

Storage Ops

Thumper tuning under Solaris

Tried various zfs+pool configurations to find best IO throughput with multiple read/write processes

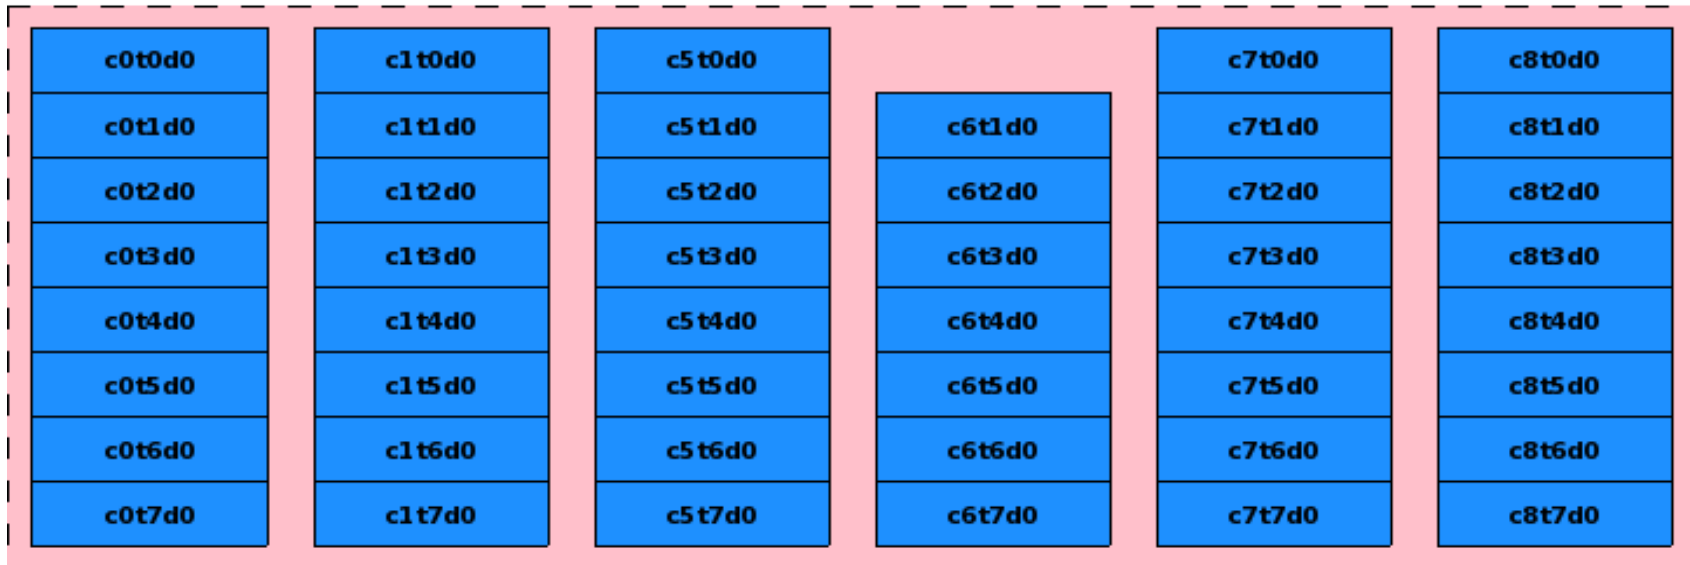
First: Reinstall with Solaris10 u5

Sun x4500 Thum per
6 Disk Controllers
8 disks per controller

c0t0d0	c1t0d0	c5t0d0	c6t0d0 (OS)	c7t0d0	c8t0d0
c0t1d0	c1t1d0	c5t1d0	c6t1d0	c7t1d0	c8t1d0
c0t2d0	c1t2d0	c5t2d0	c6t2d0	c7t2d0	c8t2d0
c0t3d0	c1t3d0	c5t3d0	c6t3d0	c7t3d0	c8t3d0
c0t4d0	c1t4d0	c5t4d0	c6t4d0	c7t4d0	c8t4d0
c0t5d0	c1t5d0	c5t5d0	c6t5d0	c7t5d0	c8t5d0
c0t6d0	c1t6d0	c5t6d0	c6t6d0	c7t6d0	c8t6d0
c0t7d0	c1t7d0	c5t7d0	c6t7d0	c7t7d0	c8t7d0

Single zfs, single pool


c6t0d0 (OS)



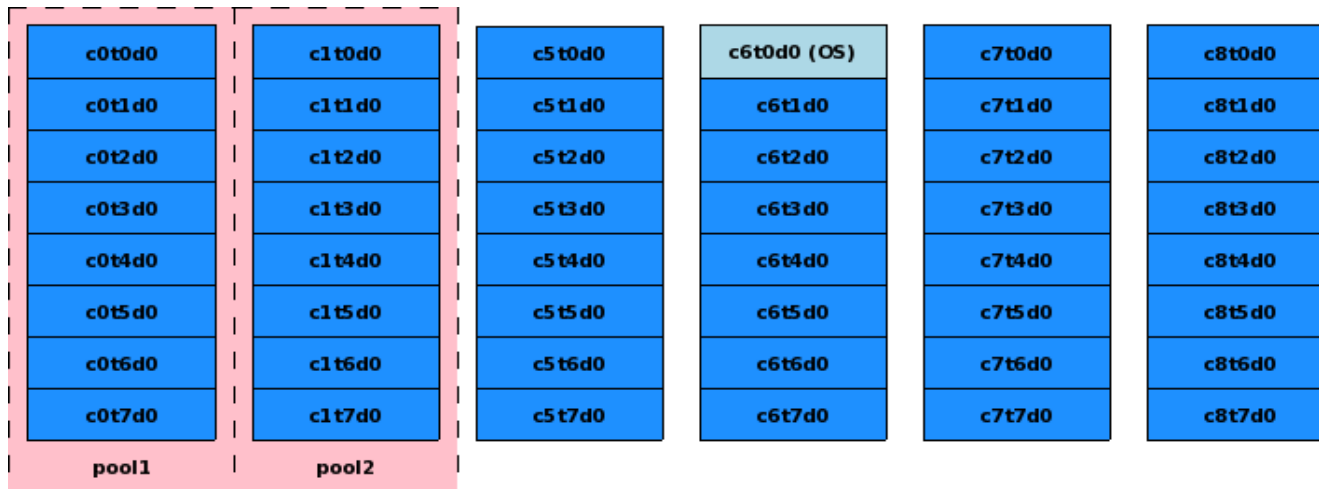
	Read(MB/s)	Write (MB/s)
raidz – 0 spare	531.6 +- 5.1	311 +- 16
raidz – 1 spare	538.0 +- 5.4	319.3 +- 8.0
raidz2 – 0 spare	509 +- 46	288 +- 13
raidz2 – 1 spare	532.2 +- 9.4	305 +- 15

Single zfs, single pool (results)

Single ZFS IO Performance




Single zfs, multiple pools (1)

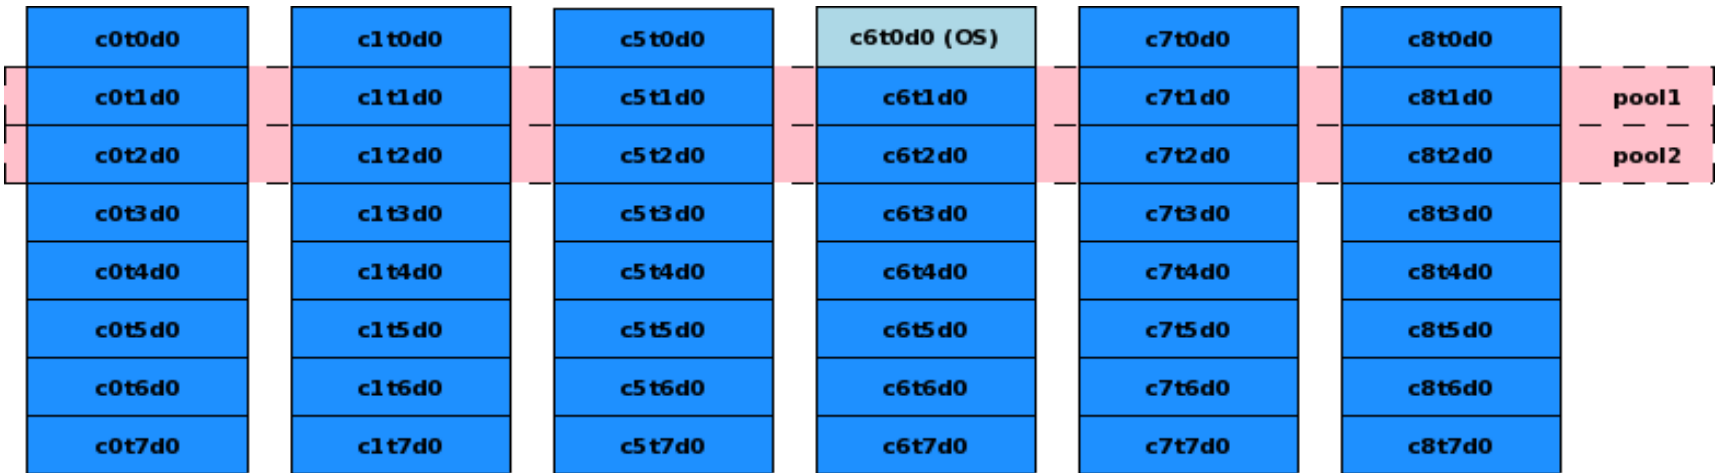


Single ZFS IO Performance

Multiple pools (controller-based)




Single zfs, multiple pools (2)



Single ZFS IO Performance

Multiple pools (sliced)



Multiple zfs

	# procs	Read(MB/s)	Write (MB/s)
6 zfs, 1 pool per controller	6	926	693
8 zfs, 1 pool per “slice”	8	953	660
3 zfs, 2 pools per zfs	3	891	644
3 zfs, 2 pools per zfs	6	925	712

Thumper conclusions

Ideal configuration: 6 zfs filesystems, each with 1 pool where all disks in the pool are located on the same controller ... or ... 3 zfs filesystems, each with 2 pools where all disks in the pool are located on the same controller

- * implies 3 or 6 dcache pools**

TODO:

- * Benchmark 6 zfs with >> 6 read/write processes**
- * Benchmark network**
- * Investigate the x4540 systems**
 - * <http://www.sun.com/servers/x64/x4540/>**