

RPM Support - Issue #3957

on_demand repositories should not allow remove_missing to be false

08/30/2018 07:21 PM - jsherril@redhat.com

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

Description

on_demand should always use remove_missing=true because:

- 1) if a package is removed upstream the on_demand repo won't be able to download it
- 2) if the checksum of the repo changes, all the old checksum'd packages will still be in the repo and won't be able to be downloaded, nor will the repodata be able to be generated

History

#1 - 08/31/2018 09:55 PM - ttereshc

I agree with the raised concern in general.
I'm not sure about details, so I'd like to discuss it here.

1. Combining repos won't be possible in case of on_demand policy.

Sync one remote repo, then change a feed and sync another remote repo into the same pulp repository. I think it's not a very rare use case for a one-time sync

I suggest at least to document alternatives if we are taking this use case from users (e.g. sync into two different pulp repos and then copy into one, original repos can be resynced on a regular basis)

2. Switching from immediate to on_demand when remove_missing is disabled.

Should this case be disallowed? - Likely, yes. The old content is all downloaded, no worries there, but later the same problems can be encountered with on_demand content.

Does Katello switch policies back and forth during a repair of RPM content? Would that feature be affected by this restriction?

3. Importer option remove_missing is disabled by default

on_demand should always use remove_missing=true

If we decide to go with the proposal, I think the remove_missing option should be set explicitly on the importer to avoid confusion that with on_demand policy a default for that option is different.

4. What to do with existing on_demand repos with remove_missing=false?

Keep it as-is but complain during sync that remove_missing should be set. Thoughts?

5. Alternatives.

Pulp can't guarantee that URL (to download the content) is up to date. Even with the remove_missing enabled, there is always a time when a remote repo has been changed but hasn't yet been synced into Pulp. This state is fixed with the upcoming sync.

So alternatively we can document that to fix a repo, one should enable remove_missing on the importer. Downside: the broken state is not fixed automatically. Benefit: old behaviour and use cases are preserved.

#2 - 09/07/2018 04:37 PM - CodeHeeler

- Triaged changed from No to Yes

#3 - 04/12/2019 10:20 PM - bmbouter

- Status changed from NEW to CLOSED - WONTFIX

Pulp 2 is approaching maintenance mode, and this Pulp 2 ticket is not being actively worked on. As such, it is being closed as WONTFIX. Pulp 2 is still accepting contributions though, so if you want to contribute a fix for this ticket, please reopen or comment on it. If you don't have permissions to reopen this ticket, or you want to discuss an issue, please reach out via the [developer mailing list](#).

#4 - 04/15/2019 10:09 PM - bmbouter

- *Tags Pulp 2 added*