

Arvados - Bug #16426

[arv-put] should calculate (worst case) manifest size up front before upload, and abort when it is too large

05/12/2020 02:52 PM - Ward Vandewege

Status:	New	Start date:	
Priority:	Normal	Due date:	
Assigned To:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			

Description

The API server has a configurable limit to the manifest size it will accept. When arv-put is asked to upload a set of files and directories, it just starts uploading until it has all the blocks written, and then calculates the manifest and tries to save it.

When that fails because the manifest is too large, the user has wasted a lot of time and network bandwidth. They will need to make a smaller version of the collection, and then reupload all the blocks. The arv-put block cache will not help them because the collection to be uploaded will be different.

Add a feature to arv-put to calculate the (worst case) manifest size up front, and refuse to upload when it is too large. This should happen before arv-put reads all the data that it needs to upload, it should error out as early as possible.

History

#1 - 05/12/2020 02:59 PM - Ward Vandewege

- Description updated

#2 - 05/12/2020 03:00 PM - Ward Vandewege

- Description updated

#3 - 07/06/2021 09:08 PM - Peter Amstutz

- Target version deleted (To Be Groomed)