

## Arvados - Bug #14488

### [keepstore] Do not try to trash blocks that were already trashed in previous sweep

11/14/2018 07:55 PM - Tom Clegg

<b>Status:</b>	New	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assigned To:</b>		<b>% Done:</b>	0%
<b>Category:</b>	Keep	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	To Be Groomed		
<b>Description</b>			
<p>There is a time interval -- after keep-balance retrieves keepstore's block index, and before keep-balance sends keepstore a new trash list -- during which keepstore is still processing the last trash list it received from a previous keep-balance sweep.</p> <p>Blocks that get trashed during this interval will typically appear on the new trash list, too. When keepstore runs the new trash list, it logs errors: the blocks can't be trashed because they don't exist.</p> <p>This produces logs that are hard to distinguish from real errors, and it wastes time and cloud API calls.</p> <p>Proposed solution:</p> <ul style="list-style-type: none"><li>• After a block is successfully trashed, add the block ID to a map of recent trashes.</li><li>• When a new trash list is accepted, remove any block ID that appears in the map of recent trashes, and clear the map of recent trashes.</li></ul>			