

Arvados - Bug #7573

Keepstore: very uneven distribution of blob between 2 Keepstore servers

10/15/2015 04:24 PM - Peter Grandi

Status:	Resolved
Priority:	Normal
Assigned To:	
Category:	SDKs
Target version:	
Description	
<p>Testing Keepstore. Had keep0 with a 100GB volume. Created 2 new filesystem, 03 and 04, and copied the contents of the existing one into 03 and delete the existing one, and restarted the daemon. Created keep1 with filesystems 01 and 02. All 4 filesystems are 1TiB. Registered keep1 with the API server.</p> <p>When uploading with arv-put a small number of files each of a few GB plus 1 file of 60GB the 64MiB blobs get distributed as follows:</p> <p>keep0:</p> <pre>\$ find /var/lib/keepstore/gcaml-keep-04 -type f wc -l 2264 \$ find /var/lib/keepstore/gcaml-keep-03 -type f wc -l 3268</pre> <p>keep1:</p> <pre>\$ find /var/lib/keepstore/gcaml-keep-02 -type f wc -l 3 \$ find /var/lib/keepstore/gcaml-keep-01 -type f wc -l 2</pre> <p>That seems very strange to me. I can understand the filesystem 03 has more blobs then 04 because it has the "old" blobs.</p> <p>Looking at a couple of the 5 blobs on keep1 in 01 and 02 they seem to belong to files stored almost entirely on keep0. What seems strange to me is both that:</p> <ul style="list-style-type: none">• Files are not evenly distributed between keep0 and keep1.• They are evenly distributed between 03 and 04 on keep0 but some stray blobs end up (apparently evenly distributed) on keep1.	
Related issues:	
Is duplicate of Arvados - Bug #6358: [SDKs] Python Keep client uses wrong pro...	Resolved 10/16/2015

History

#1 - 10/19/2015 07:46 PM - Tom Clegg

- Category set to SDKs

- Status changed from New to Feedback

[#6358](#) contained fixes for two different Python SDK bugs affecting block distribution. One of them is almost certainly the biggest contributor to this problem.

You should get even distribution when the writer (arv-put) is from arvados-python-client-0.1.20151019192928 or newer.

Uploads through keepproxy (including browser uploads) would not have been affected by either of those bugs, so this explanation assumes your arv-put process had direct access to keepstore servers (e.g., it was running on a shell node).

#2 - 10/21/2015 10:20 AM - Peter Grandi

our arv-put process had direct access to keepstore servers (e.g., it was running on a shell node).

Indeed; and with the update mentioned above new uploads are now more evenly distributed between the keep0 and keep1 servers, and are still evenly distributed between the two filetrees 01 and 02 on the keep1 server.

#3 - 10/21/2015 01:33 PM - Brett Smith

- Status changed from Feedback to Resolved

Peter Grandi wrote:

Indeed; and with the update mentioned above new uploads are now more evenly distributed between the keep0 and keep1 servers, and are still evenly distributed between the two filetrees 01 and 02 on the keep1 server.

Glad to hear it. We have other feedback that the fix improved distribution as well, so I'm going to mark this as resolved. Thanks for the report, and please don't hesitate to reopen this or file a new report if you see any other funny behavior.