

RPM Support - Issue #3310

Pulp can't sync from some mirrorlists

01/22/2018 10:20 PM - Ichimonji10

Status:	CLOSED - CURRENTRELEASE	Start date:	
Priority:	Normal	Due date:	
Assignee:	ttereshc	Estimated time:	0:00 hour
Category:		Groomed:	No
Sprint/Milestone:		Sprint Candidate:	No
Severity:	2. Medium	Tags:	Pulp 2
Version:		Sprint:	Sprint 31
Platform Release:	2.15.1	Quarter:	
OS:			
Triaged:	Yes		

Description

Pulp's RPM plugin supports mirrorlist files. When creating an RPM repository, you can set the importer's feed to a mirrorlist URL instead of a literal repository URL. When that repository is synced, Pulp will randomly select URLs from that mirrorlist until the sync succeeds, or until all URLs have been tried.

Currently, this support is broken. We know this because all of the mirrorlist tests are broken. You can run the relevant test yourself with `python -m unittest pulp_smash.tests.pulp2.rpm.api_v2.test_mirrorlist`. All of the tests will fail. Here's the error that's returned:

```
{'_href': '/pulp/api/v2/tasks/f1d6e49c-0d78-4a32-a560-c175a9af38d9/',
'_id': {'$oid': '5a6646a67c74869866747d31'},
'_ns': 'task_status',
'error': {'code': 'RPM1004',
          'data': {'reason': 'Not Found'},
          'description': 'Error retrieving metadata: Not Found',
          'sub_errors': []},
'exception': None,
'finish_time': '2018-01-22T20:16:45Z',
'id': '5a6646a67c74869866747d31',
'progress_report': {'yum_importer': {'comps': {'state': 'FINISHED'},
                                     'content': {'details': {'drpm_done': 0,
                                                             'drpm_total': 0,
                                                             'rpm_done': 0,
                                                             'rpm_total': 0},
                                     'error_details': [],
                                     'items_left': 0,
                                     'items_total': 0,
                                     'size_left': 0,
                                     'size_total': 0,
                                     'state': 'FINISHED'},
                  'distribution': {'error_details': [],
                                   'items_left': 0,
                                   'items_total': 0,
                                   'state': 'FINISHED'},
                  'errata': {'state': 'FINISHED'},
                  'metadata': {'state': 'FINISHED'},
                  'purge_duplicates': {'state': 'FINISHED'}}}},
'queue': 'reserved_resource_worker-1@fedora-25-pulp-2-15-nightly.dq',
'result': None,
'spawned_tasks': [],
'start_time': '2018-01-22T20:16:38Z',
'state': 'error',
'tags': ['pulp:repository:a5fd6c6a-541d-4c87-9ac1-fceb9ca18eb3',
         'pulp:action:sync'],
'task_id': 'f1d6e49c-0d78-4a32-a560-c175a9af38d9',
'task_type': 'pulp.server.managers.repo.sync.sync',
```

```
'traceback': 'Traceback (most recent call last):\n'
' File "/usr/lib/python2.7/site-packages/celery/app/trace.py", '
'line 240, in trace_task\n'
' R = retval = fun(*args, **kwargs)\n'
' File '
'"/usr/lib/python2.7/site-packages/pulp/server/async/tasks.py", '
'line 529, in __call__\n'
' return super(Task, self).__call__(*args, **kwargs)\n'
' File '
'"/usr/lib/python2.7/site-packages/pulp/server/async/tasks.py", '
'line 107, in __call__\n'
' return super(PulpTask, self).__call__(*args, **kwargs)\n'
' File "/usr/lib/python2.7/site-packages/celery/app/trace.py", '
'line 438, in __protected_call__\n'
' return self.run(*args, **kwargs)\n'
' File '
'"/usr/lib/python2.7/site-packages/pulp/server/controllers/repository.py", '
'line 769, in sync\n'
' sync_report = sync_repo(transfer_repo, conduit, '
'call_config)\n'
' File '
'"/usr/lib/python2.7/site-packages/pulp/server/async/tasks.py", '
'line 732, in wrap_f\n'
' return f(*args, **kwargs)\n'
' File '
'"/usr/lib/python2.7/site-packages/pulp_rpm/plugins/importers/yum/importer.py", '
'line 82, in sync_repo\n'
' report = self._current_sync.run()\n'
' File '
'"/usr/lib/python2.7/site-packages/pulp_rpm/plugins/importers/yum/sync.py", '
'line 334, in run\n'
' reason=self.repomd_not_found_reason)\n'
'PulpCodedException: Error retrieving metadata: Not Found\n',
'worker_name': 'reserved_resource_worker-1@fedora-25-pulp-2-15-nightly}'
```

Here's the formatted traceback:

```
Traceback (most recent call last):
  File "/usr/lib/python2.7/site-packages/celery/app/trace.py", line 240, in trace_task
    R = retval = fun(*args, **kwargs)
  File "/usr/lib/python2.7/site-packages/pulp/server/async/tasks.py", line 529, in __call__
    return super(Task, self).__call__(*args, **kwargs)
  File "/usr/lib/python2.7/site-packages/pulp/server/async/tasks.py", line 107, in __call__
    return super(PulpTask, self).__call__(*args, **kwargs)
  File "/usr/lib/python2.7/site-packages/celery/app/trace.py", line 438, in __protected_call__
    return self.run(*args, **kwargs)
  File "/usr/lib/python2.7/site-packages/pulp/server/controllers/repository.py", line 769, in sync
    sync_report = sync_repo(transfer_repo, conduit, call_config)
  File "/usr/lib/python2.7/site-packages/pulp/server/async/tasks.py", line 732, in wrap_f
    return f(*args, **kwargs)
  File "/usr/lib/python2.7/site-packages/pulp_rpm/plugins/importers/yum/importer.py", line 82, in
sync_repo
    report = self._current_sync.run()
  File "/usr/lib/python2.7/site-packages/pulp_rpm/plugins/importers/yum/sync.py", line 334, in run
    reason=self.repomd_not_found_reason)
PulpCodedException: Error retrieving metadata: Not Found
```

On the test host, there are:

- No errors logged by journalctl.
- No failed services, according to systemctl.
- No SELinux denials, according to /var/log/audit/audit.log.

This failure is reproducible with the Pulp 2.15 nightlies on all distributions. (Fedora 25, 26, 27, and RHEL 7.) Here's the packages installed on one test system:

```
[root@fedora-25-pulp-2-15-nightly ~]# rpm -qa | grep -i pulp | sort
```

pulp-admin-client-2.15.1-0.2.beta.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.fc25.noarch
pulp-deb-admin-extensions-1.6.0-1.fc25.noarch
pulp-deb-plugins-1.6.0-1.fc25.noarch
pulp-docker-admin-extensions-3.1.1-0.2.beta.git.3.b0dfae3.git.3.b0dfae3.git.3.b0dfae3.git.3.b0dfae3.fc25.noarch
pulp-docker-plugins-3.1.1-0.2.beta.git.3.b0dfae3.git.3.b0dfae3.git.3.b0dfae3.git.3.b0dfae3.fc25.noarch
pulp-ostree-admin-extensions-1.3.0-1.fc25.noarch
pulp-ostree-plugins-1.3.0-1.fc25.noarch
pulp-puppet-admin-extensions-2.15.1-0.2.beta.git.4.b0dfae3.git.4.b0dfae3.git.4.b0dfae3.git.4.b0dfae3.fc25.noarch
pulp-puppet-plugins-2.15.1-0.2.beta.git.4.b0dfae3.git.4.b0dfae3.git.4.b0dfae3.git.4.b0dfae3.fc25.noarch
pulp-puppet-tools-2.15.1-0.2.beta.git.4.b0dfae3.git.4.b0dfae3.git.4.b0dfae3.git.4.b0dfae3.fc25.noarch
pulp-python-admin-extensions-2.0.2-1.fc25.noarch
pulp-python-plugins-2.0.2-1.fc25.noarch
pulp-rpm-admin-extensions-2.15.1-0.2.beta.git.5.b0dfae3.git.5.b0dfae3.git.5.b0dfae3.git.5.b0dfae3.fc25.noarch
pulp-rpm-plugins-2.15.1-0.2.beta.git.5.b0dfae3.git.5.b0dfae3.git.5.b0dfae3.git.5.b0dfae3.fc25.noarch
pulp-selinux-2.15.1-0.2.beta.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.fc25.noarch
pulp-server-2.15.1-0.2.beta.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.fc25.noarch
python-kombu-3.0.33-8.pulp.fc25.noarch
python-pulp-bindings-2.15.1-0.2.beta.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.fc25.noarch
python-pulp-client-lib-2.15.1-0.2.beta.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.fc25.noarch
python-pulp-common-2.15.1-0.2.beta.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.fc25.noarch
python-pulp-deb-common-1.6.0-1.fc25.noarch
python-pulp-docker-common-3.1.1-0.2.beta.git.3.b0dfae3.git.3.b0dfae3.git.3.b0dfae3.git.3.b0dfae3.fc25.noarch
python-pulp-oid_validation-2.15.1-0.2.beta.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.fc25.noarch
python-pulp-ostree-common-1.3.0-1.fc25.noarch
python-pulp-puppet-common-2.15.1-0.2.beta.git.4.b0dfae3.git.4.b0dfae3.git.4.b0dfae3.git.4.b0dfae3.fc25.noarch
python-pulp-python-common-2.0.2-1.fc25.noarch
python-pulp-repoauth-2.15.1-0.2.beta.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.fc25.noarch
python-pulp-rpm-common-2.15.1-0.2.beta.git.5.b0dfae3.git.5.b0dfae3.git.5.b0dfae3.git.5.b0dfae3.fc25.noarch
python-pulp-streamer-2.15.1-0.2.beta.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.git.6.b0dfae3.fc25.noarch

Associated revisions

Revision c180a8c0 - 01/23/2018 04:22 PM - ttereshc

No trailing slash for mirrorlist case

closes #3310 <https://pulp.plan.io/issues/3310>

Revision 41579a99 - 01/23/2018 05:03 PM - ttereshc

No trailing slash for mirrorlist case

closes #3310 <https://pulp.plan.io/issues/3310>

(cherry picked from commit c180a8c09a36118c5df0a412a519eb759c691937)

History

#2 - 01/22/2018 10:52 PM - Ichimonji10

This issue doesn't occur on hosts running the recently released Pulp 2.15. It only occurs on Pulp 2.15 master.

#3 - 01/22/2018 10:55 PM - pcreech

This issue appears to be introduced with the fix for <https://pulp.plan.io/issues/3306>

The above fix appears to always add a trailing slash to the feed url. In the case of mirrorlists, the feed points to a file, not a directory, thus causing the "Error retrieving metadata: Not Found" issue

#4 - 01/23/2018 10:53 AM - ttereshc

- Subject changed from *Pulp can't sync from mirrorlists* to *Pulp can't sync from some mirrorlists*
- Status changed from *NEW* to *ASSIGNED*
- Assignee set to *ttereshc*

I tested several mirrorlists before merging [#3306](#) which I came across and they worked for me.
E.g. http://mirrorlist.centos.org/?release=6&arch=x86_64&repo=os ,
https://mirrors.fedoraproject.org/mirrorlist?repo=fedora-26&arch=x86_64.

I agree with the regression for mirrorlists which one refers to as a file. Do you have examples of such mirrorlists which can be found in the wild (not pulp-fixture ones)?

#5 - 01/23/2018 04:25 PM - ttereshc

- Status changed from *ASSIGNED* to *POST*
- Sprint/Milestone set to *53*

https://github.com/pulp/pulp_rpm/pull/1081

#6 - 01/23/2018 04:31 PM - Ichimonji10

I agree with the regression for mirrorlists which one refers to as a file. Do you have examples of such mirrorlists which can be found in the wild (not pulp-fixture ones)?

I don't have examples offhand. But why should we not support static-file mirrorlists? Should web browsers only support dynamic websites, and not static website?

#7 - 01/23/2018 04:37 PM - ttereshc

I'm not against that, and +1 for keeping the same behaviour we had. I'm just curious, where it's used. Thanks.

#8 - 01/23/2018 04:40 PM - Ichimonji10

I'm not against that, and +1 for keeping the same behaviour we had. I'm just curious, where it's used. Thanks.

I see. I was reading a bit too much into the question. To answer the question more directly: No, I don't have in-the-wild examples. :-/

#9 - 01/23/2018 04:50 PM - dalley

- Triaged changed from *No* to *Yes*

#10 - 01/23/2018 05:02 PM - ttereshc

- Status changed from *POST* to *MODIFIED*

Applied in changeset [c180a8c09a36118c5df0a412a519eb759c691937](#).

#11 - 01/23/2018 05:03 PM - ttereshc

Applied in changeset [41579a996dbbb0779a0ab3ba9bc3c0ab8628e21d](#).

#13 - 02/12/2018 08:58 PM - bmbouter

- Platform Release set to *2.15.2*

#14 - 02/14/2018 09:53 PM - pcreech

- Status changed from *MODIFIED* to *CLOSED - CURRENTRELEASE*
- Platform Release changed from *2.15.2* to *2.15.1*

#15 - 03/09/2018 12:30 AM - bmbouter

- *Sprint set to Sprint 31*

#16 - 03/09/2018 12:31 AM - bmbouter

- *Sprint/Milestone deleted (53)*

#17 - 04/15/2019 10:13 PM - bmbouter

- *Tags Pulp 2 added*