

Arvados - Story #9653

[keepstore] Limit entire process to configured amount or % of system memory (replace -max-buffers=X)

07/20/2016 08:57 PM - Brett Smith

Status: New	Start date: 07/20/2016
Priority: Normal	Due date:
Assigned To:	% Done: 0%
Category: Keep	Estimated time: 0.00 hour
Target version:	
Description	
It's hard to predict the right setting for -max-buffers.	
<ul style="list-style-type: none">• The sysadmin must coordinate configuration files with host characteristics: e.g., when upgrading to a bigger node, the sysadmin must also remember to change max-buffers at the same time.• The sysadmin must account for (and/or override) Go's garbage collection threshold, likely 100% (default) or 10% (our recommendation).• The sysadmin must predict how much memory keepstore will use for purposes other than {buffers + 10% garbage collection overhead}, and reduce max-buffers accordingly.• The sysadmin must predict how much memory will be used by other processes running on the same node (e.g., cloud/monitoring/orchestration agents).	
In addition to a static limit on the number of buffers, we could configure a maximum keepstore memory size, checking the actual process memory usage (not just bytes-in-buffers) before allocating a new buffer. This isn't perfect (it wouldn't be a hard limit) but it would be strictly safer and easier than the current behavior.	
Memory use could be limited by providing one or more of these configs:	
<pre># Maximum memory use, as a number of MiB MaxMemoryMiB: 6000 # Maximum memory use, as a percentage of total system RAM MaxMemoryPercent: 80</pre>	
Suggested default:	
<pre>MaxMemoryPercent: 80</pre>	
To address the last difficulty (predicting memory use by other processes) we could also offer a MinMemoryFree or MinMemoryAvailable config, which avoids allocating more buffers when Linux reports a low MemFree or MemAvailable.	
More than one limit can be configured at once. Keepstore should allocate new buffers only when none of the configured limits will be exceeded.	

History

#1 - 11/28/2017 07:41 PM - Tom Clegg

- Subject changed from [Keep] Improve default for keepstore -max-buffers to [keepstore] Limit entire process to configured amount or % of system memory (replace -max-buffers=X)

- Description updated

#2 - 11/28/2017 07:47 PM - Tom Clegg

- Description updated

#3 - 11/28/2017 07:53 PM - Tom Morris

- Target version set to To Be Groomed

#4 - 11/28/2017 07:59 PM - Tom Morris

I'd suggest that we bias this to work reliably at the expense of not using every last nibble of RAM.

Perhaps a default scheme along the lines of this: https://www.datacore.com/SSV-Webhelp/Changing_Cache_Size.htm with values tweaked appropriately.

- o RAM less than or equal to 8 GB; the cache will use 60%. (8 GB is the minimum requirement)
- o RAM less than or equal to 128 GB; the cache will use 65%.
- o RAM greater than 128 GB; the cache will use 75%.

#6 - 07/06/2021 09:10 PM - Peter Amstutz

- *Target version deleted (To Be Groomed)*