

Arvados - Feature #17749

[Keep] avoid AWS S3 request limits -- add option to use more prefixes on S3

06/02/2021 04:16 PM - Ward Vandewege

Status:	Resolved	Start date:	09/21/2021
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:	Keep	Estimated time:	0.00 hour
Target version:	2021-09-29 sprint		

Description

AWS has a hard request limit of 3,500 PUT/COPY/POST/DELETE or 5,500 GET/HEAD requests per second **per prefix** in an Amazon S3 bucket, cf. <https://aws.amazon.com/premiumsupport/knowledge-center/s3-request-limit-avoid-throttling/>.

Prefixes are defined as follows (cf. <https://aws.amazon.com/premiumsupport/knowledge-center/s3-prefix-nested-folders-difference/>):

A prefix is the complete path in front of the object name, which includes the bucket name. For example, if an object (123.txt) is stored as BucketName/Project/WordFiles/123.txt, the prefix is "BucketName/Project/WordFiles/". If the 123.txt file is saved in a bucket without a specified path, the prefix value is "BucketName/".

Keep currently does not store its blocks in subdirectories in the S3 buckets it uses. That means the prefix value for all blocks in a particular bucket is "BucketName/", and is subject to the request limits per bucket.

At some point, we may run into the request limits, particularly in a situation where one S3 bucket is shared along many keepstores, e.g. after [#16516](#) is implemented.

The fix would be to use more prefixes in each S3 bucket, perhaps adopting the same pattern keepstore uses when backed by POSIX filesystems.

There is another reason to do this: buckets with very large number of blocks become slow in certain (external) tools like aws sync. Getting a list of all those files on S3-compatible storage can be slow. Having an option to make Keep on S3 use a structure like we do on POSIX disks. Add a config option, default off.

- Config option could specify where you want the slashes. Prefix length, defaults to zero, three is recommended if you want to enable this feature on S3.
- The migration path for an existing S3 bucket with data is out of scope (migration could be handled with a script). We could do that in a future story.
- Will need to update both S3 drivers
- Same logic would apply to the trash folder in this scenario
- Will need some new tests

Subtasks:

Task # 18153: Review 17749-s3-prefixes **Resolved**

Related issues:

Related to Arvados Epics - Story #16516: Run Keepstore on local compute nodes **In Progress** 10/01/2021 11/30/2021

Associated revisions

Revision 5dbf7280 - 09/23/2021 01:54 PM - Tom Clegg

Merge branch '17749-s3-prefixes'

closes #17749

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

History

#1 - 06/02/2021 04:21 PM - Ward Vandewege

- Description updated

#2 - 06/02/2021 04:22 PM - Ward Vandewege

- Description updated

- Subject changed from [Keep] investigate AWS S3 request limits to [Keep] avoid AWS S3 request limits

#3 - 06/02/2021 04:22 PM - Ward Vandewege

- Related to Story #16516: Run Keepstore on local compute nodes added

#4 - 06/02/2021 04:23 PM - Ward Vandewege

- Description updated

#5 - 06/28/2021 03:40 PM - Ward Vandewege

- Description updated

#6 - 07/06/2021 09:08 PM - Peter Amstutz

- Target version deleted (To Be Groomed)

#7 - 07/26/2021 02:22 PM - Ward Vandewege

- Description updated

#8 - 07/28/2021 02:59 PM - Ward Vandewege

- Subject changed from [Keep] avoid AWS S3 request limits to [Keep] avoid AWS S3 request limits -- add option to use more prefixes on S3

#9 - 07/28/2021 04:50 PM - Ward Vandewege

- Story points set to 2.0

- Description updated

#10 - 07/28/2021 04:50 PM - Ward Vandewege

- Target version set to 2021-09-01 sprint

#11 - 08/10/2021 05:36 PM - Peter Amstutz

- Target version changed from 2021-09-01 sprint to 2021-09-15 sprint

#12 - 09/01/2021 03:25 PM - Peter Amstutz

- Target version changed from 2021-09-15 sprint to 2021-09-29 sprint

#13 - 09/15/2021 03:09 PM - Tom Clegg

- Assigned To set to Tom Clegg

#14 - 09/21/2021 02:21 AM - Tom Clegg

- Status changed from New to In Progress

- Category set to Keep

I'm not sure I've made the docs clear enough, particularly the part about when/why not to change PrefixLength.

17749-s3-prefixes @ [adccfe35ccc68a865a2fd2356ca2b81e0366a4b4](https://ci.arvados.org/view/Developer/job/developer-run-tests/2699/) -- <https://ci.arvados.org/view/Developer/job/developer-run-tests/2699/>

#15 - 09/21/2021 06:52 PM - Ward Vandewege

Tom Clegg wrote:

I'm not sure I've made the docs clear enough, particularly the part about when/why not to change PrefixLength.

17749-s3-prefixes @ [adccfe35ccc68a865a2fd2356ca2b81e0366a4b4](https://ci.arvados.org/view/Developer/job/developer-run-tests/2699/) -- <https://ci.arvados.org/view/Developer/job/developer-run-tests/2699/>

LGTM, thanks!

#16 - 09/23/2021 02:46 PM - Tom Clegg

- Status changed from In Progress to Resolved

Applied in changeset arvados-private:commit:arvados|5dbf72803717f58b4848b6a6490375450916e84d.