

Katello - Bug #23322

[Performance improvement] On CV publish/promote, call Pulp::Repository::* dynflow steps concurrently

04/19/2018 01:03 AM - Justin Sherrill

| | |
|---|---------------------------|
| Status: Rejected | |
| Priority: Normal | |
| Assignee: Samir Jha | |
| Category: Content Views | |
| Target version: Katello Recycle Bin | |
| Difficulty: | Fixed in Releases: |
| Triaged: Yes | Found in Releases: |
| Bugzilla link: 1554282 | Red Hat JIRA: |
| Pull request: https://github.com/Katello/katello/pull/7343 | |
| Description | |
| Cloned from https://bugzilla.redhat.com/show_bug.cgi?id=1554282 | |
| Description of problem: | |
| When publishing a CV or promoting it to a LE, a content of a pulp repo (or two repos, sequentially) is copied to another one. That is achieved via Pulp::Repository::* dynflow steps that each copies a different type of content units (rpm, errata, srpm,..). | |
| Currently, these dynflow steps are run sequentially, so the call flow is: | |
| - trigger Pulp::Repository::CopySrpm | |
| - wait until it completes - by periodical checking of the pulp task status | |
| - after it completes, trigger Pulp::Repository::CopyRpm | |
| - wait until it completes | |
| - trigger .. | |
| .. | |
| These dynflow steps can be executed concurrently - pulp engineering confirmed me this is safe to trigger those pulp tasks concurrently. Internally in pulp, they will be performed sequentially either way, but simply we dont require the extra "locking" on dynflow steps level. Since the waiting for task completion adds redundant delay too many times. | |
| Preliminary tests showed observable-to-significant improvement in CV publish time, but these will definitely vary on many factors (i.e. # and size of repos in the CV, # of pulp workers, other dynflow tasks etc.). | |
| Version-Release number of selected component (if applicable): | |
| Sat 6.3.0 | |
| How reproducible: | |
| 100% | |
| Steps to Reproduce: | |
| 1. Publish or promote a CV with yum repos. | |
| 2. Count the time it took | |
| Actual results: | |
| 2. is slower when https://github.com/Katello/katello/blob/master/app/lib/actions/katello/repository/clone_yum_content.rb#L22 is "sequence do" | |
| Expected results: | |
| 2. is faster if https://github.com/Katello/katello/blob/master/app/lib/actions/katello/repository/clone_yum_content.rb#L22 is "concurrence do" | |
| Additional info: | |
| (*) Not sure if the above change improves also CV promote (it should, havent tested). It doesnt improve incremental update where my proposal isnt applicable. | |
| (*) I will provide a script to test performance gain | |

History

#1 - 04/25/2018 05:55 PM - Jonathon Turel

- Assignee set to Samir Jha

- translation missing: en.field_release set to 349

#2 - 04/27/2018 08:27 PM - The Foreman Bot

- Pull request <https://github.com/Katello/katello/pull/7343> added

#3 - 06/11/2018 03:53 PM - Samir Jha

- Status changed from New to Rejected

#4 - 08/07/2018 07:18 PM - Samir Jha

- Target version changed from Katello 3.8.0 to Katello Recycle Bin