I run PerformanceEvaluation tool on our cluster which has 50+ nodes. It is running hdfs and hbase. To simplify the test, I only start one RS so that all the regions are on one RS and I can test the write throughput per RS. (datanodes are running on all the 50+ nodes)

Original case: it is configured with 6 threads that is shared between read and write traffic. I tested the cases with and without read traffic impact:

- Write only traffic. We can see the write throughput is very steady up to 35k/sec.
- Generating heavy read request at the same time using 50 read client threads. We can see the write throughput can only go up to 14k/sec, even when I increase the write client threads to 100.

After fixed case: I still use total of 6 threads, but assign 4 to writer pool and 2 to read pool.

- Write only traffic. Only 4 threads are used for the write requests.
- Large read traffic generated the same way as for the original case above. We can see the write throughput is not impacted, and can still keep at the 32k/sec.

# of write client thread	**After fix**		**Original**	
	Write only traffic	With large read traffic	Write only traffic	With large read traffic
1	5.88235	5.6338	3.93082	0.512623
5	28.9778	26.1901	32.9145	2.90894
10	34.2575	31.3551	30.8559	5.52995
15	31.3272	32.1894	33.3006	7.9688
20	34.7778	32.733	33.4492	11.2739
30	33.9563	33.0369	34.212	13.7353
40	35.6357	31.9733	32.4189	14.6948
50	35.7793	31.3109	36.1704	14.0027
100	37.2343	32.5468	37.0535	14.4268