

Arvados - Story #8181

[Keepstore] Support a portable data challenge API

01/11/2016 04:27 PM - Brett Smith

Status:	New	Start date:	01/11/2016
Priority:	Normal	Due date:	
Assigned To:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	To Be Groomed		

Description

Keepstore gives the client a salt. The client returns some hash including that salt plus the original data to prove it has the original data. If the client does so successfully, Keepstore discloses whether or not it has the block, and returns a newly signed block locator. Upload clients can use this to avoid uploading blocks that already exist on the destination.

Alternatively: add a 'trust keep clients' flag in config.yml, which site admins could enable to allow client-side hashing, and let the keepstores trust the client-provided hashes.

The ideal solution here is probably a slider that is a % of the data that the client needs to prove it has, before keepstore will trust its hashes. It could be set to zero in a situation where there is a lot of trust, and probably to some reasonable % by default (25? 50?).

What we gain from this story would be avoiding the need to re-upload data if it already exists in Keep.

History

#1 - 01/12/2016 06:52 PM - Brett Smith

- Target version set to Arvados Future Sprints

#2 - 05/12/2020 03:06 PM - Ward Vandewege

- Target version changed from Arvados Future Sprints to To Be Groomed

#3 - 05/12/2020 03:09 PM - Ward Vandewege

- Description updated

#4 - 05/12/2020 03:29 PM - Ward Vandewege

- Description updated