

RPM Support - Issue #9329

Regression: Syncing mirrorlist based remotes do fail again since 3.14.1

09/01/2021 06:48 PM - cityofships

Status:	CLOSED - CURRENTRELEASE	Start date:	
Priority:	High	Due date:	
Assignee:	dalley	Estimated time:	0:00 hour
Category:		Groomed:	No
Sprint/Milestone:	3.16.0	Sprint Candidate:	No
Severity:	3. High	Tags:	
Version:		Sprint:	Sprint 105
Platform Release:		Quarter:	
OS:			
Triaged:	Yes		
Description			
Very much related to: https://pulp.plan.io/issues/8981			
Affected versions are at least 3.14.1, 3.14.2 & 3.14.3			
Related issues:			
Copied to RPM Support - Backport #9394: Backport #9329 "Regression: Syncing m...		CLOSED - CURRENTRELEASE	
Copied to RPM Support - Issue #9464: Regression: Syncing username+password au...		CLOSED - CURRENTRELEASE	

Associated revisions

Revision d328e0ab - 09/22/2021 07:34 PM - dalley

Fix a regression in mirrorlist sync support

closes: #9329 <https://pulp.plan.io/issues/9329>

History

#1 - 09/02/2021 03:42 AM - dalley

- Priority changed from Normal to High
- Severity changed from 2. Medium to 3. High
- Sprint set to Sprint 104

#2 - 09/02/2021 05:55 AM - dalley

@cityofships could you describe what goes wrong? Error messages, incorrect metadata generated, etc? The pastebin link linked on the original issue is expired.

This issue was just filed, it could be the same - although I don't know if you're seeing different symptoms. <https://pulp.plan.io/issues/9328>

#3 - 09/02/2021 05:55 AM - dalley

- Status changed from NEW to ASSIGNED
- Assignee set to dalley

#4 - 09/02/2021 04:39 PM - ggainey

[dalley](#) - this sequence fails:

```
pulp rpm remote create --name test --url "http://mirrorlist.centos.org/?arch=x86_64&release=8&repo=opstools-collectd-5"
pulp rpm repository create --name test --remote test
pulp rpm repository sync --name test
pulp rpm repository sync --name test
Started background task /pulp/api/v3/tasks/blf34e98-7ea2-4994-84b5-58ea3ee4b1e5/
.Error: Task /pulp/api/v3/tasks/blf34e98-7ea2-4994-84b5-58ea3ee4b1e5/ failed: '404, message='Not Found', url=URL('http://mirrorlist.centos.org/%3Farch=x86_64&release=8&repo=opstools-collectd-5')'
```

This sequence works fine:

```
pulp rpm remote create --name fmirror --url "https://mirrors.fedoraproject.org/mirrorlist?repo=epel-modular-8&arch=x86_64&infra=stock&content=centos"
pulp rpm repository create --name fmirror --remote fmirror
pulp rpm repository sync --name fmirror
pulp rpm repository sync --name fmirror
Started background task /pulp/api/v3/tasks/ef10b45f-0dda-4a4e-b2c3-51e48fe2e11e/
.....Done.
```

(The URL above is what gets used in our test

pulp_rpm.tests.functional.api.test_sync.BasicSyncTestCase.test_sync_from_valid_mirror_list_feed_with_params , which passes)

But this sequence fails:

```
pulp rpm remote create --name fmirror2 --url "https://mirrors.fedoraproject.org/mirrorlist/?repo=epel-modular-8&arch=x86_64&infra=stock&content=centos"
pulp rpm repository create --name fmirror2 --remote fmirror2
pulp rpm repository sync --name fmirror2
Started background task /pulp/api/v3/tasks/83b29ef6-e5d1-430c-bfd6-a187ee86576c/
.Error: Task /pulp/api/v3/tasks/83b29ef6-e5d1-430c-bfd6-a187ee86576c/ failed: '404, message='Not Found', url=URL('https://mirrors.fedoraproject.org/mirrorlist/?repo=epel-modular-8&arch=x86_64&infra=stock&content=centos')
```

Note the only difference between 2 and 3 is "https.../mirrorlist?..." and "https.../mirrorlist/?..." . Clearly, tripping over a(nother) URL-parsing-subtlety.

#5 - 09/02/2021 05:11 PM - dalley

- Status changed from ASSIGNED to NEW

- Assignee deleted (dalley)

#6 - 09/02/2021 05:11 PM - dalley

- Triaged changed from No to Yes

#8 - 09/09/2021 06:16 AM - dalley

- Status changed from NEW to ASSIGNED

- Assignee set to dalley

#9 - 09/09/2021 05:34 PM - dalley

Discussion on matrix here: [https://riot.im/app/#/room/#pulp-rpm:matrix.org/\\$YdlA42DKgiCpJF2Ngz7k1y07ZrZelSnPJgf4g1F2ipY](https://riot.im/app/#/room/#pulp-rpm:matrix.org/$YdlA42DKgiCpJF2Ngz7k1y07ZrZelSnPJgf4g1F2ipY)

And here: [https://riot.im/app/#/room/#pulp-rpm:matrix.org/\\$8b98jRZv1khE2Ps3Q_gWcgdeNRpMt_ykyFlbDICXC9E](https://riot.im/app/#/room/#pulp-rpm:matrix.org/$8b98jRZv1khE2Ps3Q_gWcgdeNRpMt_ykyFlbDICXC9E)

#10 - 09/10/2021 12:28 AM - rchan

- Sprint changed from Sprint 104 to Sprint 105

#11 - 09/10/2021 02:36 AM - dalley

[ggainey](#), example 3 doesn't seem to work even from a browser, or from httpie, so I'm not sure that's a Pulp issue.

```
(pulp) [vagrant@pulp3-source-fedora34 devel]$ http "https://mirrors.fedoraproject.org/mirrorlist/?repo=epel-modular-8&arch=x86_64&infra=stock&content=centos"
HTTP/1.1 404 Not Found
```

Example 1 however is a legit bug.

#12 - 09/10/2021 03:58 AM - dalley

- Status changed from ASSIGNED to POST

#13 - 09/10/2021 04:17 AM - pulpbot

PR: https://github.com/pulp/pulp_rpm/pull/2120

#14 - 09/10/2021 11:05 PM - dalley

@cityofships FYI we'll push a new release as soon as we fix this second mirrorlist-related issue: <https://pulp.plan.io/issues/9328>

#15 - 09/11/2021 07:11 AM - dalley

- Sprint/Milestone set to 3.14.4

#16 - 09/13/2021 10:13 PM - dalley

- Sprint/Milestone changed from 3.14.4 to 3.16.0

#17 - 09/13/2021 10:14 PM - dalley

- Copied to Backport #9394: Backport #9329 "Regression: Syncing mirrolist based remotes do fail again since 3.14.1" to 3.14.z added

#18 - 09/22/2021 07:34 PM - dalley

- Status changed from POST to MODIFIED

Applied in changeset [d328e0ab975ea654950c050c3c642c3002c946ed](#).

#19 - 09/24/2021 01:12 PM - martin

We have issue with something similar:

```
FAILED - RETRYING: Sync vendor-product-develop-8 repository from remote (15 retries left).Result was: {
  "attempts": 1,
  "changed": false,
  "invocation": {
    "module_args": {
      "mirror": true,
      "password": "VALUE_SPECIFIED_IN_NO_LOG_PARAMETER",
      "pulp_url": "http://localhost:24817",
      "refresh_api_cache": false,
      "remote": "vendor-product-develop-8-remote",
      "repository": "vendor-product-develop-8",
      "username": "admin",
      "validate_certs": true
    }
  },
  "msg": "
Task failed to complete. (failed; https://username_with_underscores-and-dashes:somepasswordstring%40downloads.
vendor.url.com/repo/product/rhel/8/x86_64/current/repodata/repomd.xml) ",
  "retries": 16
}
```

Is it related or should I open separate issue? Seems that escaping @ to %40 breaks the URL and repo sync.

#20 - 09/24/2021 03:51 PM - dalley

@martin Which version of pulp_rpm? If it is the same bug, then it should be fixed in 3.14.4

It's possible that it's a separate bug however.

#21 - 09/27/2021 12:14 PM - martin

[dalley](#), it is happening on latest pulpcore 3.14.6 and pulp_rpm 3.14.4

#22 - 09/27/2021 10:13 PM - dalley

:/

@martin Could you please provide your remote configuration? You can use "pulp rpm remote show --name \$remote_name"

#23 - 09/28/2021 06:15 PM - johng

Not sure if this is the correct place to ask but will 3.15.x get a similar fix?

#24 - 09/28/2021 06:23 PM - dalley

Well, if it is still broken in some situations even with this patch, then we'll need new releases for everyone.

Probably we will not do a 3.15.z because 3.16 will be releasing in a week or so. And most users are either using Katello RPMs, which are on the 3.14 branch, or the container images, which use the latest upstream releases available. But I'm open to hearing out anyone that has a particular need for a 3.15.z release.

#25 - 09/28/2021 06:37 PM - johng

I'm attempting to use the container release as described <https://pulpproject.org/pulp-in-one-container/>

docker.io/pulp/pulp:latest which ships with 3.15.0

```
{"versions":[{"component":"core","version":"3.15.2"},{"component":"rpm","version":"3.15.0"}].....
```

The 3.14 tag is a month old so presumably doesn't include this fix either.

#26 - 09/28/2021 08:31 PM - dalley

You're right about the 3.14 tag being out of date, I just pushed a new one (and also probably document this as something we need to automate or trigger more frequently).

But if it turns out it *is* still broken then it won't help. I'm trying to reproduce that currently, I have not yet been able to do so.

As for the "latest" tag, it should just use the newest version, whether that is pulp_rpm 3.15.1, or pulp_rpm 3.16.0. So it won't matter too much what type of release we do so long as we do one soon.

#27 - 09/28/2021 09:33 PM - johng

Ok thanks, I'll await 3.16.0 then.

I've partially applied the fix locally in the meantime so I can sync the repo I was having trouble with.

#28 - 09/29/2021 06:33 AM - dalley

@martin Please file a separate issue

#29 - 09/29/2021 10:55 AM - martin

- Copied to Issue #9464: Regression: Syncing username+password authenticated remotes fails added

#30 - 10/20/2021 09:07 PM - pulpbot

- Status changed from MODIFIED to CLOSED - CURRENTRELEASE