

Arvados - Bug #10084

keep-balance reporting large numbers of replicas lost

09/16/2016 10:22 AM - Joshua Randall

Status:	New	Start date:	09/16/2016
Priority:	Normal	Due date:	
Assigned To:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Arvados Future Sprints		

Description

The number of replicas lost as reported by keep-balance runs is inconsistent. It seems that we have somehow actually lost 12 replicas, as I haven't seen a reported number lower than that in a very long time. However, if we look at the relevant keep-balance output from successive runs, the number of replicas reported lost bounces around in what seems like an impossible way, if the replicas are actually being lost.

```
run lost_replicas lost_blocks
```

```
1 12 6
2 35858 17929
3 12 6
4 8350 4175
5 17336 8868
6 27850 13925
7 48752 24376
8 43244 21622
```

I am concerned both with the fact that over the past 5 runs, the number of lost replicas has increased substantially (at peak on the 7th run, keep-balance reports 900GiB of lost data!) and also that somehow the number of replicas lost on consecutive runs can decrease (if they were "lost" then where did they reappear from?).

Hopefully this is just a problem with the reporting and doesn't actually represent a loss of nearly a TB of data from our keep system!

I'm going to suspend running keep-balance for now until this is better understood.

Snippets of the block/replica reports from the last eight keep-balance runs are below.

```
2016/09/10 21:17:26 ===
2016/09/10 21:17:26 12 replicas (6 blocks, 402731660 bytes) lost (0=have<want)
2016/09/10 21:17:26 0 replicas (0 blocks, 0 bytes) underreplicated (0<have<want)
2016/09/10 21:17:26 12486046 replicas (6241453 blocks, 593523262818830 bytes) just right (have=want)
2016/09/10 21:17:26 302579 replicas (297058 blocks, 14565325858873 bytes) overreplicated (have>want>0)
2016/09/10 21:17:26 698116 replicas (233640 blocks, 34465722439557 bytes) unreferenced (have>want=0, new)
2016/09/10 21:17:26 0 replicas (0 blocks, 0 bytes) garbage (have>want=0, old)
2016/09/10 21:17:26 ===
2016/09/10 21:17:26 13080174 replicas (6538517 blocks, 622142436959520 bytes) total commitment (excluding unreferenced)
2016/09/10 21:17:26 14080857 replicas (6772151 blocks, 671173082526290 bytes) total usage
2016/09/10 21:17:26 ===

2016/09/11 07:57:18 ===
2016/09/11 07:57:18 35858 replicas (17929 blocks, 2384199550518 bytes) lost (0=have<want)
2016/09/11 07:57:18 0 replicas (0 blocks, 0 bytes) underreplicated (0<have<want)
2016/09/11 07:57:18 12477260 replicas (6237060 blocks, 592942085158030 bytes) just right (have=want)
2016/09/11 07:57:18 306972 replicas (301451 blocks, 14855914689273 bytes) overreplicated (have>want>0)
2016/09/11 07:57:18 788737 replicas (263911 blocks, 40429324676316 bytes) unreferenced (have>want=0, new)
2016/09/11 07:57:18 0 replicas (0 blocks, 0 bytes) garbage (have>want=0, old)
2016/09/11 07:57:18 ===
```

2016/09/11 07:57:18 13116020 replicas (6556440 blocks, 624526233778378 bytes) total commitment (excluding unreferenced)
2016/09/11 07:57:18 14175871 replicas (6802422 blocks, 677427273593449 bytes) total usage
2016/09/11 07:57:18 ===

2016/09/12 23:24:10 ===
2016/09/12 23:24:10 12 replicas (6 blocks, 402731660 bytes) lost (0=have<want)
2016/09/12 23:24:10 0 replicas (0 blocks, 0 bytes) underreplicated (0<have<want)
2016/09/12 23:24:10 12337486 replicas (6167173 blocks, 583587380748384 bytes) just right (have=want)
2016/09/12 23:24:10 414865 replicas (408364 blocks, 22071603958238 bytes) overreplicated (have>want>0)
2016/09/12 23:24:10 1125047 replicas (379880 blocks, 62479753709589 bytes) unreferenced (have>want=0, new)
2016/09/12 23:24:10 0 replicas (0 blocks, 0 bytes) garbage (have>want=0, old)
2016/09/12 23:24:10 ===
2016/09/12 23:24:10 13154226 replicas (6575543 blocks, 627087699904130 bytes) total commitment (excluding unreferenced)
2016/09/12 23:24:10 14694126 replicas (6955417 blocks, 711638654840297 bytes) total usage
2016/09/12 23:24:10 ===

2016/09/13 06:35:02 ===
2016/09/13 06:35:02 8350 replicas (4175 blocks, 557200440406 bytes) lost (0=have<want)
2016/09/13 06:35:02 0 replicas (0 blocks, 0 bytes) underreplicated (0<have<want)
2016/09/13 06:35:02 12349946 replicas (6173403 blocks, 584419791304374 bytes) just right (have=want)
2016/09/13 06:35:02 493147 replicas (488392 blocks, 27285268340857 bytes) overreplicated (have>want>0)
2016/09/13 06:35:02 869789 replicas (293676 blocks, 45472970848285 bytes) unreferenced (have>want=0, new)
2016/09/13 06:35:02 0 replicas (0 blocks, 0 bytes) garbage (have>want=0, old)
2016/09/13 06:35:02 ===
2016/09/13 06:35:02 13335080 replicas (6665970 blocks, 639138581087192 bytes) total commitment (excluding unreferenced)
2016/09/13 06:35:02 14689666 replicas (6955471 blocks, 711339619835928 bytes) total usage
2016/09/13 06:35:02 ===

2016/09/14 03:46:39 ===
2016/09/14 03:46:39 17736 replicas (8868 blocks, 66121291282 bytes) lost (0=have<want)
2016/09/14 03:46:39 21 replicas (21 blocks, 80233266 bytes) underreplicated (0<have<want)
2016/09/14 03:46:39 12289800 replicas (6143330 blocks, 579689770164706 bytes) just right (have=want)
2016/09/14 03:46:39 594537 replicas (589750 blocks, 34055701327228 bytes) overreplicated (have>want>0)
2016/09/14 03:46:39 833745 replicas (283444 blocks, 42245362929706 bytes) unreferenced (have>want=0, new)
2016/09/14 03:46:39 0 replicas (0 blocks, 0 bytes) garbage (have>want=0, old)
2016/09/14 03:46:39 ===
2016/09/14 03:46:39 13487078 replicas (6741969 blocks, 647454343037666 bytes) total commitment (excluding unreferenced)
2016/09/14 03:46:39 14897603 replicas (7016545 blocks, 723689205770052 bytes) total usage
2016/09/14 03:46:39 ===

2016/09/14 12:34:20 ===
2016/09/14 12:34:20 27850 replicas (13925 blocks, 310138167078 bytes) lost (0=have<want)
2016/09/14 12:34:20 7 replicas (7 blocks, 24480509 bytes) underreplicated (0<have<want)
2016/09/14 12:34:20 12316442 replicas (6156651 blocks, 580247088000684 bytes) just right (have=want)
2016/09/14 12:34:20 663770 replicas (658955 blocks, 38667388242734 bytes) overreplicated (have>want>0)
2016/09/14 12:34:20 696088 replicas (239915 blocks, 31755141460443 bytes) unreferenced (have>want=0, new)
2016/09/14 12:34:20 0 replicas (0 blocks, 0 bytes) garbage (have>want=0, old)
2016/09/14 12:34:20 ===
2016/09/14 12:34:20 13662216 replicas (6829538 blocks, 657475181978554 bytes) total commitment (excluding unreferenced)
2016/09/14 12:34:20 14994217 replicas (7055528 blocks, 727587549034144 bytes) total usage

```

2016/09/14 12:34:20 ===
2016/09/14 22:23:53 ===
2016/09/14 22:23:53 48752 replicas (24376 blocks, 966035657236 bytes) lost (0=have<want)
2016/09/14 22:23:53 4 replicas (4 blocks, 9041825 bytes) underreplicated (0<have<want)
2016/09/14 22:23:53 12350286 replicas (6173573 blocks, 580706189583764 bytes) just right (have=want)
2016/09/14 22:23:53 794131 replicas (789247 blocks, 47354217458606 bytes) overreplicated (have>want>0)
2016/09/14 22:23:53 382703 replicas (138850 blocks, 8882917972357 bytes) unreferenced (have>want=0, new)
2016/09/14 22:23:53 0 replicas (0 blocks, 0 bytes) garbage (have>want=0, old)
2016/09/14 22:23:53 ===
2016/09/14 22:23:53 13977540 replicas (6987200 blocks, 675954605614158 bytes) total commitment (excluding unreferenced)
2016/09/14 22:23:53 15105618 replicas (7101674 blocks, 731225696346060 bytes) total usage
2016/09/14 22:23:53 ===

2016/09/15 22:41:38 ===
2016/09/15 22:41:38 43244 replicas (21622 blocks, 188294330854 bytes) lost (0=have<want)
2016/09/15 22:41:38 6 replicas (6 blocks, 24225695 bytes) underreplicated (0<have<want)
2016/09/15 22:41:38 12478760 replicas (6237810 blocks, 580976748480534 bytes) just right (have=want)
2016/09/15 22:41:38 851503 replicas (846515 blocks, 51179554094430 bytes) overreplicated (have>want>0)
2016/09/15 22:41:38 325370 replicas (140686 blocks, 173433770682 bytes) unreferenced (have>want=0, new)
2016/09/15 22:41:38 58466 replicas (19897 blocks, 19734728 bytes) garbage (have>want=0, old)
2016/09/15 22:41:38 ===
2016/09/15 22:41:38 14215046 replicas (7105953 blocks, 683084182066762 bytes) total commitment (excluding unreferenced)
2016/09/15 22:41:38 15407135 replicas (7244914 blocks, 734248871110053 bytes) total usage
2016/09/15 22:41:38 ===

```

History

#1 - 09/16/2016 10:32 AM - Joshua Randall

I should say, I do have a guess at what could be causing this - crunch jobs have been running during many of these runs, and it takes a very long time (~10 hours) for keep-balance to retrieve the full set of collections. Since keep indices are downloaded first and take only a minute or two, and the collection data not completed until 10 hours later, a TB worth of new collections could easily have appeared during that time period.

If that is the problem, [#9998](#) will probably help, but it might also be a good idea to record the time at which the keep indices are retrieved and to report metrics on collections updated more recently than that time differently (I would not suggest any different behaviour in terms of balancing the replicas, just that the reporting be stratified by collections modified before the oldest keep index and those modified after the oldest keep index).

In other words, you could have an extra set of counters to aggregate stats for collections modified more recently than the earliest keep index retrieval time, and use that to modify the reports (to say something like "possibly lost" or "recent replicas" or "references in collections newer than indices" rather than simply "lost").

#2 - 08/29/2017 01:41 PM - Tom Morris

- Target version set to Arvados Future Sprints