

RPM Support - Issue #4741

Migration of modular content fails for very old Fedora repos (upgrade from <=2.16 to 2.17+)

04/26/2019 01:45 PM - ttereshc

Status:	CLOSED - CURRENTRELEASE	Start date:	
Priority:	Normal	Due date:	
Assignee:	ttereshc	Estimated time:	0:00 hour
Category:		Sprint Candidate:	No
Sprint/Milestone:	2.19.1	Tags:	Pulp 2
Severity:	2. Medium	QA Contact:	
Version:	2.17.0	Complexity:	
Platform Release:	2.19.1	Smash Test:	
Blocks Release:		Verified:	No
OS:		Verification Required:	No
Backwards Incompatible:	No	Sprint:	Sprint 52
Triaged:	Yes		
Groomed:	No		
Description			
<p>The error is seen for repos with f26 content, those repos don't exist any more: http://modularity.fedorainfracloud.org/modularity/compat/fedora-os-stream/ http://modularity.fedorainfracloud.org/modularity/compat/fedora-modular/</p> <pre>ValidationError (Modulemd:938e503c-9913-48e0-8ebd-c68791db80d0) (Field is required: ['context']) Traceback (most recent call last): File "/usr/lib/python2.7/site-packages/pulp/server/db/manage.py", line 239, in main return _auto_manage_db(options) File "/usr/lib/python2.7/site-packages/pulp/server/db/manage.py", line 306, in _auto_manage_db migrate_database(options) File "/usr/lib/python2.7/site-packages/pulp/server/db/manage.py", line 135, in migrate_database update_current_version=not options.test) File "/usr/lib/python2.7/site-packages/pulp/server/db/migrate/models.py", line 189, in apply_migration migration.migrate() File "/usr/lib/python2.7/site-packages/pulp_rpm/plugins/migrations/0500_upstream_0043_add_modulemd_modulemd-defaults.py", line 343, in migrate add_modulemds(repository, modulemds, working_dir) File "/usr/lib/python2.7/site-packages/pulp_rpm/plugins/migrations/0500_upstream_0043_add_modulemd_modulemd-defaults.py", line 218, in add_modulemds add_modulemd(repository, modulemd, model, working_dir) File "/usr/lib/python2.7/site-packages/pulp_rpm/plugins/migrations/0500_upstream_0043_add_modulemd_modulemd-defaults.py", line 164, in add_modulemd model.save_and_import_content(path) File "/usr/lib/python2.7/site-packages/pulp/server/db/model/__init__.py", line 921, in save_and_import_content self.save() File "/usr/lib/python2.7/site-packages/mongoengine/document.py", line 304, in save self.validate(clean=clean) File "/usr/lib/python2.7/site-packages/mongoengine/base/document.py", line 413, in validate raise ValidationError(message, errors=errors) ValidationError: ValidationError (Modulemd:938e503c-9913-48e0-8ebd-c68791db80d0) (Field is required: ['context'])</pre>			

Modules in those repos have no context which is a part of unit_key for a modulemd content so it cant' be imported into Pulp.

Possible solutions:

- provide list of content for user to manually remove before migration
- skip this broken content and proceed with migration
 - skipping it will mean that this content will be removed from a repo and become orphaned.
- anything else? I'm hesitant to keep bad modules.yaml files as unparsed yum_repo_metadata_file content for consistency reasons.

No data how many users might be affected. Hopefully not many since this issue didn't blow up earlier.

Associated revisions

Revision e43c39ee - 04/29/2019 09:24 AM - ttereshc

Skip malformed modular content during migration

Some users might have old experimental modular content in Pulp. Such content doesn't contain all the necessary info to be imported in Pulp. E.g. no `context` which is a part of unit_key.

closes #4741

<https://pulp.plan.io/issues/4741>

Revision 78a47006 - 05/16/2019 08:01 PM - ttereshc

Skip malformed modular content during migration

Some users might have old experimental modular content in Pulp. Such content doesn't contain all the necessary info to be imported in Pulp. E.g. no `context` which is a part of unit_key.

closes #4741

<https://pulp.plan.io/issues/4741>

(cherry picked from commit e43c39ee16d1787df587024cd8dfb077b11a8f8d)

History

#1 - 04/26/2019 01:46 PM - ttereshc

- Description updated

#2 - 04/26/2019 01:54 PM - ttereshc

This patch is for skipping broken content.

```
diff --git a/plugins/pulp_rpm/plugins/migrations/0043_add_modulemd_modulemd-defaults.py b/plugins/pulp_rpm/plu
gins/migrations/0043_add_modulemd_modulemd-defaults.py
index 20612c57..17bd9fd4 100644
--- a/plugins/pulp_rpm/plugins/migrations/0043_add_modulemd_modulemd-defaults.py
+++ b/plugins/pulp_rpm/plugins/migrations/0043_add_modulemd_modulemd-defaults.py
@@ -6,7 +6,7 @@ from uuid import uuid4
     from hashlib import sha256
```

```

import bson
-from mongoengine import Q, NotUniqueError
+from mongoengine import Q, NotUniqueError, ValidationError

from pulp.plugins.loader import api as plugin_api
from pulp.server.controllers import repository as repository_controller
@@ -164,6 +164,11 @@ def add_modulemd(repository, modulemd, model, working_dir):
    model.save_and_import_content(path)
    except NotUniqueError:
        model = Modulemd.objects.get(**model.unit_key)
+    except ValidationError as e:
+        _logger.warn("Modulemd content import failure. Repo: %s, content: %s, "
+            "error: %s" % (repository.repo_id, model.unit_key, e))
+    return
+
    repository_controller.associate_single_unit(repository, model)

```

NOTE: it prints the message for every bad modulemd and user sees it in the output of pulp-manage-db.

#3 - 04/26/2019 06:17 PM - ttereshc

- Description updated

#4 - 04/26/2019 06:36 PM - dkliban@redhat.com

We should just skip the content and let it get orphaned. We should provide release notes that clearly state that content from F26 will stop working after the migration runs.

#5 - 04/29/2019 09:30 AM - ttereshc

- Status changed from NEW to POST

- Triaged changed from No to Yes

- Sprint set to Sprint 52

https://github.com/pulp/pulp_rpm/pull/1333

#6 - 04/29/2019 09:30 AM - ttereshc

- Assignee set to ttereshc

#7 - 04/30/2019 11:34 AM - ttereshc

- Status changed from POST to MODIFIED

Applied in changeset [e43c39ee16d1787df587024cd8dfb077b11a8f8d](#).

#8 - 05/02/2019 02:55 PM - dkliban@redhat.com

- Platform Release set to 2.19.1

#9 - 05/03/2019 09:33 PM - dkliban@redhat.com

- Sprint/Milestone set to 2.19.1

#10 - 05/21/2019 02:15 PM - ttereshc

Applied in changeset [78a470060a84c6be5c19fc90800f1f82dcfcedad](#).

#11 - 05/22/2019 09:26 PM - dkliban@redhat.com

- Status changed from MODIFIED to ON_QA

#12 - 05/30/2019 06:42 PM - dkliban@redhat.com

- Status changed from ON_QA to CLOSED - CURRENTRELEASE