162.el8.x86_64

```

```

- package glibc-langpack-en-2.28-162.el8.x86_64 requires glibc-common = 2.28-162.el8, but none of the providers can be installed
- cannot install the best update candidate for package systemd-libs-239-48.el8.x86_64
- cannot install the best update candidate for package glibc-langpack-en-2.28-161.el8.x86_64
Problem 8: package tuned-2.16.0-1.el8.noarch requires /usr/libexec/platform-python, but none of the providers can be installed
- platform-python-3.6.8-38.el8.i686 has inferior architecture
- platform-python-3.6.8-37.el8.i686 has inferior architecture
- cannot install both platform-python-3.6.8-39.el8.x86_64 and platform-python-3.6.8-38.el8.x86_64
- cannot install both platform-python-3.6.8-38.el8.x86_64 and platform-python-3.6.8-39.el8.x86_64
- cannot install both platform-python-3.6.8-37.el8.x86_64 and platform-python-3.6.8-39.el8.x86_64
- cannot install the best update candidate for package tuned-2.16.0-1.el8.noarch
- cannot install the best update candidate for package platform-python-3.6.8-38.el8.x86_64
Problem 9: package subscription-manager-1.28.19-1.el8.x86_64 requires /usr/libexec/platform-python, but none of the providers can be installed
- package platform-python-3.6.8-38.el8.i686 requires python3-libs(x86-32) = 3.6.8-38.el8, but none of the providers can be installed
- package platform-python-3.6.8-37.el8.i686 requires python3-libs(x86-32) = 3.6.8-37.el8, but none of the providers can be installed
- python3-libs-3.6.8-38.el8.i686 has inferior architecture
- package platform-python-3.6.8-38.el8.x86_64 requires python3-libs(x86-64) = 3.6.8-38.el8, but none of the providers can be installed
- python3-libs-3.6.8-37.el8.i686 has inferior architecture
- package platform-python-3.6.8-37.el8.x86_64 requires python3-libs(x86-64) = 3.6.8-37.el8, but none of the providers can be installed
- cannot install both python3-libs-3.6.8-39.el8.x86_64 and python3-libs-3.6.8-38.el8.x86_64
- cannot install both python3-libs-3.6.8-38.el8.x86_64 and python3-libs-3.6.8-39.el8.x86_64
- cannot install both python3-libs-3.6.8-37.el8.x86_64 and python3-libs-3.6.8-39.el8.x86_64
- cannot install the best update candidate for package subscription-manager-1.28.19-1.el8.x86_64
- cannot install the best update candidate for package python3-libs-3.6.8-38.el8.x86_64
(try to add '--allowerase' to command line to replace conflicting packages or '--skip-broken' to skip uninstalleable packages or '--nobest' to use not only best candidate packages)

```

These are local on-demand mirrors of CentOS 8 Stream repositories. Taking a closer look at some of the reported issues, `bash-4.4.20-2.el8.x86_64` is present in the BaseOS repository, and the katello webui is showing the filelist for this package. This system is subscribed to the Library rather than a versioned content view. I have regenerated repository metadata for all repositories within the CentOS 8 Stream product, and done `dnf clean all` on the client.

For the specific case of `bash`, I *think* the problem is that while the UI is showing the filelist on `bash-4.4.20-2`, the yum client is not. Previous versions claim ownership of this file, but `4.4.20-2` does not appear to be.

```

[root@test-centos3 ~]# dnf clean all
Updating Subscription Management repositories.
subscription-manager plugin disabled 6 system repositories with respect of configuration in /etc/dnf/plugins/subscription-manager.conf
75 files removed

[root@test-centos3 ~]# yum whatprovides /usr/bin/sh
Updating Subscription Management repositories.
subscription-manager plugin disabled 6 system repositories with respect of configuration in /etc/dnf/plugins/subscription-manager.conf
CentOS 8 Stream PowerTools (x86_64)
 5.9 MB/s | 3.4 MB    00:00
CentOS 8 Stream Extras (x86_64)
 42 kB/s | 15 kB     00:00
CentOS 8 Stream AppStream (x86_64)
 12 MB/s | 15 MB     00:01
CentOS 8 Stream Base OS (x86_64)
 25 MB/s | 11 MB     00:00
centos-8-fish
 7.2 kB/s | 2.0 kB   00:00
Choria (Source)
 1.7 kB/s | 905 B    00:00
Choria (x86_64)
 9.8 kB/s | 3.0 kB   00:00
CentOS 8 Stream SIG Ansible-2.9
 45 kB/s | 25 kB     00:00
Puppet Tools (EL8)
 20 MB/s | 25 MB     00:01
Puppet 6 (EL8)
 28 MB/s | 29 MB     00:01
bash-4.4.19-14.el8.x86_64 : The GNU Bourne Again shell
Repo                       : sihnon_centos-8-stream_centos-8-stream-baseos-x86_64
Matched from:
Filename                   : /usr/bin/sh

```

```
bash-4.4.20-1.el8_4.x86_64 : The GNU Bourne Again shell
Repo      : @System
Matched from:
Filename  : /usr/bin/sh
```

```
bash-4.4.20-1.el8_4.x86_64 : The GNU Bourne Again shell
Repo      : sihnon-centos-8-stream-centos-8-stream-baseos-x86_64
Matched from:
Filename  : /usr/bin/sh
```

repoquery appears to be telling me the package contains no files, but simultaneously does contain files?

```
[root@test-centos3 ~]# yum repoquery -l bash
Updating Subscription Management repositories.
subscription-manager plugin disabled 6 system repositories with respect of configuration in /etc/dnf/plugins/s
ubscription-manager.conf
CentOS 8 Stream PowerTools (x86_64)
 85 kB/s | 3.0 kB    00:00
CentOS 8 Stream Extras (x86_64)
 58 kB/s | 2.0 kB    00:00
CentOS 8 Stream AppStream (x86_64)
 49 kB/s | 3.0 kB    00:00
CentOS 8 Stream Base OS (x86_64)
 54 kB/s | 2.7 kB    00:00
centos-8-fish
 8.0 kB/s | 2.0 kB    00:00
Choria (Source)
 6.8 kB/s | 1.7 kB    00:00
Choria (x86_64)
 6.7 kB/s | 1.7 kB    00:00
CentOS 8 Stream SIG Ansible-2.9
 50 kB/s | 2.0 kB    00:00
Puppet Tools (EL8)
 6.8 kB/s | 1.7 kB    00:00
Puppet 6 (EL8)
 6.7 kB/s | 1.7 kB    00:00
Package bash-4.4.20-2.el8.x86_64 contains no files
```

```
/etc/skel/.bash_logout
/etc/skel/.bash_profile
/etc/skel/.bashrc
/usr/bin/alias
/usr/bin/bash
/usr/bin/bashbug
/usr/bin/bashbug-64
/usr/bin/bg
/usr/bin/cd
/usr/bin/command
/usr/bin/fc
/usr/bin/fg
/usr/bin/getopts
/usr/bin/hash
/usr/bin/jobs
/usr/bin/read
/usr/bin/sh
/usr/bin/type
/usr/bin/ulimit
/usr/bin/umask
/usr/bin/unalias
/usr/bin/wait
/usr/lib/.build-id
/usr/lib/.build-id/7f
/usr/lib/.build-id/7f/63144f207f3372a7515e04e7595db6adc24677
/usr/share/doc/bash
/usr/share/doc/bash/FAQ
/usr/share/doc/bash/INTRO
/usr/share/doc/bash/RBASH
/usr/share/doc/bash/README
/usr/share/doc/bash/bash.html
/usr/share/doc/bash/bashref.html
/usr/share/info/bash.info.gz
/usr/share/licenses/bash
/usr/share/licenses/bash/COPYING
/usr/share/locale/af/LC_MESSAGES/bash.mo
```

/usr/share/locale/bg/LC_MESSAGES/bash.mo
/usr/share/locale/ca/LC_MESSAGES/bash.mo
/usr/share/locale/cs/LC_MESSAGES/bash.mo
/usr/share/locale/da/LC_MESSAGES/bash.mo
/usr/share/locale/de/LC_MESSAGES/bash.mo
/usr/share/locale/el/LC_MESSAGES/bash.mo
/usr/share/locale/en@boldquot/LC_MESSAGES/bash.mo
/usr/share/locale/en@quot/LC_MESSAGES/bash.mo
/usr/share/locale/eo/LC_MESSAGES/bash.mo
/usr/share/locale/es/LC_MESSAGES/bash.mo
/usr/share/locale/et/LC_MESSAGES/bash.mo
/usr/share/locale/fi/LC_MESSAGES/bash.mo
/usr/share/locale/fr/LC_MESSAGES/bash.mo
/usr/share/locale/ga/LC_MESSAGES/bash.mo
/usr/share/locale/gl/LC_MESSAGES/bash.mo
/usr/share/locale/hr/LC_MESSAGES/bash.mo
/usr/share/locale/hu/LC_MESSAGES/bash.mo
/usr/share/locale/id/LC_MESSAGES/bash.mo
/usr/share/locale/it/LC_MESSAGES/bash.mo
/usr/share/locale/ja/LC_MESSAGES/bash.mo
/usr/share/locale/lt/LC_MESSAGES/bash.mo
/usr/share/locale/nb/LC_MESSAGES/bash.mo
/usr/share/locale/nl/LC_MESSAGES/bash.mo
/usr/share/locale/pl/LC_MESSAGES/bash.mo
/usr/share/locale/pt_BR/LC_MESSAGES/bash.mo
/usr/share/locale/ro/LC_MESSAGES/bash.mo
/usr/share/locale/ru/LC_MESSAGES/bash.mo
/usr/share/locale/sk/LC_MESSAGES/bash.mo
/usr/share/locale/sl/LC_MESSAGES/bash.mo
/usr/share/locale/sr/LC_MESSAGES/bash.mo
/usr/share/locale/sv/LC_MESSAGES/bash.mo
/usr/share/locale/tr/LC_MESSAGES/bash.mo
/usr/share/locale/uk/LC_MESSAGES/bash.mo
/usr/share/locale/vi/LC_MESSAGES/bash.mo
/usr/share/locale/zh_CN/LC_MESSAGES/bash.mo
/usr/share/locale/zh_TW/LC_MESSAGES/bash.mo
/usr/share/man/man1/.1.gz
/usr/share/man/man1/:.1.gz
/usr/share/man/man1/[.1.gz
/usr/share/man/man1/alias.1.gz
/usr/share/man/man1/bash.1.gz
/usr/share/man/man1/bashbug-64.1.gz
/usr/share/man/man1/bashbug.1.gz
/usr/share/man/man1/bg.1.gz
/usr/share/man/man1/bind.1.gz
/usr/share/man/man1/break.1.gz
/usr/share/man/man1/builtin.1.gz
/usr/share/man/man1/builtins.1.gz
/usr/share/man/man1/caller.1.gz
/usr/share/man/man1/cd.1.gz
/usr/share/man/man1/command.1.gz
/usr/share/man/man1/compgen.1.gz
/usr/share/man/man1/complete.1.gz
/usr/share/man/man1/comptopt.1.gz
/usr/share/man/man1/continue.1.gz
/usr/share/man/man1/declare.1.gz
/usr/share/man/man1/dirs.1.gz
/usr/share/man/man1/disown.1.gz
/usr/share/man/man1/enable.1.gz
/usr/share/man/man1/eval.1.gz
/usr/share/man/man1/exec.1.gz
/usr/share/man/man1/exit.1.gz
/usr/share/man/man1/export.1.gz
/usr/share/man/man1/fc.1.gz
/usr/share/man/man1/fg.1.gz
/usr/share/man/man1/getopts.1.gz
/usr/share/man/man1/hash.1.gz
/usr/share/man/man1/help.1.gz
/usr/share/man/man1/history.1.gz
/usr/share/man/man1/jobs.1.gz
/usr/share/man/man1/let.1.gz
/usr/share/man/man1/local.1.gz
/usr/share/man/man1/logout.1.gz
/usr/share/man/man1/mapfile.1.gz
/usr/share/man/man1/popd.1.gz

/usr/share/man/man1/pushd.1.gz
/usr/share/man/man1/read.1.gz
/usr/share/man/man1/readonly.1.gz
/usr/share/man/man1/return.1.gz
/usr/share/man/man1/set.1.gz
/usr/share/man/man1/sh.1.gz
/usr/share/man/man1/shift.1.gz
/usr/share/man/man1/shopt.1.gz
/usr/share/man/man1/source.1.gz
/usr/share/man/man1/suspend.1.gz
/usr/share/man/man1/times.1.gz
/usr/share/man/man1/trap.1.gz
/usr/share/man/man1/type.1.gz
/usr/share/man/man1/typeset.1.gz
/usr/share/man/man1/ulimit.1.gz
/usr/share/man/man1/umask.1.gz
/usr/share/man/man1/unalias.1.gz
/usr/share/man/man1/unset.1.gz
/usr/share/man/man1/wait.1.gz
/etc/skel/.bash_logout
/etc/skel/.bash_profile
/etc/skel/.bashrc
/usr/bin/alias
/usr/bin/bash
/usr/bin/bashbug
/usr/bin/bashbug-64
/usr/bin/bg
/usr/bin/cd
/usr/bin/command
/usr/bin/fc
/usr/bin/fg
/usr/bin/getopts
/usr/bin/hash
/usr/bin/jobs
/usr/bin/read
/usr/bin/sh
/usr/bin/type
/usr/bin/ulimit
/usr/bin/umask
/usr/bin/unalias
/usr/bin/wait
/usr/lib/.build-id
/usr/lib/.build-id/c5
/usr/lib/.build-id/c5/74b717bf593e9d26ca7d2b7eebf6922ec64004
/usr/share/doc/bash
/usr/share/doc/bash/FAQ
/usr/share/doc/bash/INTRO
/usr/share/doc/bash/RBASH
/usr/share/doc/bash/README
/usr/share/doc/bash/bash.html
/usr/share/doc/bash/bashref.html
/usr/share/info/bash.info.gz
/usr/share/licenses/bash
/usr/share/licenses/bash/COPYING
/usr/share/locale/af/LC_MESSAGES/bash.mo
/usr/share/locale/bg/LC_MESSAGES/bash.mo
/usr/share/locale/ca/LC_MESSAGES/bash.mo
/usr/share/locale/cs/LC_MESSAGES/bash.mo
/usr/share/locale/da/LC_MESSAGES/bash.mo
/usr/share/locale/de/LC_MESSAGES/bash.mo
/usr/share/locale/el/LC_MESSAGES/bash.mo
/usr/share/locale/en@boldquot/LC_MESSAGES/bash.mo
/usr/share/locale/en@quot/LC_MESSAGES/bash.mo
/usr/share/locale/eo/LC_MESSAGES/bash.mo
/usr/share/locale/es/LC_MESSAGES/bash.mo
/usr/share/locale/et/LC_MESSAGES/bash.mo
/usr/share/locale/fi/LC_MESSAGES/bash.mo
/usr/share/locale/fr/LC_MESSAGES/bash.mo
/usr/share/locale/ga/LC_MESSAGES/bash.mo
/usr/share/locale/gl/LC_MESSAGES/bash.mo
/usr/share/locale/hr/LC_MESSAGES/bash.mo
/usr/share/locale/hu/LC_MESSAGES/bash.mo
/usr/share/locale/id/LC_MESSAGES/bash.mo
/usr/share/locale/it/LC_MESSAGES/bash.mo
/usr/share/locale/ja/LC_MESSAGES/bash.mo

/usr/share/locale/lt/LC_MESSAGES/bash.mo
/usr/share/locale/nb/LC_MESSAGES/bash.mo
/usr/share/locale/nl/LC_MESSAGES/bash.mo
/usr/share/locale/pl/LC_MESSAGES/bash.mo
/usr/share/locale/pt_BR/LC_MESSAGES/bash.mo
/usr/share/locale/ro/LC_MESSAGES/bash.mo
/usr/share/locale/ru/LC_MESSAGES/bash.mo
/usr/share/locale/sk/LC_MESSAGES/bash.mo
/usr/share/locale/sl/LC_MESSAGES/bash.mo
/usr/share/locale/sr/LC_MESSAGES/bash.mo
/usr/share/locale/sv/LC_MESSAGES/bash.mo
/usr/share/locale/tr/LC_MESSAGES/bash.mo
/usr/share/locale/uk/LC_MESSAGES/bash.mo
/usr/share/locale/vi/LC_MESSAGES/bash.mo
/usr/share/locale/zh_CN/LC_MESSAGES/bash.mo
/usr/share/locale/zh_TW/LC_MESSAGES/bash.mo
/usr/share/man/man1/.1.gz
/usr/share/man/man1/:.1.gz
/usr/share/man/man1/[.1.gz
/usr/share/man/man1/alias.1.gz
/usr/share/man/man1/bash.1.gz
/usr/share/man/man1/bashbug-64.1.gz
/usr/share/man/man1/bashbug.1.gz
/usr/share/man/man1/bg.1.gz
/usr/share/man/man1/bind.1.gz
/usr/share/man/man1/break.1.gz
/usr/share/man/man1/builtin.1.gz
/usr/share/man/man1/builtins.1.gz
/usr/share/man/man1/caller.1.gz
/usr/share/man/man1/cd.1.gz
/usr/share/man/man1/command.1.gz
/usr/share/man/man1/compgen.1.gz
/usr/share/man/man1/complete.1.gz
/usr/share/man/man1/compgen.1.gz
/usr/share/man/man1/continue.1.gz
/usr/share/man/man1/declare.1.gz
/usr/share/man/man1/dirs.1.gz
/usr/share/man/man1/disown.1.gz
/usr/share/man/man1/enable.1.gz
/usr/share/man/man1/eval.1.gz
/usr/share/man/man1/exec.1.gz
/usr/share/man/man1/exit.1.gz
/usr/share/man/man1/export.1.gz
/usr/share/man/man1/fc.1.gz
/usr/share/man/man1/fg.1.gz
/usr/share/man/man1/getopts.1.gz
/usr/share/man/man1/hash.1.gz
/usr/share/man/man1/help.1.gz
/usr/share/man/man1/history.1.gz
/usr/share/man/man1/jobs.1.gz
/usr/share/man/man1/let.1.gz
/usr/share/man/man1/local.1.gz
/usr/share/man/man1/logout.1.gz
/usr/share/man/man1/mapfile.1.gz
/usr/share/man/man1/popd.1.gz
/usr/share/man/man1/pushd.1.gz
/usr/share/man/man1/read.1.gz
/usr/share/man/man1/readonly.1.gz
/usr/share/man/man1/return.1.gz
/usr/share/man/man1/set.1.gz
/usr/share/man/man1/sh.1.gz
/usr/share/man/man1/shift.1.gz
/usr/share/man/man1/shopt.1.gz
/usr/share/man/man1/source.1.gz
/usr/share/man/man1/suspend.1.gz
/usr/share/man/man1/times.1.gz
/usr/share/man/man1/trap.1.gz
/usr/share/man/man1/type.1.gz
/usr/share/man/man1/typeset.1.gz
/usr/share/man/man1/ulimit.1.gz
/usr/share/man/man1/umask.1.gz
/usr/share/man/man1/unalias.1.gz
/usr/share/man/man1/unset.1.gz
/usr/share/man/man1/wait.1.gz

(This is using nightly versions of Katello 4.2, with the python3-pulp-rpm 3.14.0-1 and python3-pulpcore-3.14.3-1 grabbed from Koji, so admittedly not a super clean environment)

#39 - 08/02/2021 03:31 PM - optiz0r

To confirm, the filelist downloaded by the client attached to library is still including no files for latest bash, despite the webui showing them:

```
$ zgrep -A4 '"bash"' d496cb997e3d7cc0b750cc5599f71948185fad60af4a09c201ab129086fd05e7--filelists.xml.gz
<package pkgid="dfa496aff0d2d632d7c6ea114cd57d44f61fea610ddc61d130f95ab38ee84d97" name="bash" arch="x86_64">
  <version epoch="0" ver="4.4.19" rel="14.el8"/>
  <file>/usr/share/man/man1/wait.1.gz</file>
  <file>/usr/share/man/man1/unset.1.gz</file>
  <file>/usr/share/man/man1/unalias.1.gz</file>
--
<package pkgid="d54031866d4dcf7fb9c16e105379a335fbd805b7967336900f4a3cd5934095b8" name="bash" arch="x86_64">
  <version epoch="0" ver="4.4.20" rel="2.el8"/>
</package>
<package pkgid="52c5ce6b8eef11f43c76882c143ede613854c757e52368a6c1761b32e3fb003b" name="cockpit-ws" arch="x86_64">
  <version epoch="0" ver="242" rel="1.el8"/>
--
<package pkgid="00b1627edf47ecbaabd4eceebladb81650caa0672a6a86aa2cade570eb31204" name="bash" arch="x86_64">
  <version epoch="0" ver="4.4.20" rel="1.el8_4"/>
  <file>/usr/share/man/man1/wait.1.gz</file>
  <file>/usr/share/man/man1/unset.1.gz</file>
  <file>/usr/share/man/man1/unalias.1.gz</file>
```

Is this because I need to delete the publications reported at the end of the output script? I have not done this yet, as it calls for me to run the "pulp" executable, which does not appear to be present on my katello system (or is not in the default PATH at least).

#40 - 08/02/2021 03:34 PM - dalley

Yes, that's likely the reason. You need to make sure that the new metadata has been generated and is being distributed. It's easy to do through Katello, see Jsherrill's comment here: <https://pulp.plan.io/issues/9107#note-24>

#41 - 08/02/2021 05:28 PM - optiz0r

Jsherrill's comment is what I was doing to regenerate the metadata for both the library version, and the content-view versions prior to deleting the publications.

I installed the python3-pulp-cli RPM, worked out that I need to pass --cert /etc/foreman/client_cert.pem --key /etc/foreman/client_cert.key to authenticate, and swap the --publication option for --href for compat with cli 0.10.1 and then the publications deleted successfully. Each of the deletion commands ran without producing any output, and re-running the same delete again yielded a "not found" error.

This appears to have gotten me into a worse state though. Now various actions are failing with errors relating to missing publications.

Re-generating repository metadata:

```
Error message: the server returns an error
HTTP status code: 400
Response headers: {"date"=>"Mon, 02 Aug 2021 14:17:11 GMT", "server"=>"unicorn", "content-type"=>"application/json", "vary"=>"Accept, Cookie", "allow"=>"GET, PUT, PATCH, DELETE, HEAD, OPTIONS", "x-frame-options"=>"SAMEORIGIN", "content-length"=>"62", "correlation-id"=>"ec6157e0-8df8-414b-b155-c724e67e74c3", "access-control-expose-headers"=>"Correlation-ID", "via"=>"1.1 whitefall.jellybean.sihnon.net", "connection"=>"close"}
Response body: {"publication":["Invalid hyperlink - Object does not exist."]}
```

Synchronising from upstream:

```
Error message: the server returns an error
HTTP status code: 400
Response headers: {"date"=>"Mon, 02 Aug 2021 14:19:53 GMT", "server"=>"unicorn", "content-type"=>"application/json", "vary"=>"Accept, Cookie", "allow"=>"GET, PUT, PATCH, DELETE, HEAD, OPTIONS", "x-frame-options"=>"SAMEORIGIN", "content-length"=>"62", "correlation-id"=>"c886e9a0-0294-4acd-99ac-330c9eb99dcf", "access-control-expose-headers"=>"Correlation-ID", "via"=>"1.1 whitefall.jellybean.sihnon.net", "connection"=>"close"}
Response body: {"publication":["Invalid hyperlink - Object does not exist."]}
```

Creating new content-version:

```
PulpRpmClient::ApiError: Error message: the server returns an error
HTTP status code: 400
Response headers: {"date"=>"Mon, 02 Aug 2021 15:15:35 GMT", "server"=>"unicorn", "content-type"=>"application/json", "vary"=>"Accept, Cookie", "allow"=>"GET, PUT, PATCH, DELETE, HEAD, OPTIONS", "x-frame-options"=>"SAMEORIGIN", "content-length"=>"62", "correlation-id"=>"ff3c3a08-54e3-4104-840e-41a3c0ed6e74", "access-control-expose-headers"=>"Correlation-ID", "via"=>"1.1 whitefall.jellybean.sihnon.net", "connection"=>"close"}
Response body: {"publication":["Invalid hyperlink - Object does not exist."]}
```

These tasks are in stopped/error state, and attempts to re-do the same action again are failing because resources are locked. The stuck task cannot be resumed (it returns to the stopped/error state presumably because of the same root cause), and attempts to unlock/force-unlock the tasks doesn't change the status.

I think I'm getting close to needing to wipe everything and start again. Katello/pulp are fairly new to me, so it's entirely plausible I've done something unwise along the way that's gotten me into a worse state.

#42 - 08/02/2021 06:40 PM - jsherril@redhat.com

optiz0r

Deleting publications with pulp-cli will leave katello in a bad state. You should be able to recover from it by:

1. doing an advanced sync and selecting 'full sync'
2. regenerating metadata for the repository under library

If these things aren't working, maybe there is some case that we're not accounting for. I can do a simple test to see if i can reproduce.

#43 - 08/02/2021 08:16 PM - caseybea

The updated RPM made it way to me upon returning from a brief vacation; I did the update, ran the repair script (V2), and republished metadata and content views in katello. All good on my wide of things. Thanks all!

#44 - 08/03/2021 12:51 AM - optiz0r

Thanks jsherril, dalley, finally got the canary hosts upgrading cleanly again.

To recap, for other Katello users needing to run through this, the final working sequence for me was:

- Install python3-pulpcore-3.14.3+ and python3-pulp-rpm-3.14.0+ (from Koji for now, until they hit yum.theforeman.org repos), python3-pulp-cli-0.10.1+ from normal foreman repos
- Run the V2 repair script attached to this ticket
- Delete the publications, as advised by the repair script (substituting --publication for --href, unless V2 of the script fixed that?)
- For each repo in each product run a Advanced Sync->Complete sync, then a Regenerate Repo Metadata
- For each version of each non-composite content view, regenerate repo metadata
- For each version of each composite content view, regenerate repo metadata

If at any point tasks fail with the error Invalid hyperlink - Object does not exist. (which I ran into several times where I tried to update a composite content-view before I'd successfully repaired all the constituent content views):

- Open the dynflow console for the task
- For each step showing (error), click the skip button
- Back on the task view, hit Resume and wait for the task to complete to 100% with warnings
- Check the errors to see which repository IDs triggered the error
- Navigate to [https://foreman.example.com/products/1/repositories/\\${repo_id}](https://foreman.example.com/products/1/repositories/${repo_id}) and make sure the above steps were followed in the correct sequence for this repository (following the obvious links in the foreman GUI took me to the repo under the content-view, rather than under the product, where triggering a sync+regenerate did not fix the issue)

#45 - 08/03/2021 01:18 AM - dalley

- File deleted (repair.py)

#46 - 08/03/2021 01:18 AM - dalley

- File repair.py added

#47 - 08/03/2021 01:20 AM - dalley

@optiz0r, Glad everything is working now! And thank you for that update.

The --href vs --publication was an oversight which I've now fixed. You're correct that --href is the right parameter to use.

#48 - 08/03/2021 06:15 AM - gvde

I have never deleted any publication listed in the output of the repair script.

I ran the repair script.

Then ran a complete sync on all my repositories. (Which I think isn't even necessary as the standard sync did it's job, too).

Published and promoted a new version of all my content views to all LE.

I think it's best for Katello users NOT to delete the publications from pulp but use only Katello functions...

#49 - 08/06/2021 04:11 PM - equipe_serveurs@exane.com

Hello,

I was also hit by this issue. For info trying to get the repair script running I ran into an exception, I am looking for a workaround ...:

```
UnicodeDecodeError: 'utf-8' codec can't decode byte 0xe4 in position 11: invalid continuation byte
```

The above exception was the direct cause of the following exception:

```
Traceback (most recent call last):
  File "repair.py", line 98, in <module>
    new_package = Package.createrepo_to_dict(cr_pkginfo)
  File "/srv/pulp/app/.venv/lib/python3.6/site-packages/pulp_rpm/app/models/package.py", line 280, in
  createrepo_to_dict
    PULP_PACKAGE_ATTRS.CHANGELOGS: getattr(package, CR_PACKAGE_ATTRS.CHANGELOGS, []),
SystemError: <built-in function getattr> returned a result with an error set
```

Well this is triggered by a single package (On various version) system-config-kickstart-2.8.6.2-1.el6.noarch

#50 - 08/06/2021 05:08 PM - dalley

Try applying this patch to the repair script and run it again. It's still going to fail but it will give more information about what is responsible so that I can attempt to reproduce locally.

```
diff --git a/repair.py b/repair.py
index 935dfd0d..972b4b4f 100644
--- a/repair.py
+++ b/repair.py
@@ -94,7 +94,13 @@ for ca in broken_content_artifacts_with_files.select_related("artifact", "content
     temp_file.flush()
     cr_pkginfo = cr.package_from_rpm(temp_file.name)

-    new_package = Package.createrepo_to_dict(cr_pkginfo)
+    try:
+        new_package = Package.createrepo_to_dict(cr_pkginfo)
+    except Exception as e:
+        print("Problem parsing package {}".format(cr_pkginfo.nevra()))
+        print(cr_pkginfo.checksum_type)
+        print(cr_pkginfo.changelogs)
+        raise e
     old_package = ca.content.rpm_package # not very efficient, but for a one-time script it's "fine".

     old_package.files = new_package["files"]
```

Files

repair.py	8.26 KB	08/02/2021	dalley
-----------	---------	------------	--------