

## RPM Support - Issue #4152

### Regression Pulp 2.17.1: recursive copy of RPMs does not copy partially resolvable dependencies

11/13/2018 03:18 PM - gmbnomis

<b>Status:</b>	CLOSED - CURRENTRELEASE	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	milan	<b>Estimated time:</b>	0:00 hour
<b>Category:</b>			
<b>Sprint/Milestone:</b>			
<b>Severity:</b>	2. Medium	<b>Groomed:</b>	No
<b>Version:</b>	2.17.1	<b>Sprint Candidate:</b>	No
<b>Platform Release:</b>	2.18.1	<b>Tags:</b>	Pulp 2
<b>OS:</b>		<b>Sprint:</b>	Sprint 46
<b>Triaged:</b>	Yes	<b>Quarter:</b>	

**Description**

### Problem

In contrast to e.g. Pulp 2.15.3, recursive copy in 2.17 seems to ignore dependencies that are not fully resolvable.

Packages often haven dependencies to OS packages (like "bash"). This means that "recursive" only works if the repository actually contains these system dependencies. In our use case, this is not the case (we build application repositories from a "union" of multiple repos using recursive copy).

### Reproducer

```
pulp-admin rpm repo create --repo-id pulp_2_17_stable --feed https://repos.fedorapeople.org/repos/pulp/pulp/stable/2.17/7/x86_64/

pulp-admin rpm repo sync run --repo-id pulp_2_17_stable

pulp-admin rpm repo create --repo-id test

pulp-admin rpm repo copy rpm --from-repo-id pulp_2_17_stable --to-repo-id test --recursive --stream-q=name=pulp-server
```

On Pulp 2.15.3, this yields:

Copied:

```
pulp-selinux-2.17.1-1.el7-noarch
pulp-server-2.17.1-1.el7-noarch
python-bson-3.2-2.el7-x86_64
python-gofer-2.12.1-1.el7-noarch
python-isodate-0.5.0-4.pulp.el7-noarch
python-mongoengine-0.10.5-1.el7-noarch
python-nectar-1.5.6-1.el7-noarch
python-pulp-common-2.17.1-1.el7-noarch
python-pulp-repoauth-2.17.1-1.el7-noarch
python-pymongo-3.2-2.el7-x86_64
python-pymongo-gridfs-3.2-2.el7-x86_64
```

On Pulp 2.17:

Copied:

```
pulp-server-2.17.1-1.el7-noarch
```

Additionally, it logs:

```
pulp_rpm.plugins.importers.yum.pulp_solv:WARNING: [f0572afa] Encountered problems solving: nothing provides /bin/bash needed by pulp-server-0:2.17.1-1.el7.noarch
```

**Related issues:**

Related to RPM Support - Story #3740: Implement modularity content dependency...	<b>CLOSED - CURRENTRELEASE</b>
Related to RPM Support - Task #3528: Investigate which library to use for dep...	<b>CLOSED - COMPLETE</b>
Related to RPM Support - Test #4269: Recursive and conservative recursive copy	<b>CLOSED - COMPLETE</b>
Related to RPM Support - Task #4371: Document how to use the newly added recu...	<b>CLOSED - CURRENTRELEASE</b>

**Associated revisions****Revision cb4d33c5 - 12/17/2018 07:06 PM - milan**

Fake missing dependencies with dummy solvables

This patch allows automatically retrying the dependency solving during recursive copy if missing dependencies were detected.

Opt-in for the conservative recursive copy

This allows one to switch between a conservative and a relaxed (compatible) recursive copy of units (the default). To use the conservative recursive copy, add the "conservative\_recursive\_copy": true tag into the "override\_config" dictionary of the REST API request body.

Fixes: #4152 <https://pulp.plan.io/issues/4152>

**Revision 740dff16 - 01/22/2019 03:35 AM - milan**

Fake missing dependencies with dummy solvables

This patch allows automatically retrying the dependency solving during recursive copy if missing dependencies were detected.

Opt-in for the conservative recursive copy

This allows one to switch between a conservative and a relaxed (compatible) recursive copy of units (the default). To use the conservative recursive copy, add the "conservative\_recursive\_copy": true tag into the "override\_config" dictionary of the REST API request body.

Fixes: #4152 <https://pulp.plan.io/issues/4152>

**History****#1 - 11/13/2018 04:38 PM - CodeHeeler**

- *Triaged changed from No to Yes*

- *Sprint set to Sprint 45*

**#2 - 11/13/2018 09:21 PM - milan**

Pulp (2.17) now uses libsolv as the solver to calculate dependencies when recursively copying units.

This solver is a yes-no solver rather than the best-effort solver we used to have before.

The motivation for the change is discussed on the Issue [#3528](#).

The cleanest solution would be to include multiple repositories in the dependencies calculation, i.e as gmbnomis put it: "given source repo A and target repo B, consider the union of the B, C1, C2... repos as 'installed', where C1, C2... are the repos that would be enabled on the consumer system during DNF installation." Any dependency-solving failures would ideally be propagated back thru the recursive copy tasks result too. This solution would however imply both REST API changes as well as the recursive copy workflow change so it might not be the most welcomed idea for our users. Based on some early investigation of modular depsolving (downstream product implications), similar approach might have to be reused in Issue [#3740](#): given the fact multiple source repositories have to be considered for modular content to resolve, multiple target repositories have to be considered, but the reasons for the multiple target repositories are still unclear to me now.

Another option might be to "fake" missing dependencies with dummy solvables providing those dependencies like DummySolvable(name='/bin/bash', fileprovides=['/bin/bash'], evr="", arch='noarch').

The solver might need to retry solving (multiple times?) before the units can be copied but the solver behavior would remain backwards-compatible which could be preferred by users. This option might be even considered in case of modular depsolving because modular content is kept separate from its dependencies in OS repo content, so the modular repos are "solo-uninstallable" by design.

**#3 - 11/13/2018 09:27 PM - milan**

- *Related to Story #3740: Implement modularity content dependency solving added*

**#4 - 11/13/2018 09:28 PM - milan**

- *Related to Task #3528: Investigate which library to use for dep solving added*

**#5 - 11/14/2018 01:08 PM - milan**

- *Status changed from NEW to ASSIGNED*

- *Assignee set to milan*

**#6 - 11/15/2018 09:58 AM - milan**

- Status changed from ASSIGNED to POST

Just pushed the first take on it; [https://github.com/pulp/pulp\\_rpm/pull/1226](https://github.com/pulp/pulp_rpm/pull/1226)  
gmbnomis, mind taking a look?

Thanks!  
milan

**#7 - 11/29/2018 03:31 PM - rchan**

- Sprint changed from Sprint 45 to Sprint 46

**#8 - 12/17/2018 08:57 PM - milan**

- Status changed from POST to MODIFIED

Applied in changeset [cb4d33c5f998b91b3b42ec551db21d22139dde63](#).

**#9 - 12/18/2018 03:18 PM - kersom**

- Related to Test #4269: Recursive and conservative recursive copy added

**#10 - 01/16/2019 09:09 PM - daviddavis**

- Platform Release set to 2.18.1

**#11 - 01/24/2019 08:17 PM - milan**

Applied in changeset [740dff163e5607180c7cb3cfdc4d296c6cc288ef](#).

**#12 - 01/31/2019 09:54 PM - kersom**

- Related to Task #4371: Document how to use the newly added recursive\_conservative added

**#13 - 02/07/2019 11:00 PM - jortel@redhat.com**

- Status changed from MODIFIED to 5

**#14 - 02/22/2019 11:25 PM - jortel@redhat.com**

- Status changed from 5 to CLOSED - CURRENTRELEASE

**#15 - 04/15/2019 10:07 PM - bmbouter**

- Tags Pulp 2 added