transfersize			
cached	write	read	Data set Size/ Total memory
new set	MiB/s	MiB/s	
1	891	1083	2
2	885	1084	2
4	895	1152	2
8	889	1113	2
32	881	1092	2
512	865	933	2
1024	839	1096	2

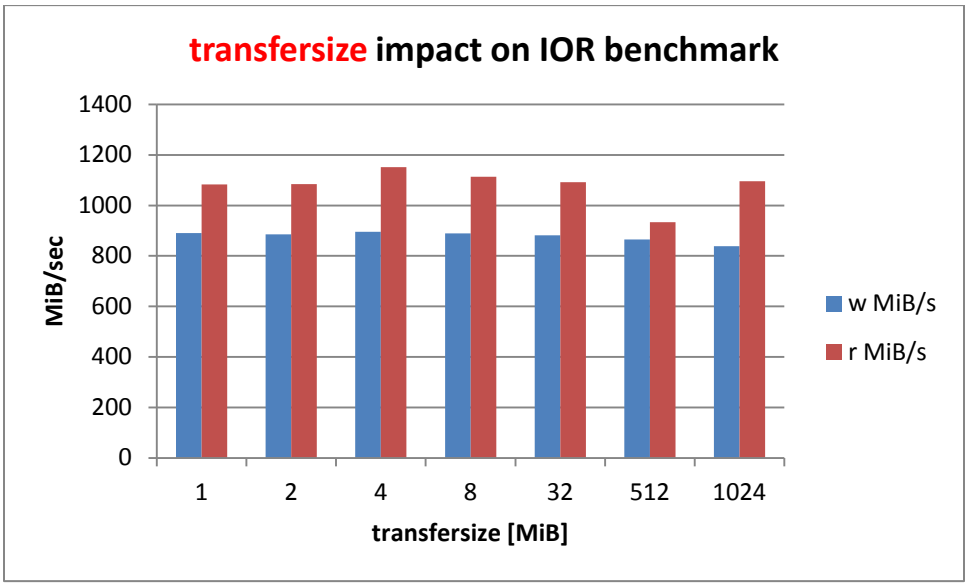

Notes: We evaluate the impact of the transfersize variable on the performance using new data sets created for each test case. The data set size was 2x total memory of the system to prevent any caching effects when running cached IOs. We made separate runs for write and read performance using new data sets, e.g. different data sets after dropping all caches, to guarantee no cached data. We can observe that both write and read peak at 4M IO. Conclusion is that transfersize performance depends on the PFS blocksize and max at the optimal size.

transfersize			
`-b 4g; -t 4m;	writye std	read std	set/mem
cached new	%BW	%BW	
3	5.2	1.1	2
6	6.1	1.3	2
10	7.3	1.5	2

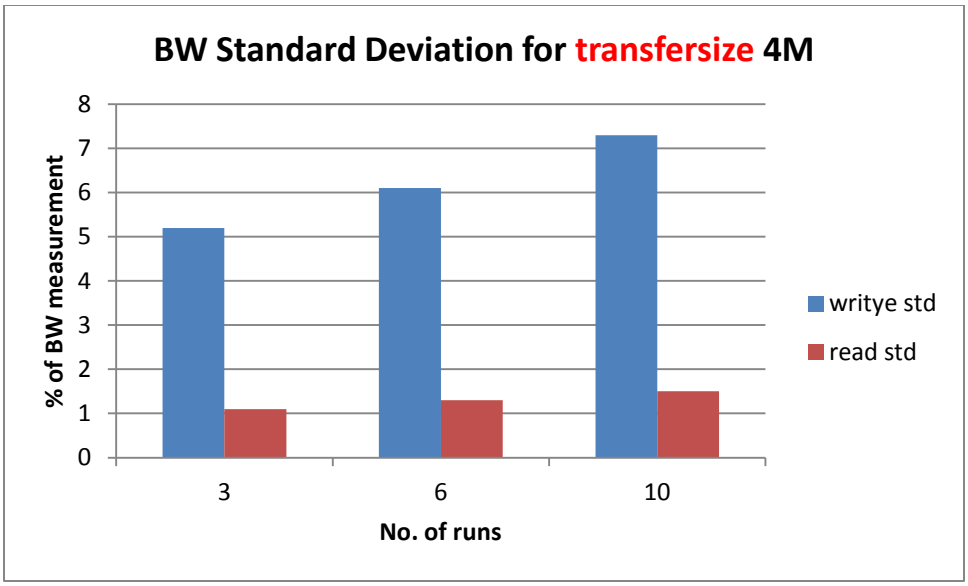

Notes: We repeated the test with optimal BW several times for both write and read BW and measured the variability of the Standard Deviation of the measurement for different numbers of repetitions using the parameter: “-i N”. We used same ratio of data set/memory of 2x and we used new data sets (uncached) as well as we flushed all the caches before each test by using the drop_caches variable of the Linux kernel. We noticed that the variability of the write BW measurement is 5-7% when using cached tests with data sets 2x than aggregate memory. The read BW performance was more stable with std under 2%. This proves that the read measurements are stable and depend little on the number of iterations.