


Analysis of the TAT of a single 4k IO



Time = avg latency of one IO (queue-depth=1) with 5x ACK

µsec	rbd		Network		Intel 910		ACK	Ceph code	
	fio write		qperf lat		fio_randwrite		msg		
	4k	8k	4k	8k	4k	8k	128	4k	8k
1 GbE	2565	2709	182	227	54	64	26	2017	2061
10 GbE	2555	2584	109	122	54	64	21	2178	2171
40 GbE	2191	2142	19	22	54	64	15	2024	1959
40 Gb IPoIB	2392	2357	29	24	54	64	18	2190	2155
56 Gb IPoIB	1848	1821	19	37	54	64	14	1686	1613

- ❑ approximately 1600 µs on a single 4k/8k IO is spend in the Ceph code
- The Ceph code has a lot of room for improvement