

RPM Support - Issue #2781

Vendor field is present in db if RPM is uploaded, it is missed in case of sync

05/30/2017 04:15 PM - ttereshc

| | |
|---|----------------------------------|
| Status: CLOSED - CURRENTRELEASE | Start date: |
| Priority: High | Due date: |
| Assignee: dalley | Estimated time: 0:00 hour |
| Category: | |
| Sprint/Milestone: 2.17.0 | |
| Severity: 2. Medium | Groomed: No |
| Version: | Sprint Candidate: Yes |
| Platform Release: 2.17.0 | Tags: Pulp 2 |
| OS: | Sprint: Sprint 40 |
| Triaged: Yes | Quarter: |
| Description There is an information about vendor for RPM unit in metadata (primary.xml). In case of upload this information is added to db and thus available via search API. In case of sync vendor is always set to None and it looks more as unfinished work than the intentional removal of data . It is not a regression, it has been this way for years. | |
| Related issues: Related to RPM Support - Task #3924: Add a migration for the "vendor" field CLOSED - CURRENTRELEASE | |

Associated revisions

Revision b1057331 - 07/19/2018 02:54 PM - dalley

Fix missing vendor metadata

closes #2781 <https://pulp.plan.io/issues/2781>

History

#1 - 05/30/2017 04:36 PM - ttereshc

- Priority changed from Normal to High

- Triaged changed from No to Yes

#2 - 06/03/2017 01:31 PM - ttereshc

- Sprint Candidate changed from No to Yes

#4 - 11/02/2017 09:01 PM - Ichimonji10

According to this issue, if one uploads an RPM with vendor metadata to Pulp, then that vendor metadata is available. As far as I can tell, that's not true. If I upload an RPM with vendor metadata to Pulp, then that vendor metadata is **not** available. This can be demonstrated with a simple script:

```
pulp-admin rpm repo create --repo-id foo
wget 'https://repos.fedorapeople.org/pulp/pulp/fixtures/rpm-with-vendor/rpm-with-vendor-1-1.fc25.noarch.rpm'
pulp-admin rpm repo uploads rpm --repo-id foo --file rpm-with-vendor-1-1.fc25.noarch.rpm
pulp-admin rpm repo content rpm --repo-id foo --str-eq 'vendor=Pulp Fixtures'
pulp-admin rpm repo content rpm --repo-id foo --str-eq 'license=Public Domain'
```

Here's some console output from when running the searches:

```
$ pulp-admin rpm repo content rpm --repo-id foo --str-eq 'vendor=Pulp Fixtures'
$ pulp-admin rpm repo content rpm --repo-id foo --str-eq 'license=Public Domain'
```

```
Arch:          noarch
Buildhost:     host-172-16-46-37.openstacklocaldocker-np-ci-rhos-20022.localdo
Checksum:      2f69f4bf6958d4c2281888cdf9df85d87f0117e39ff997c67ec4336d184d25b6
Checksumtype:  sha256
Description:   This RPM has a vendor
Epoch:        0
Filename:      rpm-with-vendor-1-1.fc25.noarch.rpm
License:       Public Domain
Name:          rpm-with-vendor
Provides:      rpm-with-vendor = 1-1.fc25-0
Release:       1.fc25
Requires:      1
Version:       1
```

If I reset Pulp and run a script with a feed URL, the same thing happens:

```
pulp-admin rpm repo create --repo-id bar --feed 'https://repos.fedorapeople.org/pulp/pulp/fixtures/rpm-with-ve
ndor/'
pulp-admin rpm repo sync run --repo-id bar
pulp-admin rpm repo content rpm --repo-id bar --str-eq 'vendor=Pulp Fixtures'
pulp-admin rpm repo content rpm --repo-id bar --str-eq 'license=Public Domain'
```

Here's some console output from when running the searches:

```
$ pulp-admin rpm repo content rpm --repo-id bar --str-eq 'vendor=Pulp Fixtures'
$ pulp-admin rpm repo content rpm --repo-id bar --str-eq 'license=Public Domain'
Arch:          noarch
Buildhost:     host-172-16-46-37.openstacklocaldocker-np-ci-rhos-20022.localdo
Checksum:      2f69f4bf6958d4c2281888cdf9df85d87f0117e39ff997c67ec4336d184d25b6
Checksumtype:  sha256
Description:   This RPM has a vendor
Epoch:        0
Filename:      rpm-with-vendor-1-1.fc25.noarch.rpm
License:       Public Domain
Name:          rpm-with-vendor
Provides:      rpm-with-vendor = 1-1.fc25-0
Release:       1.fc25
Requires:      1
Version:       1
```

This behaviour is consistent for Pulp 2.14 and master (2.15). Here's the packages installed on one of my test systems:

```
[root@rhel-7-pulp-2-14 ~]# rpm -qa | grep pulp | sort
pulp-admin-client-2.14.3-0.1.alpha.git.14.0437abb.el7.noarch
pulp-docker-admin-extensions-3.0.3-0.1.alpha.git.13.1abca9e.el7.noarch
pulp-docker-plugins-3.0.3-0.1.alpha.git.13.1abca9e.el7.noarch
pulp-ostree-admin-extensions-1.3.1-0.1.alpha.git.7.f22fcec.el7.noarch
pulp-ostree-plugins-1.3.1-0.1.alpha.git.7.f22fcec.el7.noarch
pulp-puppet-admin-extensions-2.14.3-0.1.alpha.git.7.c75cc6e.el7.noarch
pulp-puppet-plugins-2.14.3-0.1.alpha.git.7.c75cc6e.el7.noarch
pulp-python-admin-extensions-2.0.3-0.1.alpha.git.6.8ffc1e8.el7.noarch
pulp-python-plugins-2.0.3-0.1.alpha.git.6.8ffc1e8.el7.noarch
pulp-rpm-admin-extensions-2.14.3-0.1.alpha.git.6.8d330c9.el7.noarch
pulp-rpm-plugins-2.14.3-0.1.alpha.git.6.8d330c9.el7.noarch
pulp-selinux-2.14.3-0.1.alpha.git.14.0437abb.el7.noarch
pulp-server-2.14.3-0.1.alpha.git.14.0437abb.el7.noarch
python-isodate-0.5.0-4.pulp.el7.noarch
python-kombu-3.0.33-8.pulp.el7.noarch
python-pulp-bindings-2.14.3-0.1.alpha.git.14.0437abb.el7.noarch
python-pulp-client-lib-2.14.3-0.1.alpha.git.14.0437abb.el7.noarch
python-pulp-common-2.14.3-0.1.alpha.git.14.0437abb.el7.noarch
python-pulp-docker-common-3.0.3-0.1.alpha.git.13.1abca9e.el7.noarch
python-pulp-oid_validation-2.14.3-0.1.alpha.git.14.0437abb.el7.noarch
python-pulp-ostree-common-1.3.1-0.1.alpha.git.7.f22fcec.el7.noarch
python-pulp-puppet-common-2.14.3-0.1.alpha.git.7.c75cc6e.el7.noarch
python-pulp-python-common-2.0.3-0.1.alpha.git.6.8ffc1e8.el7.noarch
```

```
python-pulp-repoauth-2.14.3-0.1.alpha.git.14.0437abb.e17.noarch
python-pulp-rpm-common-2.14.3-0.1.alpha.git.6.8d330c9.e17.noarch
python-pulp-streamer-2.14.3-0.1.alpha.git.14.0437abb.e17.noarch
```

systemctl status doesn't show anything unusual.

#5 - 11/04/2017 02:17 AM - ttereshc

Interesting catch :)

I believe [this commit](#) which landed in 2.14.0 introduced the change of the upload behavior. Upload now works in the same way as sync does. That explains why you no longer can see the vendor field for uploaded RPMs. Before that change upload implementation of parsing rpm metadata was custom and vendor field was imported and saved in DB. So my expectation is that if you try to upload rpm with vendor data using Pulp 2.13, you'll see it.

I'm not sure how to treat it. Is it a regression since Pulp 2.14.0? The feature itself (having vendor information in Pulp) has never been fully implemented (only upload worked this way) and thus it was almost not possible to use it.

Thoughts?

#6 - 07/18/2018 04:36 PM - dkliban@redhat.com

- Sprint set to Sprint 40

#7 - 07/19/2018 04:15 AM - dalley

- Status changed from NEW to ASSIGNED

- Assignee set to dalley

#8 - 07/19/2018 02:56 PM - dalley

- Status changed from ASSIGNED to POST

PR: https://github.com/pulp/pulp_rpm/pull/1138

#9 - 07/19/2018 03:23 PM - dalley

- Status changed from POST to MODIFIED

Applied in changeset [b10573314e28f34f68bfd7e8289f0083acf9e431](#).

#10 - 07/24/2018 03:30 PM - dalley

- Platform Release set to 2.17.0

#11 - 08/07/2018 06:14 PM - ipanova@redhat.com

- Sprint/Milestone set to 2.17.0

#12 - 08/07/2018 06:31 PM - ipanova@redhat.com

- Platform Release changed from 2.17.0 to master

#13 - 08/07/2018 06:31 PM - ipanova@redhat.com

- Platform Release deleted (master)

#14 - 08/07/2018 06:36 PM - ipanova@redhat.com

- Platform Release set to 2.17.0

#15 - 08/14/2018 11:17 PM - ipanova@redhat.com

- Status changed from MODIFIED to 5

#16 - 08/16/2018 03:16 PM - dalley

- Related to Task #3924: Add a migration for the "vendor" field added

#17 - 08/31/2018 10:45 PM - ipanova@redhat.com

- Status changed from 5 to CLOSED - CURRENTRELEASE

#18 - 04/15/2019 10:17 PM - bmbouter

- Tags Pulp 2 added