Debian Support - Issue #9277

Katello::Errors::Pulp3Error: value too long for type character varying(255)

08/23/2021 06:10 PM - lfu

Status:	CLOSED - CURRENTRELEASE	**Start date:**	
Priority:	Normal	**Due date:**	
Assignee:		**Estimated time:**	0:00 hour
Category:			
Sprint/Milestone:	2.16.0		
Severity:	2. Medium	**Groomed:**	No
Version - Debian:		**Sprint Candidate:**	No
Platform Release:		**Tags:**	
Target Release - Debian:		**Sprint:**	
OS:		**Quarter:**	
Triaged:	No		
Description			
https://community.theforeman.org/t/katello-value-too-long-for-type-character-varying-255/24953			

Associated revisions

Revision 00fe4226 - 10/28/2021 01:21 PM - quba42

Allow more than 256 characters for certain DB fields

closes #9277 https://pulp.plan.io/issues/9277

The relevant fields containing a whitespace seperated list of values, in particular any fields named "distributions", "components", or "architectures".

History

#1 - 08/24/2021 04:37 PM - fao89

- Project changed from Pulp to Debian Support

#2 - 10/26/2021 12:34 PM - quba42

I went and had a look at all fields in pulp_deb that have max_length=255 and I don't think any of them are likely to have caused this sync failure. (Given what I know about the repo being synced).

Without the ability to test the upstream repo, I do not know how to proceed...

#3 - 10/26/2021 02:43 PM - quba42

Ok, I was given the relevant release file, and the problem is that the repo has uncharacteristically many components, thus exceeding the max length of the field in pulp_deb. The following is from the upstream release file:

```
Components: main mysql postgresql ldap cloud-s3 cloud-glacier cloud-azure cloud-google cloud-oracle shstore de
lta ndmp vsphere demo demo-support bweb netapp-hfc html-manuals kvm snapshot dedup aligned single-item-restore
 sybase kubernetes docker rhv db2 xenserver hdfs m365
```

#4 - 10/26/2021 04:23 PM - quba42

Several possibilities for improving components fields:

- Increase the max_length?
- Chane to TextField(null=True)?
- Change to ArrayField(CharField(max_length=255), null=True)?

Also update architectures and distributions fields?

#5 - 10/28/2021 01:19 PM - quba42

I decided against the array option, since I want to do as little as possible processing of upstream metadata, before storing it in the DB. In other words, I want to keep upstream fields containing space separated string lists as a single string (containing space separated list values) in the DB.

#6 - 10/28/2021 01:24 PM - pulpbot

- Status changed from NEW to POST

PR: https://github.com/pulp/pulp_deb/pull/364

#7 - 10/28/2021 04:00 PM - quba42

- Status changed from POST to MODIFIED

Applied in changeset 00fe4226e8b79bc8de335414963d25a0fbb23672.

#8 - 10/28/2021 04:39 PM - quba42

- Sprint/Milestone set to 2.16.0

#9 - 10/28/2021 05:00 PM - pulpbot

- Status changed from MODIFIED to CLOSED - CURRENTRELEASE