

## Migration Plugin - Issue #5374

### re-run of migration tooling has incorrect counts

08/30/2019 04:08 PM - jsherril@redhat.com

<b>Status:</b> NEW	<b>Start date:</b>
<b>Priority:</b> Low	<b>Due date:</b>
<b>Assignee:</b>	<b>Estimated time:</b> 0:00 hour
<b>Category:</b>	
<b>Sprint/Milestone:</b>	
<b>Severity:</b> 2. Medium	<b>Sprint Candidate:</b> No
<b>Platform Release:</b>	<b>Tags:</b> Katello
<b>OS:</b>	<b>Sprint:</b>
<b>Triaged:</b> Yes	<b>Quarter:</b>
<b>Groomed:</b> No	

**Description**

I synced a repo with ~14K files in pulp2 and then migrated all content to pulp3:

```
{
  "_created": "2019-08-29T18:59:50.347635Z",
  "_href": "/pulp/api/v3/tasks/1a1a93c2-dfd6-4ba6-b54e-c74735a3ce5d/",
  "created_resources": [],
  "error": null,
  "finished_at": "2019-08-29T19:01:37.544664Z",
  "name": "pulp_2to3_migrate.app.tasks.migrate.migrate_from_pulp2",
  "non_fatal_errors": [],
  "parent": null,
  "progress_reports": [
    {
      "done": 14444,
      "message": "Pre-migrating Pulp 2 ISO content (general info)",
      "state": "completed",
      "suffix": null,
      "total": 14447
    },
    {
      "done": 14444,
      "message": "Pre-migrating Pulp 2 ISO content (detail info)",
      "state": "completed",
      "suffix": null,
      "total": 14447
    },
    {
      "done": 14051,
      "message": "Migrating ISO content to Pulp 3",
      "state": "completed",
      "suffix": null,
      "total": 14453
    },
    {
      "done": 14453,
      "message": "Migrating content to Pulp 3",
      "state": "completed",
      "suffix": null,
      "total": 14453
    }
  ],
  "reserved_resources_record": [
    "pulp_2to3_migration"
  ],
}
```

```
"spawned_tasks": [],
"started_at": "2019-08-29T18:59:50.474079Z",
"state": "completed",
"worker": "/pulp/api/v3/workers/75176664-9633-46f5-8989-456923ba333a/"
}
```

the task counts here looked correct. I then re-ran the migration without any changes, and it looked like this:

```
{
  "_created": "2019-08-29T19:24:08.615486Z",
  "_href": "/pulp/api/v3/tasks/eee0440a-0d46-48f8-8d3b-2bf711b7b0c4/",
  "created_resources": [],
  "error": null,
  "finished_at": "2019-08-29T19:24:11.410935Z",
  "name": "pulp_2to3_migrate.app.tasks.migrate.migrate_from_pulp2",
  "non_fatal_errors": [],
  "parent": null,
  "progress_reports": [
    {
      "done": 390,
      "message": "Migrating ISO content to Pulp 3",
      "state": "completed",
      "suffix": null,
      "total": 402
    },
    {
      "done": 402,
      "message": "Migrating content to Pulp 3",
      "state": "completed",
      "suffix": null,
      "total": 402
    },
    {
      "done": 0,
      "message": "Pre-migrating Pulp 2 ISO content (general info)",
      "state": "completed",
      "suffix": null,
      "total": 0
    },
    {
      "done": 0,
      "message": "Pre-migrating Pulp 2 ISO content (detail info)",
      "state": "completed",
      "suffix": null,
      "total": 0
    }
  ],
  "reserved_resources_record": [
    "pulp_2to3_migration"
  ],
  "spawned_tasks": [],
  "started_at": "2019-08-29T19:24:08.732119Z",
  "state": "completed",
  "worker": "/pulp/api/v3/workers/75176664-9633-46f5-8989-456923ba333a/"
}
```

the counts for that first progress\_report appears incorrect

## History

#1 - 08/30/2019 04:34 PM - amacdona@redhat.com

- Triaged changed from No to Yes
- Sprint set to Sprint 58

## #2 - 08/30/2019 04:44 PM - ttereshc

Just for the context:

Migration plugin looks at the last\_updated time and chooses what to migrate based on that. last\_updated precision is 1 sec.

Since it can be re-run N times on a working system, it's possible that some content from that second has been migrated and some has not.

So only in the process of migration we figure out if something was migrated before or not. It's likely some counters need to be adjusted accordingly.

It's done for the pre-migration part, maybe the later parts need the same.

Maybe it's a wrong guess ^ but the context might be useful.

However, +1 the first report also shows incorrect counters.

## #3 - 09/17/2019 09:02 PM - rchan

- Sprint deleted (Sprint 58)

## #4 - 01/14/2020 05:53 PM - ttereshc

- Project changed from Pulp to Migration Plugin

## #5 - 01/25/2020 01:00 AM - dalley

At least for the first report, it no longer seems to be an issue, at least using the 4 standard test ISO repos. I can't vouch for the more comprehensive repos Justin is using.

The second time around, the numbers are still different.

```
{
  "code": "premigrating.content.detail",
  "done": 0,
  "message": "Pre-migrating Pulp 2 ISO content (detail info)",
  "state": "completed",
  "suffix": null,
  "total": 0
},
{
  "code": "migrating.content",
  "done": 253,
  "message": "Migrating content to Pulp 3",
  "state": "completed",
  "suffix": null,
  "total": 253
},
{
  "code": "migrating.iso.content",
  "done": 266,
  "message": "Migrating iso content to Pulp 3 iso",
  "state": "completed",
  "suffix": null,
  "total": 266
},
{
  "code": "migrating.distributors",
  "done": 0,
  "message": "Migrating distributors to Pulp 3",
  "state": "completed",
  "suffix": null,
  "total": 4
}
```

**#6 - 05/08/2020 07:34 PM - ggaaney**

- *Priority changed from Normal to Low*

**#7 - 05/08/2020 07:44 PM - ggaaney**

- *Tags Katello added*

- *Tags deleted (Katello-P3)*

**#8 - 06/09/2020 03:30 PM - ttereshc**

- *Sprint/Milestone set to 0.2.0*

**#9 - 06/11/2020 04:34 PM - ttereshc**

Needs to be retested

**#10 - 08/20/2020 06:32 PM - ttereshc**

- *Sprint/Milestone deleted (0.2.0)*