

Arvados - Feature #16585

[keep-exercise] improvements

07/06/2020 01:22 PM - Ward Vandewege

Status:	Resolved	Start date:	07/17/2020
Priority:	Normal	Due date:	
Assigned To:	Ward Vandewege	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	2020-08-12 Sprint		
Description			
Add --repeat option to have keep-exercise repeat experiments N times.			
Make sure the read threads never starve; if there are no new blocks to read, re-read the previous block.			
Add warmup phase so that statistics only start when there is something to do for all read threads.			
Subtasks:			
Task # 16608: review 16585-keep-exercise-improvements			Resolved

Associated revisions

Revision 6e5b24e8 - 07/22/2020 08:09 PM - Ward Vandewege

Merge branch 'master' into 16585-keep-exercise-improvements

refs #16585

Arvados-DCO-1.1-Signed-off-by: Ward Vandewege <ward@curii.com>

Revision 715b5afb - 07/22/2020 08:09 PM - Ward Vandewege

Merge branch '16585-keep-exercise-improvements'

closes #16585

Arvados-DCO-1.1-Signed-off-by: Ward Vandewege <ward@curii.com>

History

#1 - 07/06/2020 01:22 PM - Ward Vandewege

- Status changed from New to In Progress

#2 - 07/06/2020 08:24 PM - Ward Vandewege

- Description updated

Ready for review at [90b3f51e0462234c54322f1381a8e7bc230938a4](#) on branch 16585-keep-exercise-improvements

#3 - 07/07/2020 04:08 PM - Ward Vandewege

Simplified version ready for review at [199d58a4df0dba1ee71c6816bc3a9d9d439cfd7e](#) on branch 16585-keep-exercise-improvements

#4 - 07/15/2020 03:51 PM - Ward Vandewege

- Target version changed from 2020-07-15 to 2020-08-12 Sprint

#5 - 07/16/2020 04:13 PM - Ward Vandewege

Added -useIndex flag to rule out caching issues. Ready for review at [34cb31d5191ce17c37ddd9d344f4d77125af64a8](#) on branch 16585-keep-exercise-improvements

#6 - 07/20/2020 08:41 PM - Tom Clegg

recommend changing -useIndex flag to -use-index to match other flags

fmt.Print("\r") after ^C should probably print to stderr instead of stdout

typo UseIdx in comment

seems weird to see stderr as the name of a logger -- how about log or lgr?

getIndexLocators() and doIndexReads() - instead of putting the entire func in an if ctx.Err()==nil, just put if ctx.Err() != nil { return } at the top and un-indent the rest

Similarly you could do for i := 0; i < *Repeat && ctx.Err() == nil; i++ { in main()

At the end of main, maybe it's better *not* to skip the summary when ctx.Err()!=nil, i.e., exiting on a signal?

Having two copies of doIndexReads/doReads doesn't seem ideal. Maybe something like this could handle both cases, and the caller would just provide a nil indexLocatorChan if !*UseIndex?

```
for ctx.Err() == nil {
    var locator string
    if indexLocatorChan != nil {
        select {
            case locator = <- indexLocatorChan:
            case <-ctx.Done():
                return
        }
    } else {
        locator = nextLocator.Load().(string)
    }
    // ...
}
```

Is it worth deduplicating the keepstore index results, to improve cache-avoidance a bit?

#7 - 07/21/2020 09:15 PM - Ward Vandewege

Tom Clegg wrote:

...

I've implemented all these changes, ready for another look at [866decb73a9cf07b28e26c5028a1d42a5ef243a7](https://github.com/Arvados/arvados/commit/866decb73a9cf07b28e26c5028a1d42a5ef243a7) on branch 16585-keep-exercise-improvements.

#8 - 07/22/2020 08:08 PM - Tom Clegg

LGTM, thanks

#9 - 07/22/2020 08:22 PM - Ward Vandewege

- % Done changed from 0 to 100

- Status changed from In Progress to Resolved

Applied in changeset [arvados|715b5afbe744a85893eca159109658bb8cddc97f](https://github.com/Arvados/arvados/commit/arvados|715b5afbe744a85893eca159109658bb8cddc97f).

#10 - 08/04/2020 08:48 PM - Ward Vandewege

- Release set to 25