

Arvados - Bug #18547

[keep-balance] Avoid redundant indexing when multiple keepstore servers use a single NFS mount

12/03/2021 03:04 PM - Tom Clegg

Status:	Resolved	Start date:	12/06/2021
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	2021-12-08 sprint		
Description			
Background: Currently keep-balance detects when a storage device like an S3 bucket is used by multiple keepstore servers, and arbitrarily chooses one of them to get the index. However, this relies on the "device ID" returned by keepstore, which is			
<ul style="list-style-type: none">• s3://endpoint/bucketname, if the volume is an S3 bucket• block device UUID, if the volume is a local filesystem• empty, if the volume is a network-mounted filesystem			
When the device ID is empty, the volumes might be different, so keep-balance indexes all of them.			
Now that each keepstore server uses the same configuration file, each configured volume has a unique UUID, and the volume UUID is returned in the list of mounts reported by keepstore (none of which were true when the "device ID" approach started), keep-balance should detect identical/duplicate mounts by comparing volume UUIDs instead of device IDs.			
(Note this will confuse keep-balance if the config uses a single volume UUID to mount a local disk like "/data" on multiple keepstore machines. But the install docs explicitly describe <i>not</i> doing that, and it is not a kind of configuration we want to support. Worst outcome is that someone with this kind of wonky config would see a lot of blocks misreported as underreplicated or missing by keep-balance until they fix their config.)			
Subtasks:			
Task # 18561: Review 18547-use-volume-uuid-not-device-id			Resolved
Related issues:			
Related to Arvados - Bug #18376: [keepstore] Avoid long-lived readdirent cook...		Resolved	11/16/2021
Blocks Arvados - Story #18518: Release Arvados 2.3.2		Resolved	12/06/2021

Associated revisions

Revision 92030788 - 12/07/2021 08:07 PM - Tom Clegg

Merge branch '18547-use-volume-uuid-not-device-id'

fixes #18547

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <tom@curii.com>

History

#1 - 12/03/2021 03:10 PM - Tom Clegg

- Related to Bug #18376: [keepstore] Avoid long-lived readdirent cookies in filesystem driver added

#2 - 12/03/2021 04:27 PM - Tom Clegg

- Status changed from New to In Progress

#3 - 12/03/2021 04:28 PM - Tom Clegg

18547-use-volume-uuid-not-device-id @ [e16866d0f398f6c61f11e2ecdf473d47100329c0](https://github.com/Arvados/arvados/commit/e16866d0f398f6c61f11e2ecdf473d47100329c0) -- developer-run-tests: #2824
[icon?job=developer-run-tests&build=2824](https://github.com/Arvados/arvados/commit/e16866d0f398f6c61f11e2ecdf473d47100329c0)

#4 - 12/06/2021 04:23 PM - Tom Clegg

Wondering whether we want a more limited version of this for 2.3.2 ("only use UUID if deviceID is empty") just in case the full version affects non-NFS setups in an unexpected way...

#5 - 12/06/2021 04:50 PM - Tom Clegg

TODO: keep-balance should error out if two volumes return the same non-empty DeviceID.

#6 - 12/06/2021 05:10 PM - Peter Amstutz

- Release set to 48

#7 - 12/06/2021 09:31 PM - Peter Amstutz

- Blocks Story #18518: Release Arvados 2.3.2 added

#8 - 12/07/2021 12:34 AM - Tom Clegg

18547-use-volume-uuid-not-device-id @ [24f140f9ed1a2180541c0c7ceb7572c5155fe27](#) -- [developer-run-tests: #2829 icon?job=developer-run-tests&build=2829](#)

- error out if two volumes return the same non-empty DeviceID

#9 - 12/07/2021 07:00 PM - Peter Amstutz

Tom Clegg wrote:

18547-use-volume-uuid-not-device-id @ [24f140f9ed1a2180541c0c7ceb7572c5155fe27](#) -- [developer-run-tests: #2829 icon?job=developer-run-tests&build=2829](#)

- error out if two volumes return the same non-empty DeviceID

This LGTM. Could you please merge into both main and 2.3-dev?

#10 - 12/07/2021 08:08 PM - Tom Clegg

- Status changed from *In Progress* to *Resolved*

Applied in changeset arvados-private:commit:arvados|920307882b3fe52a08b366a1c81e62f44ee639b9.

#11 - 12/07/2021 08:09 PM - Tom Clegg

Merged, and cherry-picked e16866d0f and 24f140f9e onto 2.3-dev as 11864d817 and 56c37ef9b respectively.