

Arvados - Bug #15148

keep-balance incorrectly accounts for blocks in collections with null `modified_at` field

04/24/2019 05:48 PM - Tom Morris

Status: Resolved	Start date: 04/26/2019
Priority: Normal	Due date:
Assigned To: Tom Clegg	% Done: 100%
Category:	Estimated time: 0.00 hour
Target version: 2019-05-08 Sprint	
Description In certain circumstances, when collections have a null modified_at field (which should normally never happen), keep-balance can mark blocks for deletion even though there are still references to them. A check for this will be included in 1.4 and 1.3.2: keep-balance will refuse to run at all if this situation is detected. An API server bug which causes the modified_at field to be null was introduced by #13561 and released with Arvados v1.3. The API server bug was fixed as a side effect of #14595 and this fix will also be included in 1.4 and 1.3.2. Related: Recovering lost data	
Subtasks: Task # 15156: Review 15148-lost-collection-pdh Resolved	
Related issues: Related to Arvados - Feature #13561: [API] Store, and add APIs to retrieve, p... Resolved 10/04/2018	

Associated revisions

Revision 96be8f8e - 04/24/2019 08:42 PM - Tom Clegg

Merge branch '15112-save-lost-blocks-file'

refs #15112
refs #15148

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <tclegg@veritasgenetics.com>

Revision c2ed4aab - 04/24/2019 08:52 PM - Tom Clegg

Merge branch '15112-save-lost-blocks-file'

refs #15112
refs #15148

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <tclegg@veritasgenetics.com>

Revision 48d2a213 - 04/26/2019 08:39 PM - Tom Clegg

Merge branch '15148-lost-collection-pdh'

refs #15148

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <tclegg@veritasgenetics.com>

History

#1 - 04/24/2019 05:58 PM - Tom Morris

Anyone running 1.3 should do these things ASAP:

- Wherever keep-balance is running: systemctl stop keep-balance (or whatever it takes to disable keep-balance) until a fixed version is released+installed
- On every keepstore node: add TrashCheckInterval: 87600h to /etc/arvados/keepstore/keepstore.yml, and then systemctl restart keepstore (or equivalent) to avoid deleting any trashed blocks that are still recoverable

Then

- Install a fixed version of keep-balance and arvados-api-server ($\geq 1.3.2$ or ≥ 1.4)
- Enable keep-balance
- Use the keepstore untrash API to recover any blocks that were trashed but not yet deleted (details TBD)

- Delete/revert TrashCheckInterval in keepstore configs and restart keepstore processes

Any system with containers that finished while running Arvados 1.3 will need a migration to fix the collections table for the output collections of those containers.

The following fixes have been made:

- Fix to refuse to run if any modified_at fields are null 15112-dont-trash-needed-replica @ 8b9ea19ebde9f4653d6adc145ef6fcbd36d2aace
- Database migration to repair any null `modified_at`s:
 - 15112-migration @ 243130b8c5a8558d6bd132d4a062483be93ef7bc
 - 15112-migration-1.3 @ a2c3d1ffd627974c8daa3bff300c0ad96f07d3a0
 - Update migration to handle empty database case @ 7bc8e8add
- Cherry pick [#14595](#) - \$ git cherry-pick 2aa58f31ac8fc696361214a05ab9ba75a5140b084e32f0b140ec0ec7f96c1f9eaae00950c176ff03

#2 - 04/24/2019 06:27 PM - Tom Clegg

- Description updated

#4 - 04/26/2019 03:34 PM - Tom Morris

- Release set to 23

#5 - 04/26/2019 05:09 PM - Tom Clegg

- Description updated

#6 - 04/26/2019 05:46 PM - Tom Clegg

- Install a fixed version of keep-balance and arvados-api-server ($\geq 1.3.2$ or ≥ 1.4)
- Enable keep-balance
- Use the keepstore untrash API to recover any blocks that were trashed but not yet deleted (details TBD)

Details: [Untrashing lost blocks](#)

Any system with containers that finished while running Arvados 1.3 will need a migration to fix the collections table for the output collections of those containers.

This migration runs during the upgrade to arvados $\geq 1.3.2$ or ≥ 1.4 .

#7 - 04/26/2019 08:09 PM - Tom Clegg

15148-lost-collection-pdh @ [6c5852fb18c0b6422c079c6fee66891a273ad089](#) -- <https://ci.curoverse.com/view/Developer/job/developer-run-tests/1219/>

#8 - 04/26/2019 08:37 PM - Peter Amstutz

Tom Clegg wrote:

15148-lost-collection-pdh @ [6c5852fb18c0b6422c079c6fee66891a273ad089](#) -- <https://ci.curoverse.com/view/Developer/job/developer-run-tests/1219/>

Waiting on jenkins but this LGTM.

#9 - 04/26/2019 08:45 PM - Tom Clegg

This change is in master ([48d2a213b](#), destined for 1.4) and 1.3-dev ([675237bec](#), destined for 1.3.3).

Each line of the "lost blocks" file will now be "BLOCKHASH PDH1 PDH2 ..." where PDH* are all collections that refer to BLOCKHASH. From here you can get a complete list of affected collection PDHs:

```
cut -d" " -f2- < lost-blocks.txt | tr " " "\n" | sort -u > lost-collections.txt
```

#10 - 05/08/2019 05:22 PM - Tom Clegg

- Status changed from In Progress to Resolved

#11 - 05/08/2019 06:17 PM - Tom Morris

- Description updated

#12 - 05/13/2019 08:30 PM - Tom Morris

- Release changed from 23 to 24

#13 - 05/13/2019 08:31 PM - Tom Morris

- Description updated

#14 - 09/23/2019 02:55 PM - Tom Morris

- Related to Feature #13561: [API] Store, and add APIs to retrieve, previous versions of collection objects added