

OSTree Support - Issue #3999

Publishing incorrect branch head.

09/11/2018 11:02 PM - jortel@redhat.com

Status:	CLOSED - CURRENTRELEASE	Start date:	
Priority:	Normal	Due date:	
Assignee:	jortel@redhat.com	Estimated time:	0:00 hour
Category:			
Sprint/Milestone:			
Severity:	2. Medium	Groomed:	No
Version - OSTree:		Sprint Candidate:	No
Platform Release:	2.18.1	Tags:	Pulp 2
Target Release - OSTree:		Sprint:	Sprint 45
OS:		Quarter:	
Triaged:	Yes		

Description

Problem

The ostree plugin creates a `Branch` content (units) for each commit in each branch history. Each `Branch` has a `_created` timestamp that is used by the publisher to determine the branch HEAD for each branch. Because Branch._created is the timestamp when pulp created the record in the DB instead of the timestamp of the commit. As a result, subsequent syncs with a deeper history (depth increased), a newer Branch content (unit) can be created for an older commit. This causes the publisher, which is sorting on _created, to incorrectly determine the branch HEAD. Publishing the branch with a HEAD that is not the latest commit in the history means the history is not deep enough to satisfy the depth requirement.

Example:

depth=3 branch=A

Fetch (PULL) history for branch=A:

commit-9
commit-8
commit-7

Create content (unit):

commit-7 _created: 2018-01-01
commit-8 _created: 2018-01-02
commit-9 _created: 2018-01-03

Change the depth to 4:

Fetch (PULL) history for branch=A:

commit-9
commit-8
commit-7
commit-6

Create content (unit):

commit-6 _created: 2018-01-04

List content (sorted by _created):

commit-7 _created: 2018-01-01
commit-8 _created: 2018-01-02
commit-9 _created: 2018-01-03
commit-6 _created: 2018-01-04 <-- uh-oh

Now the publisher (sorting by `_created`) incorrectly determines commit-6 is the branch HEAD and tries to do a PULL-LOCAL on commit-6 with `depth=4`. libostree will attempt to pull commit-6's parent (commit-5) but it's not in pulp's storage (ostree) repository because the `depth=4` keep it from being fetched.

```
GLib.Error('Importing commit-5.commit: linkat: No such file or directory', 'g-io-error-quark', 1) (Katello::Errors::PulpError)
```

Proposed Solution

Discontinue storing and using

```
Branch._created
```

to determine branch commit ordering. Then, add

```
Branch.parent
```

that contains the parent commit ID. Basically, model the parent/child commit chain just as OSTree does. This provides a deterministic method of maintaining the commit chain that will support rebasing and changing traversal depth.

The importer would pull the history (as it does currently) and set the commit and parent when creating a Branch content (unit). Further, it would need to update any existing Branch content (units) that have been re-parented by either rebase or an increase in traversal depth.

The publisher would determine the branch HEAD based on walking the commit chain.

This will require a migration that would need to iterate all of the Branch content (units) and *unset*

```
Branch._created
```

and *set*

```
Branch.parent
```

using the history.

Questions

1. When a commit is (re)parented, the `Branch.parent` will need to be updated. What problems might this cause?

Related issues:

Related to Pulp - Test #4258: Publishing incorrect branch head.

CLOSED - COMPLETE

Has duplicate OSTree Support - Issue #4276: Ostree repo with custom depth set...

CLOSED - DUPLICATE

Associated revisions

Revision `dfe2d07e` - 11/26/2018 05:00 PM - jortel@redhat.com

Publish correct branch head and depth. closes #3999

History

#1 - 09/12/2018 05:39 PM - jortel@redhat.com

- Description updated

#2 - 09/12/2018 05:40 PM - jortel@redhat.com

- Description updated

#3 - 09/14/2018 04:41 PM - CodeHeeler

- Triaged changed from No to Yes

#4 - 11/14/2018 06:25 PM - jortel@redhat.com

- Sprint set to Sprint 46

#5 - 11/14/2018 06:25 PM - jortel@redhat.com

- Sprint changed from Sprint 46 to Sprint 45

#6 - 11/26/2018 05:05 PM - jortel@redhat.com

- Status changed from NEW to POST

https://github.com/pulp/pulp_ostree/pull/104

#7 - 11/26/2018 05:09 PM - jortel@redhat.com

- Status changed from POST to MODIFIED

Applied in changeset [dfe2d07e11093c38576f2abd594a1d7ec7540ec6](#).

#8 - 12/04/2018 07:47 PM - amacdona@redhat.com

- Assignee set to jortel@redhat.com

#9 - 12/11/2018 02:33 PM - bherring

- Related to Test #4258: Publishing incorrect branch head. added

#10 - 12/18/2018 10:25 PM - davidDavis

- Has duplicate Issue #4276: Ostree repo with custom depth set returns error added

#12 - 01/22/2019 09:12 PM - jortel@redhat.com

- Platform Release set to 2.18.1

#13 - 02/07/2019 11:00 PM - jortel@redhat.com

- Status changed from MODIFIED to 5

#14 - 02/22/2019 11:25 PM - jortel@redhat.com

- Status changed from 5 to CLOSED - CURRENTRELEASE

#15 - 04/15/2019 10:08 PM - bmbouter

- Tags Pulp 2 added