

RPM Support - Issue #3020

Distribution files mismatch

09/19/2017 02:16 PM - beny

Status:	CLOSED - WORKSFORME	Start date:	
Priority:	High	Due date:	
Assignee:		Estimated time:	0:00 hour
Category:			
Sprint/Milestone:			
Severity:	2. Medium	Groomed:	No
Version:	2.8.6	Sprint Candidate:	No
Platform Release:		Tags:	
OS:	CentOS 7	Sprint:	
Triaged:	Yes	Quarter:	
Description			
<p>We mirror CentOS repositories in our Pulp and we've hit the issue described in [1] with releases 7.2 and 7.4. It affects all files in the LiveOS and images directories as follows (note the 7.2 and 7.4 symlink pointing to the same file):</p> <pre>:/var/lib/pulp] find ./published/yum/master/yum_distributor/centos-7* -iname vmlinuz -ls 11163912 4 lrwxrwxrwx 1 apache apache 129 Sep 19 12:17 ./published/yum/master/yum_distributor/centos-72/1505816224.61/images/pxeboot/vmlinuz -> /var/lib/pulp/content/units/distribution/de/d64602f1a2165af46f6cd4c8c2645a0422607173fb34617458c9f419e2c7bf/images/pxeboot/vmlinuz 34740815 4 lrwxrwxrwx 1 apache apache 129 Sep 19 13:16 ./published/yum/master/yum_distributor/centos-73/1505819770.72/images/pxeboot/vmlinuz -> /var/lib/pulp/content/units/distribution/8d/d927134b78ab15e3c803b686a6baf564c5ccb97154da48f9ead7041eb53188/images/pxeboot/vmlinuz 11021555 4 lrwxrwxrwx 1 apache apache 129 Sep 19 08:39 ./published/yum/master/yum_distributor/centos-74/1505803143.84/images/pxeboot/vmlinuz -> /var/lib/pulp/content/units/distribution/de/d64602f1a2165af46f6cd4c8c2645a0422607173fb34617458c9f419e2c7bf/images/pxeboot/vmlinuz</pre> <p>Even when the repos are deleted and created again, pulp still thinks the 7.2 and 7.4 distribution files are the same. When downloaded manually from upstream we feed our pulp from [2, 3], the files indeed differ.</p> <pre>md5sum 72/vmlinuz 74/vmlinuz 418a1a0074554d7d7d63b514990c2d1f 72/vmlinuz 43a395b89485ec908e6c0622273299c4 74/vmlinuz</pre> <p>I use a workaround of removing the feed and sync schedule of 7.4 to have up-to-date 7.2 distribution files. This allows me to serve both 7.2 and 7.4 rpms, but I'm unable to mirror both 7.2 and 7.4 distribution files simultaneously</p> <p>[1] https://www.redhat.com/archives/pulp-list/2016-January/msg00003.html [2] http://vault.centos.org/centos/7.2.1511/os/x86_64/ [3] http://mirror.centos.org/centos/7.4.1708/os/x86_64/</p>			

History

#1 - 09/19/2017 04:24 PM - mhrivnak

This comes down to how Pulp defines uniqueness among different "distribution" units. It does so based on these fields taken from the treeinfo file:

- family
- variant
- version
- arch

With RHEL we expect to see versions such as "7.2" or "7.4", but with the CentOS repos you referenced, I see the version as just "7" for all of them.

It also appears that for CentOS 7.3, the "family" is arbitrarily different from 7.2 and 7.4, which would have allowed it to be seen as a unique distribution unit.

Here are the three treeinfo files I'm looking at:

http://vault.centos.org/centos/7.2.1511/os/x86_64/treeinfo

http://mirror.centos.org/centos/7.3.1611/os/x86_64/treeinfo
http://mirror.centos.org/centos/7.4.1708/os/x86_64/treeinfo

This is a general problem that Pulp should find a better way to establish uniqueness for these units, but it could be difficult to make that big of a data model change with Pulp 2.

#2 - 09/19/2017 04:38 PM - ttereshc

- Project changed from Pulp to RPM Support
- Priority changed from Normal to High
- Triaged changed from No to Yes
- Tags Pulp 3 added

Candidate to be fixed in Pulp3

#3 - 09/25/2017 04:46 PM - beny

Thanks for the root cause analysis. Reported the treeinfo value mismatch to CentOS. <https://bugs.centos.org/view.php?id=13904>

#4 - 04/26/2019 10:37 PM - bmbouter

- Tags deleted (Pulp 3)

#5 - 11/05/2020 05:28 PM - ttereshc

- Status changed from NEW to CLOSED - WORKSFORME

Worked in Pulp 3. Feel free to open a bug if you encounter it again.