

RPM Support - Story #7806

Support publish of v3 modulemd documents in Pulp 2.x

11/09/2020 05:30 AM - rmcgover

Status:	CLOSED - WONTFIX	Start date:	
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0:00 hour
Sprint/Milestone:		Tags:	Pulp 2
Platform Release:		Sprint:	
Groomed:	No	Quarter:	
Sprint Candidate:	No		

Description

In libmodulemd, it is planned to introduce a version 3 of modulemd documents which differ slightly from the current v2.

This is a request to have Pulp 2.x support upload and publishing of these documents within a yum repo, similarly as v2 documents can be uploaded and published today.

Must have

- It must remain possible to upload & publish a YAML stream of modulemd v2 documents with identical behavior as seen today.
- It must be possible to upload & publish a YAML stream of (modulemd/modulemd-stream) v3 documents, with same behavior as for v2 documents (i.e. uploaded YAML stream is stored, and reproduced verbatim - while concatenated with all others in the same repo - at publish time).

Nice to have

- It would be nice if v3 documents are saved as "modulemd" units (i.e. not introducing a new type of unit).
- It would be nice to avoid any changes to the unit key on "modulemd" units.

Not required

- We don't require any kind of migration or upgrades between v2 and v3 document types.
- We don't require support for the modulemd-packager document type.
- We don't require support for storing both (v2 and v3) documents for the same module (NSVCA), across the entire Pulp installation.

Additional info

Most up-to-date info on modulemd v3 documents at time of writing was:
https://github.com/fedora-modularity/libmodulemd/blob/b7bb222030b7c31359f8f7940d7540b644b019ab/yaml_specs/modulemd_stream_v3.yaml

This issue was written only to cover the use-cases relating to "publish modulemd documents to CDN". There may be other work required for modulemd v3 documents which is not described here.

History

#1 - 11/24/2020 01:10 PM - ttereshc

Do you use copy operation for modulemd documents? or do you just upload them? If you use copy, it's a simple one and not a recursive, correct?

Do you need to have a field to know the version of modulemd document? E.g. to search by it.

#2 - 11/25/2020 06:23 AM - rmcgover

We do copy/associate modulemd units and we don't need recursive copy.

I'm not aware of any requirement to have the modulemd version stored as a field.

#3 - 01/08/2021 06:23 PM - dalley

So we need to modify the fixtures to include this new "static context: bool" bit of data. I gather that this is each individual snippet of modulemd metadata that needs to be tagged this way? Do we want a mix of (no static context), (static context = true), (static context = false) or just flipped one way?

#4 - 01/08/2021 06:33 PM - ipanova@redhat.com

dalley wrote:

So we need to modify the fixtures to include this new "static context: bool" bit of data. I gather that this is each individual snippet of modulemd metadata that needs to be tagged this way? Do we want a mix of (no static context), (static context = true), (static context = false) or just flipped one way?

I think we should test module snippets that 1) do not contain static context 2) contain static context. Both snippets upload, publish and confirm no errors are encountered and both modules are consumable by dnf client.

#5 - 01/08/2021 06:36 PM - ttereshc

It's most likely that for this specific request we need to be sure that everything works with no static_context at all (like the modules we have now), or with static_context set to true. That would be my expectation.

However, to be on the safe side, I'm +1 to have a repo with no static_context at all and another repo with a mixture of all 3 that you listed.

#6 - 01/15/2021 07:21 AM - dalley

Summary of the problems:

- Pulp 2 is actually using libmodulemd1, as in the literal libmodulemd1 package:
https://github.com/pulp/pulp_rpm/blob/2-master/plugins/pulp_rpm/plugins/importers/yum/repomd/modules.py#L22
 - GObject basically hides the source of the library behind the "require version" import. The libmodulemd1 package provides 1.0, the libmodulemd-2 package provides 2.0
- Libmodulemd1 will not parse module metadata containing the static_context key (the value does not matter). So these modules will be silently ignored and not uploaded or synced.
- v2 makes some breaking changes to the API and we hit at least one of them, presumably because the original function signatures don't provide any errors (which is why it happens silently). Probably however there is a nearly identical function that *does* return the errors:
<https://github.com/fedora-modularity/libmodulemd/pull/58>

I didn't really go further than this, so perhaps there are more issues or perhaps not. But Pulp 2 won't work "as-is" with any module metadata containing the static context key.

#7 - 04/27/2021 08:01 PM - dalley

- Status changed from NEW to CLOSED - WONTFIX

Addressed by EXD patches.