Arvados - Story #11016

Document how to choose a suitable blob signature TTL

01/31/2017 08:03 PM - Tom Clegg

Status:	Resolved	**Start date:**	10/01/2019
Priority:	Normal	**Due date:**	
Assigned To:	Tom Clegg	**% Done:**	100%
Category:		**Estimated time:**	0.00 hour
Target version:	2019-10-09 Sprint		

Description

Subtasks:

Task # 15482: Review 11016-doc-signing-ttl	**Resolved**

Related issues:

Related to Arvados - Story #15697: [doc] explain lifecycle of Keep blocks, an...	**New**

Associated revisions

Revision fd38b59a - 10/03/2019 03:39 PM - Tom Clegg

Merge branch '11016-doc-signing-ttl'

refs #11016

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

History

#1 - 08/28/2017 09:27 PM - Tom Morris

- Target version set to Arvados Future Sprints

#2 - 07/08/2019 05:30 PM - Nico César

I am re-using this ticket, since we had a lot of inquiries about this. This is an example of it:

> We are trying to delete data from the keepstore. But after we run "arv collection delete --uuid=zzzzz-4zz18-7p9s7j1qa" it trashes it, but sets the final delete date 2 weeks in the future.
>
> I ask because on our test server we have filled up the keepstore and can't delete data to continue testing. It's put a stop to the project.
>
> ```
> "delete_at":"2019-07-11T20:45:30.557593000Z",
> "trash_at":"2019-06-27T20:45:30.557593000Z",
> "is_trashed":true,
> ```
>
> Even if we change the trash times in " /etc/arvados/keepstore/keepstore.yml" and reinstall keepstore + restart the service.
>
> \# How often to check for (and delete) trashed blocks whose
>
> ```
> # TrashLifetime has expired.
> TrashCheckInterval: 1h0m0s
> ```
>
> ```
> # Time duration after a block is trashed during which it can be
> # recovered using an /untrash request.
> TrashLifetime: 1h0m0s
> ```

this was my answer:

> Just to make sure we are all in the same page in terms of terminology, there is a good explanation here: https://doc.arvados.org/user/tutorials/tutorial-keep-collection-lifecycle.html
>
> And the method you used is "delete" in collections, from https://doc.arvados.org/v1.4/api/methods/collections.html
> ----------
> delete
>
> Put a Collection in the trash. This sets the trash_at field to now and delete_at field to now + token TTL. A trashed collection is invisible to most API calls unless the include_trash parameter is true.

As you can see the "token TTL" expressed there is set by default in 2 weeks, this is the Collections->BlobSigningTTL and Collections->DefaultThrashLifetime parameters in the configuration. Here is description from in https://doc.arvados.org/v1.4/admin/config.html

```
--------

# Lifetime (in seconds) of blob permission signatures generated by
# the API server. This determines how long a client can take (after
# retrieving a collection record) to retrieve the collection data
# from Keep. If the client needs more time than that (assuming the
# collection still has the same content and the relevant user/token
# still has permission) the client can retrieve the collection again
# to get fresh signatures.
#
# This must be exactly equal to the -blob-signature-ttl flag used by
# keepstore servers.  Otherwise, reading data blocks and saving
# collections will fail with HTTP 403 permission errors.
#
# Modifying blob_signature_ttl invalidates existing signatures; see
# blob_signing_key note above.
#
# The default is 2 weeks.
BlobSigningTTL: 336h

# Default lifetime for ephemeral collections: 2 weeks. This must not
# be less than blob_signature_ttl.
DefaultTrashLifetime: 336h

------
```

This assumes that you have the central configuration in /etc/arvados/config.yml, keep-balance.service up and running.

As you can see we have 3 different places with pieces of the information. And usually **our test server we have filled up the keepstore** is the reason they need a quick "delete all this" process without having to wait 2 weeks.

#3 - 07/08/2019 06:37 PM - Tom Morris

One of the main things that's missing from https://doc.arvados.org/user/tutorials/tutorial-keep-collection-lifecycle.html (which is really conceptual documentation, not a tutorial)
is "when do I get my disk space back?" and the associated Keep store pieces of the data lifecycle.

#4 - 07/08/2019 06:39 PM - Tom Morris

- Target version changed from Arvados Future Sprints to 2019-07-31 Sprint

#5 - 07/17/2019 03:06 PM - Tom Morris

- Assigned To set to Tom Clegg

#6 - 07/31/2019 03:00 PM - Tom Clegg

- Target version changed from 2019-07-31 Sprint to 2019-08-14 Sprint

#7 - 08/14/2019 01:30 PM - Tom Clegg

- Status changed from New to In Progress

BlobSigningTTL determines the minimum lifetime of transient data, i.e., blocks that are not referenced by collections. Unreferenced blocks exist for two reasons:

- A data block must be written to a disk/cloud backend device before a collection can be created/updated with a reference to it.
- Deleting or updating a collection can have the effect of removing the last remaining reference to a data block.

If BlobSigningTTL is too short, long-running processes/containers will fail when they take too long (a) between writing blocks and writing collections that reference them, or (b) between reading collections and reading the referenced blocks.

If BlobSigningTTL is too long, data will still be stored long after the referring collections are deleted, and you will needlessly fill up disks or waste money on cloud storage.

#8 - 08/14/2019 03:15 PM - Tom Clegg

- Target version changed from 2019-08-14 Sprint to 2019-08-28 Sprint

#9 - 08/28/2019 02:40 PM - Tom Morris

- Target version changed from 2019-08-28 Sprint to 2019-09-11 Sprint

#10 - 09/11/2019 02:46 PM - Tom Clegg

- Target version changed from 2019-09-11 Sprint to 2019-09-25 Sprint

#11 - 09/25/2019 03:15 PM - Tom Clegg

- Target version changed from 2019-09-25 Sprint to 2019-10-09 Sprint

#12 - 10/01/2019 02:10 PM - Tom Clegg

11016-doc-signing-ttl @ f902c301c8c40707bb25f8dae01537f8a50be286 -- https://ci.curoverse.com/view/Developer/job/developer-run-tests/1571/

#13 - 10/01/2019 05:01 PM - Tom Morris

Suggestions/comments:

- add "(ie Time-To-Live)" after "BlobSigningTTL determines the minimum lifetime" to help explain TTL. I'm not sure it's a familiar term outside of the networking space.
- it seems like there are pieces of the lifecycle for blocks and signed block references missing from https://doc.arvados.org/user/tutorials/tutorial-keep-collection-lifecycle.html unless we have them covered somewhere else (ie the two cases mentioned in the config file description)

#14 - 10/08/2019 07:31 PM - Tom Clegg

- Related to Story #15697: [doc] explain lifecycle of Keep blocks, and how it affects storage backend usage/cost added

#15 - 10/08/2019 07:32 PM - Tom Clegg

- Status changed from In Progress to Resolved

#16 - 01/22/2020 02:45 PM - Peter Amstutz

- Release set to 22

- Status changed from Resolved to Rejected

#17 - 01/22/2020 02:45 PM - Peter Amstutz

- Status changed from Rejected to Resolved