RPM Support - Story #8673

Auto-publishing should be more fault-tolerant

04/30/2021 04:56 PM - sskracic@redhat.com

Status:	NEW	**Start date:**	
Priority:	Low	**Due date:**	
Assignee:		**% Done:**	0%
Category:		**Estimated time:**	0:00 hour
Sprint/Milestone:			
Platform Release:		**Tags:**	
Groomed:	No	**Sprint:**	Sprint 110
Sprint Candidate:	No	**Quarter:**	

Description

I admit the title is a bit vague.

During auto-publishing sync of a very large repository (rhel-7-server-rpms), the rq process got killed by oom-killer sometime in the middle of the publishing step. So the new repository version (1) got created, but accompanying publication did not.

On the subsequent sync runs, the repository did not get published, as no new content was available to sync, hence a new repository version was not created, which in turn should trigger publication and distribution update.

Of course, the repo can still be published and distributed using the non-autopublishing REST API, but I wonder whether the behavior was intended.

Related issues:

Related to Pulp - Task #8740: [EPIC] Publication based plugins should use eit...	**NEW**
Related to RPM Support - Story #6353: As a user, I can mirror RPM repository ...	**CLOSED - CURRENTRELEASE**

History

#1 - 05/12/2021 11:07 PM - dalley

- Related to Task #8740: [EPIC] Publication based plugins should use either `distribution.repository` or `distribution.publication` but not both added

#2 - 05/13/2021 03:38 PM - dalley

- Related to Story #6353: As a user, I can mirror RPM repository content and metadata added

#3 - 05/13/2021 03:40 PM - dalley

#6353 should help with this, since metadata mirroring avoids the need for a publishing step - even via auto-publishing.

Although of course for custom repos, auto-publish is necessary to make repo modifications convenient.

#4 - 05/27/2021 04:49 AM - dalley

@sskracic I would call this "an expected but unfortunate corner case". In this situation, which would you prefer to happen?

- The current behavior (with maybe some additional logging of warnings)
- Have entire task fail and all changes reverted, including the destruction of the repository version that was created. This means that the operation that created the repository version would need to be re-done in its entirety.

#5 - 05/27/2021 03:20 PM - sskracic@redhat.com

dalley wrote:

> @sskracic I would call this "an expected but unfortunate corner case". In this situation, which would you prefer to happen?
>
> - The current behavior (with maybe some additional logging of warnings)
> - Have entire task fail and all changes reverted, including the destruction of the repository version that was created. This means that the operation that created the repository version would need to be re-done in its entirety.

If feasible, I would prefer the latter. My understanding is that with the artifacts already being downloaded, only their associations to the new repo version will need to be redone on the next sync run.

#6 - 07/21/2021 09:41 PM - dalley

- Sprint set to Sprint 101

#7 - 08/02/2021 07:44 PM - ipanova@redhat.com

- Sprint changed from Sprint 101 to Sprint 102

#8 - 08/12/2021 05:23 PM - rchan

- Sprint changed from Sprint 102 to Sprint 103

#9 - 08/27/2021 05:08 PM - rchan

- Sprint changed from Sprint 103 to Sprint 104

#10 - 09/10/2021 12:28 AM - rchan

- Sprint changed from Sprint 104 to Sprint 105

#11 - 09/11/2021 04:30 AM - dalley

- Priority changed from Normal to Low

#12 - 09/23/2021 11:54 PM - rchan

- Sprint changed from Sprint 105 to Sprint 106

#13 - 10/08/2021 03:16 PM - rchan

- Sprint changed from Sprint 106 to Sprint 107

#14 - 10/21/2021 06:35 PM - rchan

- Sprint changed from Sprint 107 to Sprint 108

#15 - 11/04/2021 10:21 PM - rchan

- Sprint changed from Sprint 108 to Sprint 109

#16 - 11/19/2021 09:39 PM - rchan

- Sprint changed from Sprint 109 to Sprint 110