

## RPM Support - Issue #6214

### Module dependency metadata is stored in an overly-simplified format

02/20/2020 11:31 PM - dalley

<b>Status:</b>	CLOSED - CURRENTRELEASE	<b>Start date:</b>	
<b>Priority:</b>	High	<b>Due date:</b>	
<b>Assignee:</b>	ppicka	<b>Estimated time:</b>	0:00 hour
<b>Category:</b>		<b>Groomed:</b>	No
<b>Sprint/Milestone:</b>	Pulp RPM 3.2.0	<b>Sprint Candidate:</b>	No
<b>Severity:</b>	2. Medium	<b>Tags:</b>	
<b>Version:</b>		<b>Sprint:</b>	Sprint 67
<b>Platform Release:</b>		<b>Quarter:</b>	
<b>OS:</b>			
<b>Triaged:</b>	Yes		

#### Description

The dependency segment of the module metadata is stored in the database as follows:

```
"dependencies": {  
  "ant": [  
    "1.10"  
  ],  
  "maven": [  
    "3.5"  
  ],  
  "platform": [  
    "f31"  
  ]  
},
```

(example from the eclipse module from fedora 31 modular release repo)

I *think* this is over-simplified slightly. I think module stream is allowed to have different dependencies on different platforms. The "platform" key of the dictionary is used as a tag for the dictionary of dependencies, meaning "the dependencies listed on this dict apply to each of these platforms".

This means that the outer level of the dependencies should be a list, not a dict. This is the way it is stored in Pulp 2.

For this same module, Pulp 2 stores it as:

```
"dependencies": [
  {
    "maven": [
      "3.5"
    ],
    "platform": [
      "f31"
    ],
    "ant": [
      "1.10"
    ]
  },
]
```

I *think* that flattening out the dependencies as we have done is incorrect. It's difficult to say for sure because I can't find any modules that have different dependencies on different platforms, or even different values for platform to begin with.

In any case, part of the Pulp 2 depsolving algorithm (which I am copying to Pulp 3) is built around this format, and it was originally copied from Module Build Service code, which also works this way. But since a data migration to un-flatten this might be painful (depending on whether multi-platform module streams actually exist), and changing the algorithm is probably unwise, we should pause and try to figure out how exactly the metadata is supposed to work (if anybody knows).

---

## Associated revisions

### Revision 50b410ae - 02/26/2020 05:42 PM - ppicka

Modulemd dependencies fix

Modulemd dependencies is now stored correctly in database. Django migration for dependencies added.

closes: #6214 <https://pulp.plan.io/issues/6214>

---

## History

### #1 - 02/21/2020 10:47 AM - ttereshc

I agree with you that the structure is oversimplified and it doesn't cover the case when there are different dependencies for different platforms.

Here are some examples for different platforms:

- Docs describing that (see "A complex example"):  
[https://docs.fedoraproject.org/en-US/modularity/making-modules/defining-modules/#\\_advanced\\_dependencies\\_optional](https://docs.fedoraproject.org/en-US/modularity/making-modules/defining-modules/#_advanced_dependencies_optional)
- From our fixtures <https://repos.fedorapeople.org/pulp/pulp/fixtures/rpm-with-modules/>:

### module duck 0 (0.8)

```
dependencies:
- buildrequires:
  walrus: [0.71, 5.21]
  platform: [f28]
  requires:
  walrus: [0.71, 5.21]
  platform: [f28]
- buildrequires:
  platform: [el8]
  requires:
  platform: [el8]
```

Notice that **module duck 0 (0.7)** has also multiple platforms but they have the same set of dependencies:

```
dependencies:
- buildrequires:
  walrus: []
  platform: [-f29, -f30]
  requires:
  walrus: []
  platform: [-f29, -f30]
```

I hope migration is not terribly painful. We can't unflatten the structure based on the data from DB only if there are multiple platforms specified. But in such cases we can look at its artifact where original yaml snippet is stored and just recreate this field.

### #2 - 02/21/2020 03:22 PM - dalley

- Triaged changed from No to Yes
- Sprint set to Sprint 67

**#3 - 02/24/2020 05:05 PM - ppicka**

- Status changed from NEW to ASSIGNED
- Assignee set to ppicka

**#4 - 02/24/2020 05:53 PM - ppicka**

- Status changed from ASSIGNED to POST

[https://github.com/pulp/pulp\\_rpm/pull/1611](https://github.com/pulp/pulp_rpm/pull/1611)

**#5 - 02/26/2020 05:42 PM - ppicka**

- Status changed from POST to MODIFIED

Applied in changeset [pulp\\_rpm:50b410ae2b8dbeb2c3bd70f08f04b29e82d51e7c](#).

**#6 - 03/02/2020 01:43 PM - ttereshc**

- Project changed from Pulp to RPM Support
- Status changed from MODIFIED to CLOSED - CURRENTRELEASE
- Sprint/Milestone set to Pulp RPM 3.2.0