

RPM Support - Issue #7030

Pulp does not resync yum metadata files on change

06/22/2020 07:02 PM - ttereshc

Status:	CLOSED - CURRENTRELEASE	Start date:	
Priority:	Normal	Due date:	
Assignee:	ppicka	Estimated time:	0:00 hour
Category:		Groomed:	No
Sprint/Milestone:	3.8.0	Sprint Candidate:	No
Severity:	2. Medium	Tags:	
Version:		Sprint:	Sprint 85
Platform Release:		Quarter:	
OS:			
Triaged:	Yes		
Description			
<p>This is for Pulp 3. Since we introduced sync optimizations, the issue we had in Pulp 2 is now also relevant to Pulp 3.</p> <p>Pulp uses revision numbers in repomd.xml to determine if contents need to be updated on sync. However "modifyrepo" does not generate new revision numbers for non rpm data.</p> <p>Steps:</p> <p>1) Setup the following repo</p> <pre>\$ mkdir /tmp/my-data \$ cd /tmp/my-data \$ wget https://partha.fedorapeople.org/test-repos/rpm-with-productid/elephant-0.3-0.8.noarch.rpm \$ createrepo . \$ echo "100000" >> productid \$ modifyrepo --mdtype=productid productid repodata \$ grep revision repodata/repomd.xml <revision>1554217257</revision></pre> <p>2) Sync this repo</p> <p>3) Now update the repo</p> <pre>\$ cd /tmp/my-data \$ echo "100001" >> productid \$ modifyrepo --mdtype=productid productid repodata \$ grep revision repodata/repomd.xml <revision>1554217257</revision></pre> <p>Notice that the revision number did not change even though a metadata file got updated. Try resyncing this change and notice that the productid change will get ignored.</p>			
Related issues:			
Copied from RPM Support - Issue #4627: Pulp does not resync yum metadata file...		CLOSED - CURRENTRELEASE	

Associated revisions

Revision f81d7730 - 11/10/2020 05:07 PM - ppicka

Repomd checksum

Check checksum of repomd file (not only revision) to recognize changes in repository metadata files as createrepo_c doesn't change revision number on change.

closes: #7030 <https://pulp.plan.io/issues/7030>

History

#1 - 06/22/2020 07:02 PM - ttereshc

- Copied from Issue #4627: Pulp does not resync yum metadata files on change added

#2 - 06/22/2020 07:04 PM - ttereshc

- Description updated
- Triaged changed from Yes to No

#3 - 06/22/2020 11:13 PM - dalley

Is this a known bug with modifyrepo? (or modifyrepo_c which is aliased). If not, maybe we should report it.

#4 - 06/25/2020 05:48 PM - ttereshc

- Triaged changed from No to Yes

#5 - 08/03/2020 09:23 PM - ttereshc

It's not considered a bug, and, in theory, user can set an arbitrary number for a revision, so it's safer to rely on the checksum :/

#6 - 10/01/2020 04:49 PM - ttereshc

- Sprint set to Sprint 82

#7 - 10/02/2020 06:03 PM - rchan

- Sprint changed from Sprint 82 to Sprint 83

#8 - 10/09/2020 04:42 PM - ppicka

- Status changed from NEW to ASSIGNED
- Assignee set to ppicka

#9 - 10/15/2020 02:29 PM - pulpbot

- Status changed from ASSIGNED to POST

PR: https://github.com/pulp/pulp_rpm/pull/1874

#10 - 10/15/2020 11:48 PM - rchan

- Sprint changed from Sprint 83 to Sprint 84

#11 - 10/30/2020 04:56 PM - rchan

- Sprint changed from Sprint 84 to Sprint 85

#12 - 11/10/2020 05:07 PM - ppicka

- Status changed from POST to MODIFIED

Applied in changeset [f81d7730ba6c4c7a3783300fad6949ce8bdbca5a](#).

#13 - 11/13/2020 01:06 AM - dalley

- Sprint/Milestone set to 3.8.0

#14 - 11/13/2020 07:22 PM - pulpbot

- Status changed from MODIFIED to CLOSED - CURRENTRELEASE