

Debian Support - Issue #4008

python-debpkg does not structure Repositories correctly (in some case)

09/13/2018 06:16 PM - quba42

Status: CLOSED - CURRENTRELEASE	Start date:
Priority: Normal	Due date:
Assignee: quba42	Estimated time: 0:00 hour
Category:	
Sprint/Milestone: 2.19.0	
Severity: 2. Medium	Groomed: No
Version - Debian:	Sprint Candidate: No
Platform Release: 2.19.0	Tags: Pulp 2
Target Release - Debian:	Sprint:
OS:	Quarter:
Triaged: No	

Description

python-debpkg (a dependency of Pulp Debian Support) does not structure Debian repositories quite right:

The Debian repository structure (<https://wiki.debian.org/DebianRepository/Format>) is built from the following schema:

```
<repo_base_url>/dists/<distribution>/<plain_component>/Packages
<repo_base_url>/dists/<distribution>/Release
```

Here, <distribution> may contain additional "/" (like "stretch/updates"). Even were components are listed in the Release file as "updates/main", only the "main" part is the "<plain_component>". (slightly simplified.)

However, the debpkg data structure of the relevant object is built entirely from a <repo_base_url> and a Release file. Therefore, debpkg has no way of knowing what the <distribution> part of the paths should be. It builds the following paths instead:

```
<repo_base_url>/dists/<Releasefile.codename>/<Releasefile.component>/Packages
<repo_base_url>/dists/<Releasefile.codename>/Release
```

For most official debian repos this happens to work since <Releasefile.codename>=<distribution>. However for debian security the <Releasefile.codename> is for example "stretch" and the <distribution> is "stretch/updates". The reason we can still synchronize debian security is because <Releasefile.component> happens to be "update/stretch" while <plain_component> is "stetch". (Two wrongs do make a right...)

However the downstream repository in pulp will have a different structure than the upstream Debian repository, making it impossible to reuse this downstream repository as a new upstream repository. In other words, pulp to pulp synchronization is broken.

Related issues:

Related to Debian Support - Issue #4138: Handling publish paths correctly **CLOSED - CURRENTRELEASE**

Associated revisions

Revision c3831093 - 02/11/2019 09:16 PM - quba42

Account for components with an extra path prefix

The components in the repo metadata object ('Release' file) may be prefixed. This prefix must be stripped from certain comparisons. See also: <https://wiki.debian.org/DebianRepository/Format#Components>

Fixes #4008 <https://pulp.plan.io/issues/4008>

History

#1 - 09/24/2018 01:52 PM - quba42

I now have a pull request to resolve this issue:

https://github.com/pulp/pulp_deb/pull/52

This pull request is in turn dependent on the following changes in python-debpgkr:

<https://github.com/sassoftware/python-debpgkr/pull/13>

#2 - 11/08/2018 06:55 PM - mdellweg

- Assignee set to quba42

#3 - 11/13/2018 12:50 PM - mdellweg

- Related to Issue #4138: Handling publish paths correctly added

#4 - 02/11/2019 09:16 PM - quba42

- Status changed from NEW to MODIFIED

Applied in changeset [c383109308ac44aed452fdf35d770d655764f874](#).

#5 - 03/18/2019 05:59 PM - ttereshc

- Platform Release set to 2.19.0

#6 - 03/18/2019 06:00 PM - ttereshc

- Sprint/Milestone set to 2.19.0

#7 - 03/20/2019 10:02 AM - ttereshc

- Status changed from MODIFIED to 5

#8 - 04/02/2019 10:52 PM - ttereshc

- Status changed from 5 to CLOSED - CURRENTRELEASE

#9 - 04/15/2019 10:08 PM - bmbouter

- Tags Pulp 2 added