GEOM theoretical performance and scalability limits

Based on ZERO GEOM class.

```
# kldload geom_zero
# sysctl kern.geom.zero.clear=0
# diskinfo -v /dev/gzero
/dev/gzero
        512                     # sectorsize
        1152921504606846976     # mediasize in bytes (1.0E)
        2251799813685248        # mediasize in sectors
```

```
CPU Intel(R) Xeon(R) CPU X5272 @ 3.40GHz
FreeBSD/SMP: Multiprocessor System Detected: 4 CPUs
FreeBSD/SMP: 2 package(s) x 2 core(s)
real memory = 8589934592 (8192 MB)
```

```
static void
g_zero_start(struct bio *bp)
{
        int error;

        switch (bp->bio_cmd) {
        case BIO_READ:
                if (g_zero_clear)
                        memset(bp->bio_data, g_zero_byte, bp->bio_length);
                /* FALLTHROUGH */
        case BIO_WRITE:
        case BIO_DELETE:
                bp->bio_completed = bp->bio_length;
                error = 0;
                break;
        default:
                error = EOPNOTSUPP;
                break;
        }
        g_io_deliver(bp, error);
}
```

kern.geom.collectstats=1

raidtest test -r -d /dev/gzero -n <NPROC>

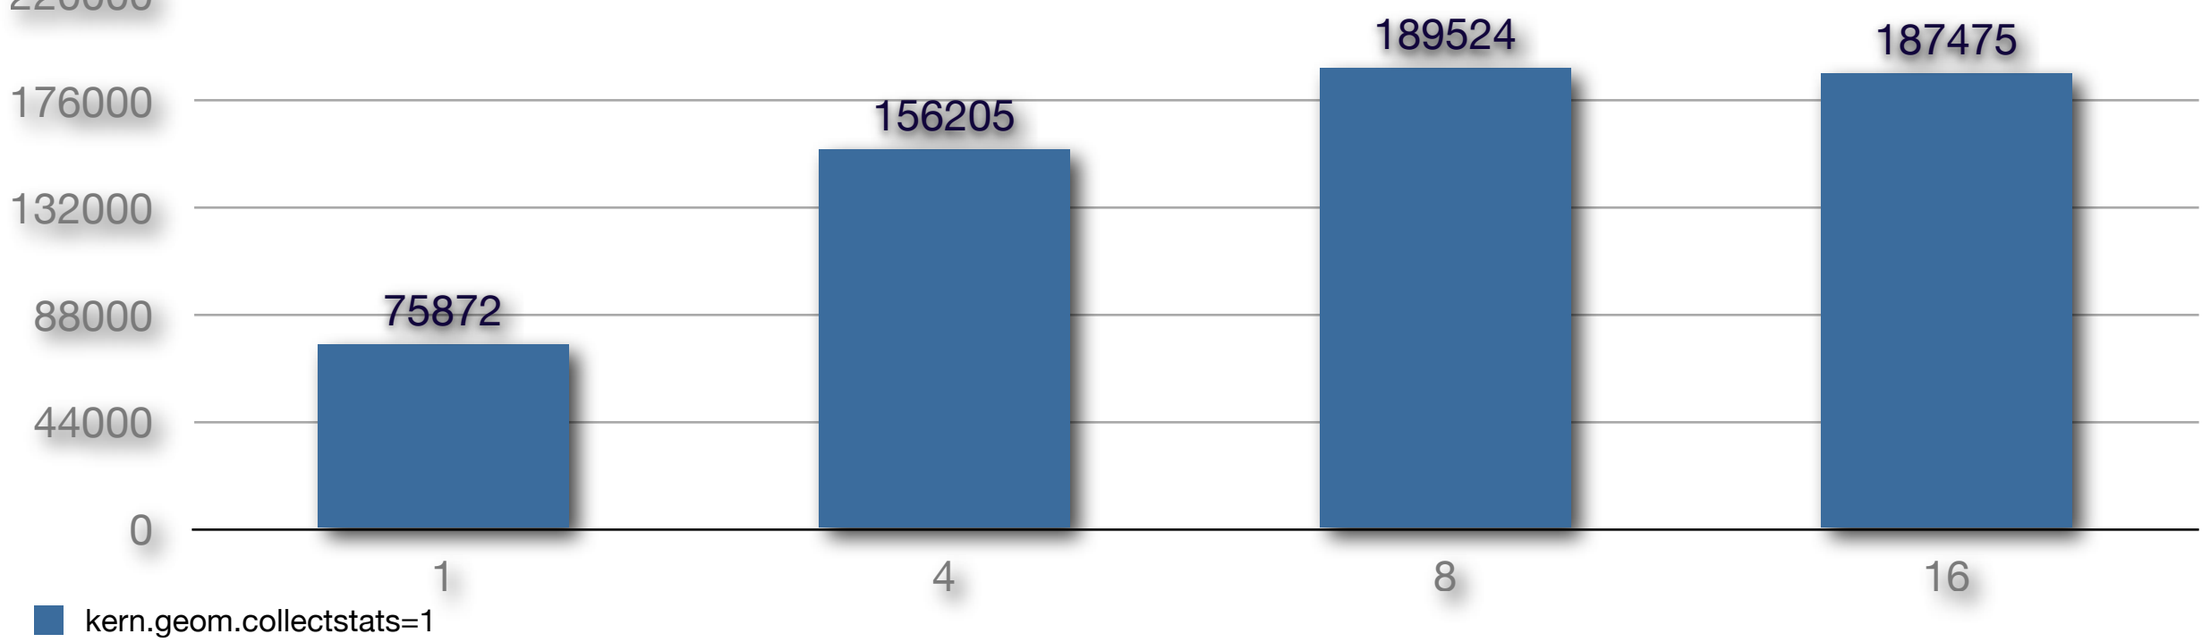

kern.geom.collectstats=0

raidtest test -r -d /dev/gzero -n <NPROC>

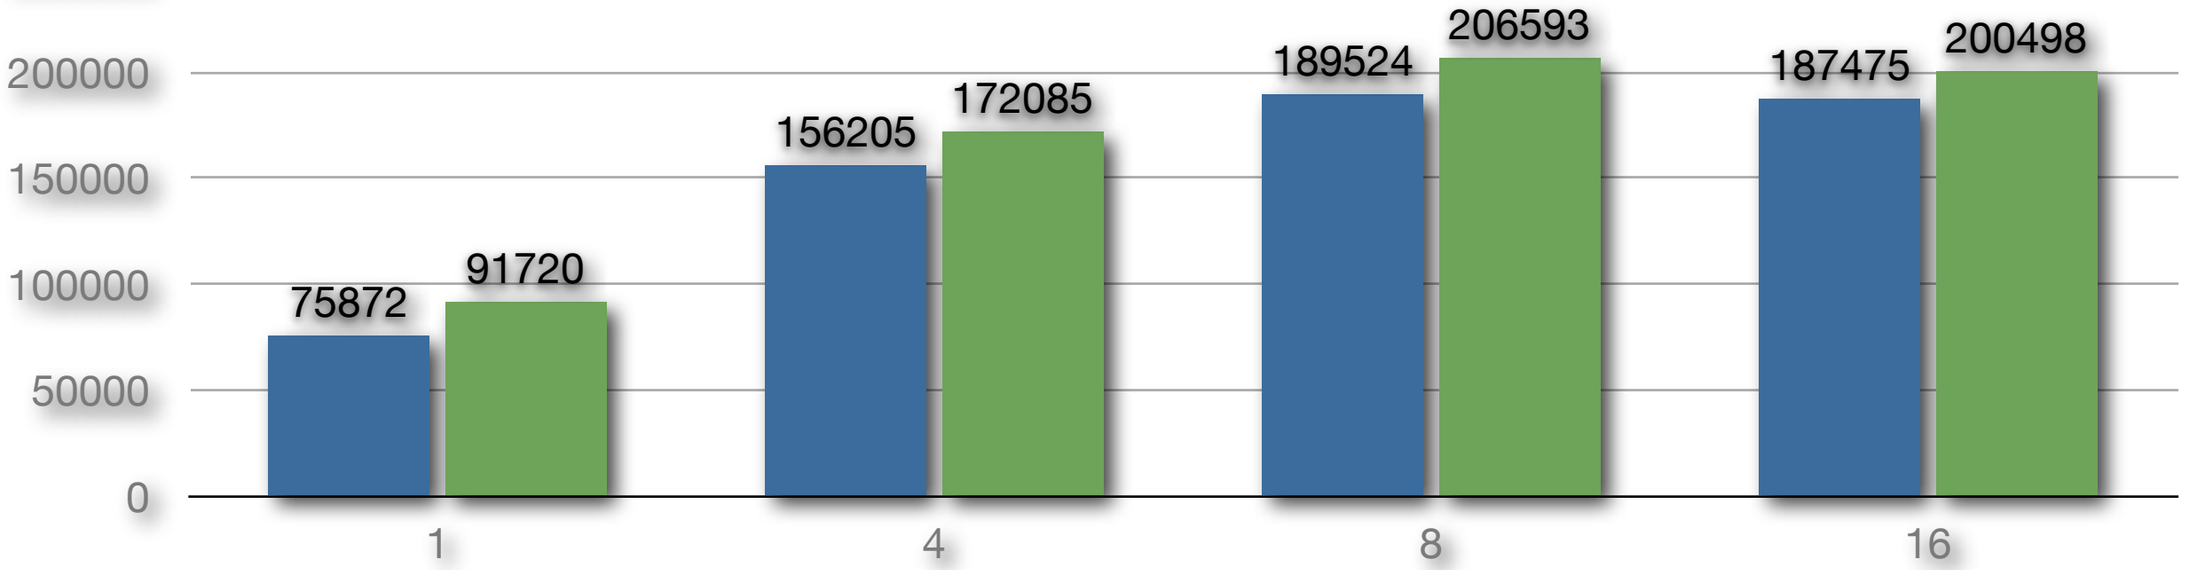

debug.geomperf_direct_io=1

Eliminates the g_up/g_down threads and passes the I/O requests directly in both directions.

raidtest test -r -d /dev/gzero -n <NPROC>

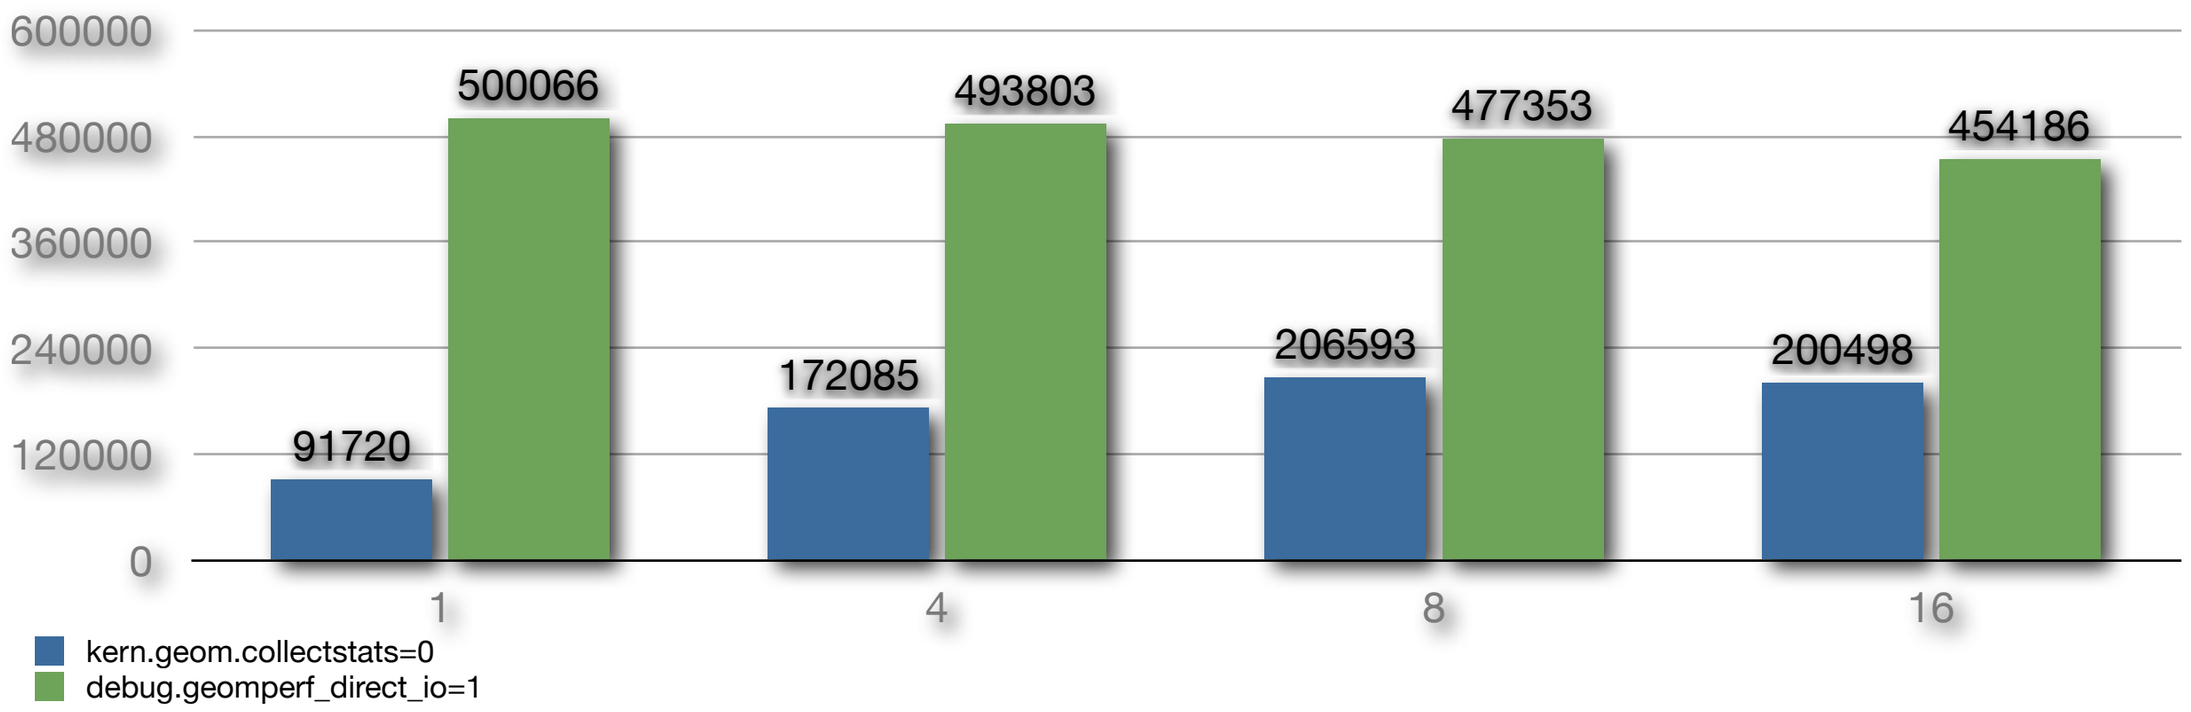

debug.geomperf_dev_eternal=1

Adds the MAKEDEV_ETERNAL flag to make_dev_p() call in GEOM DEV class.
This avoids contentions on the global devmtx mutex.

raidtest test -r -d /dev/gzero -n <NPROC>

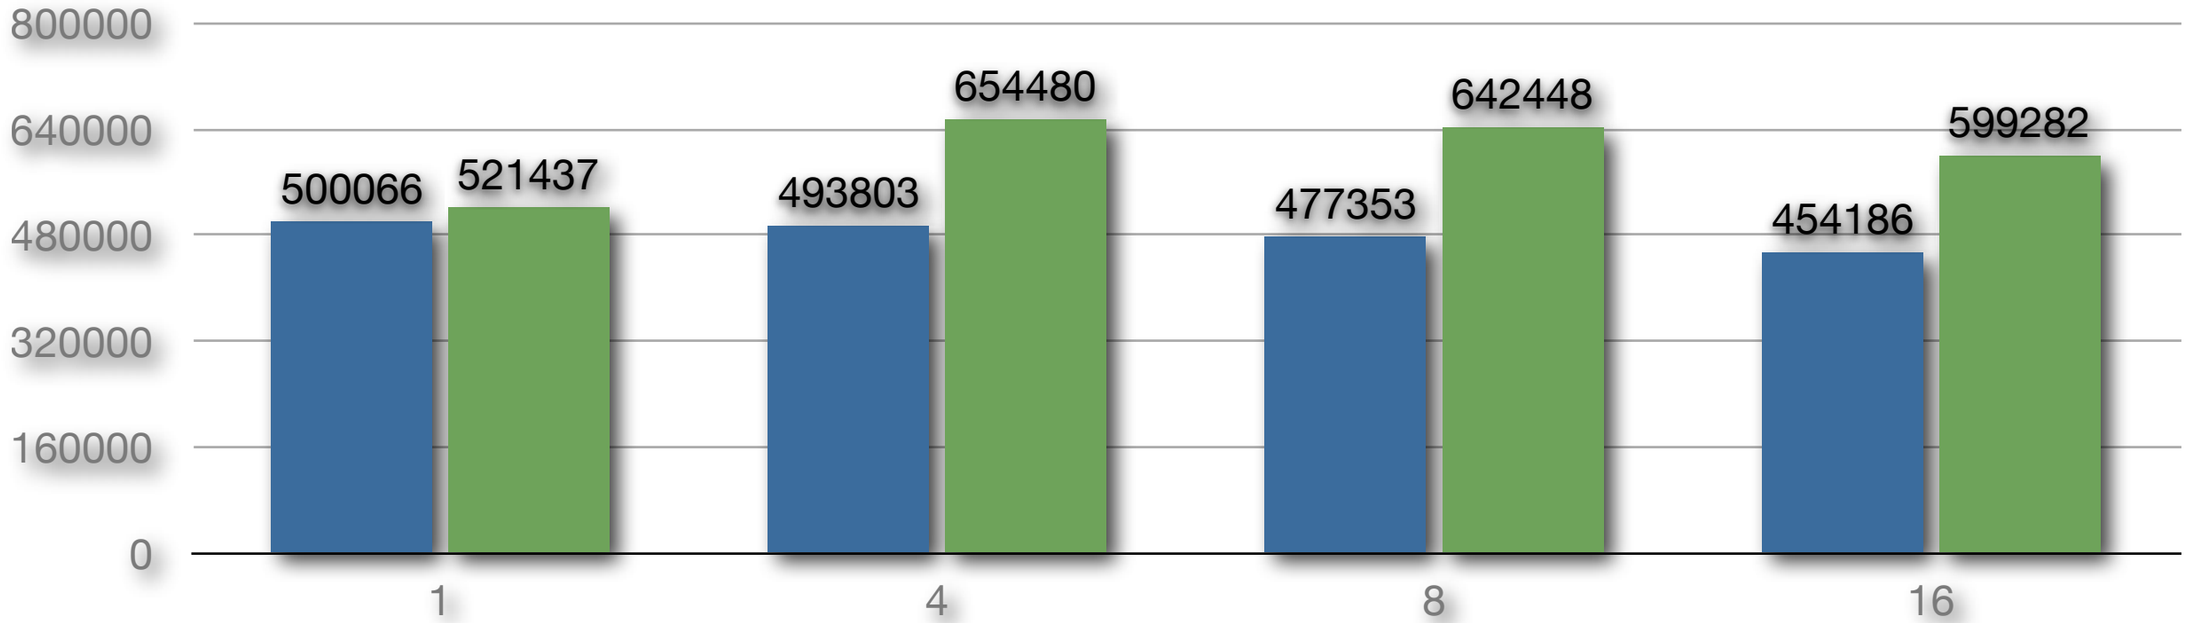

debug.geomperf_percpu_pbufs=1

Uses per-CPU pbufs.
This avoids contentions on the global pbuf_mtx mutex.

raidtest test -r -d /dev/gzero -n <NPROC>

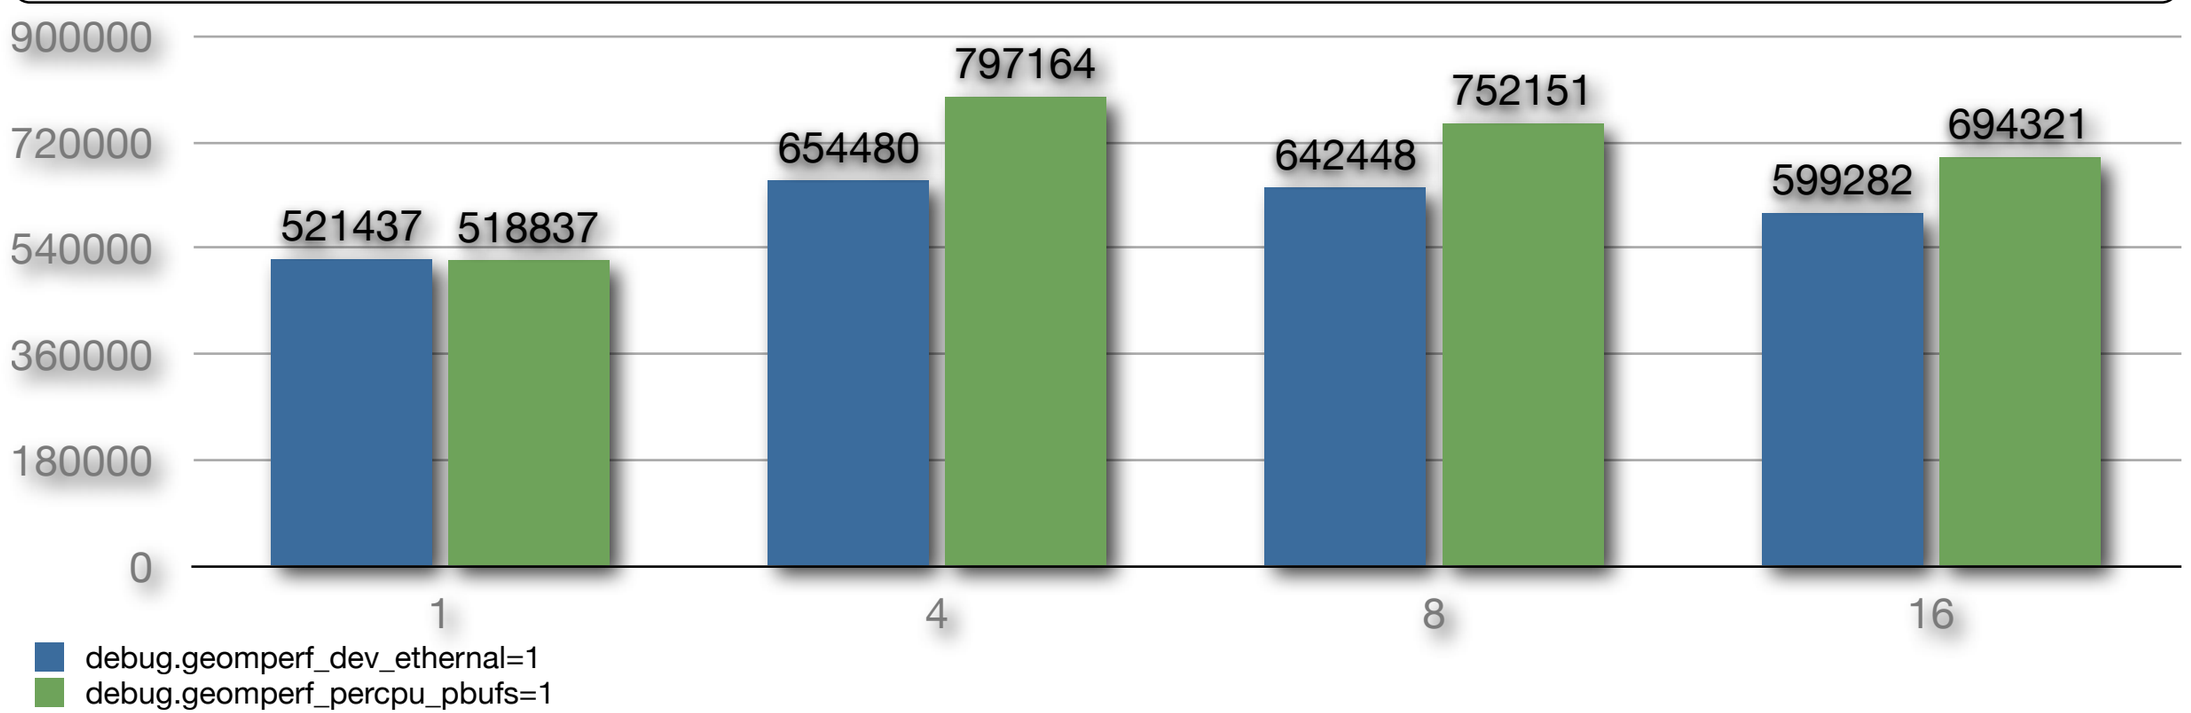

-#define MAX_POOL_SLEEP_SIZE 128
+#define MAX_POOL_SLEEP_SIZE 1024

This reduces contention on mtxpool.

raidtest test -r -d /dev/gzero -n <NPROC>

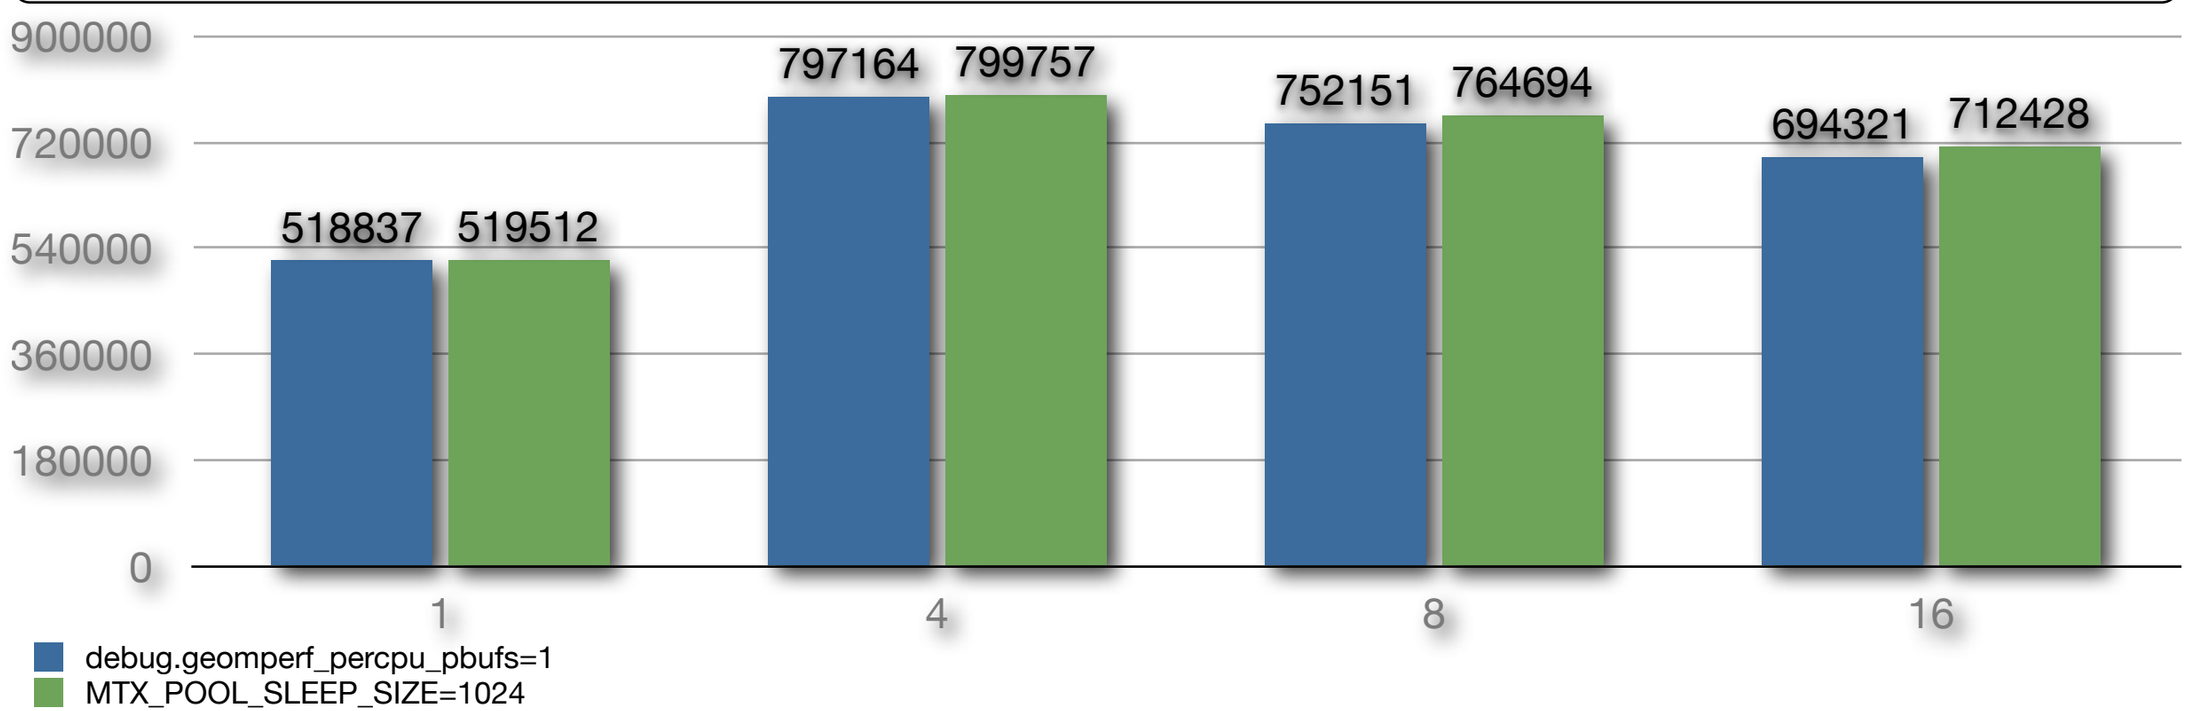

-#define pa_index(pa) ((pa) >> PDRSHIFT)
+#define pa_index(pa) ((pa) >> 17)

This reduces contention on vm_page lock.

raidtest test -r -d /dev/gzero -n <NPROC>

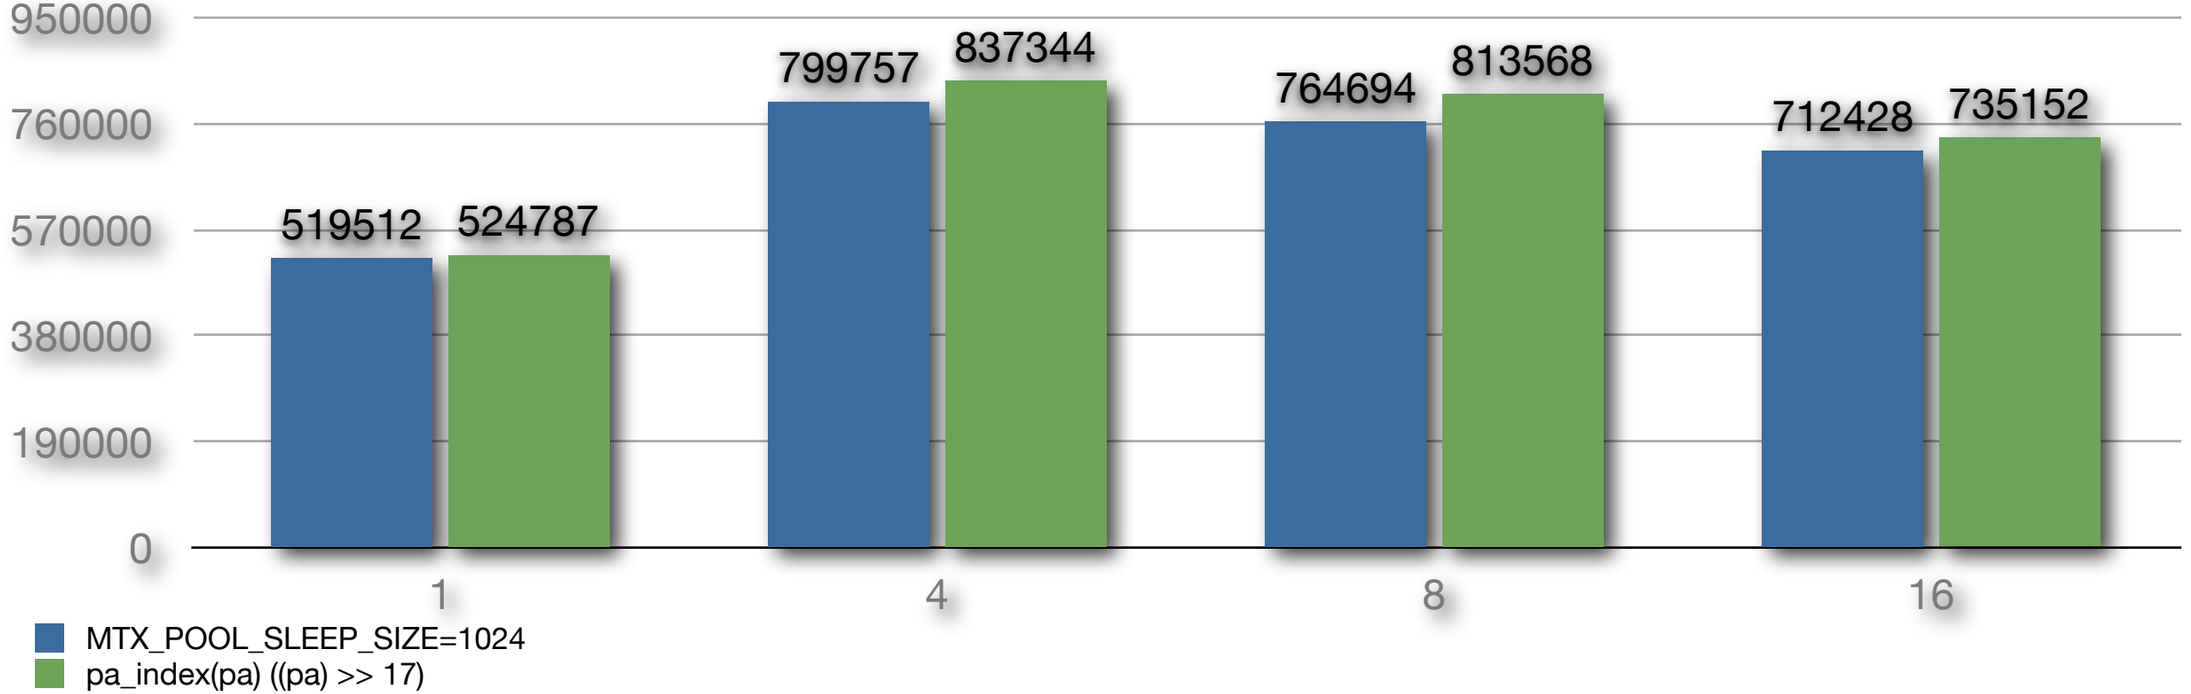

raidtest test -r -d /dev/gzero -n <NPROC>

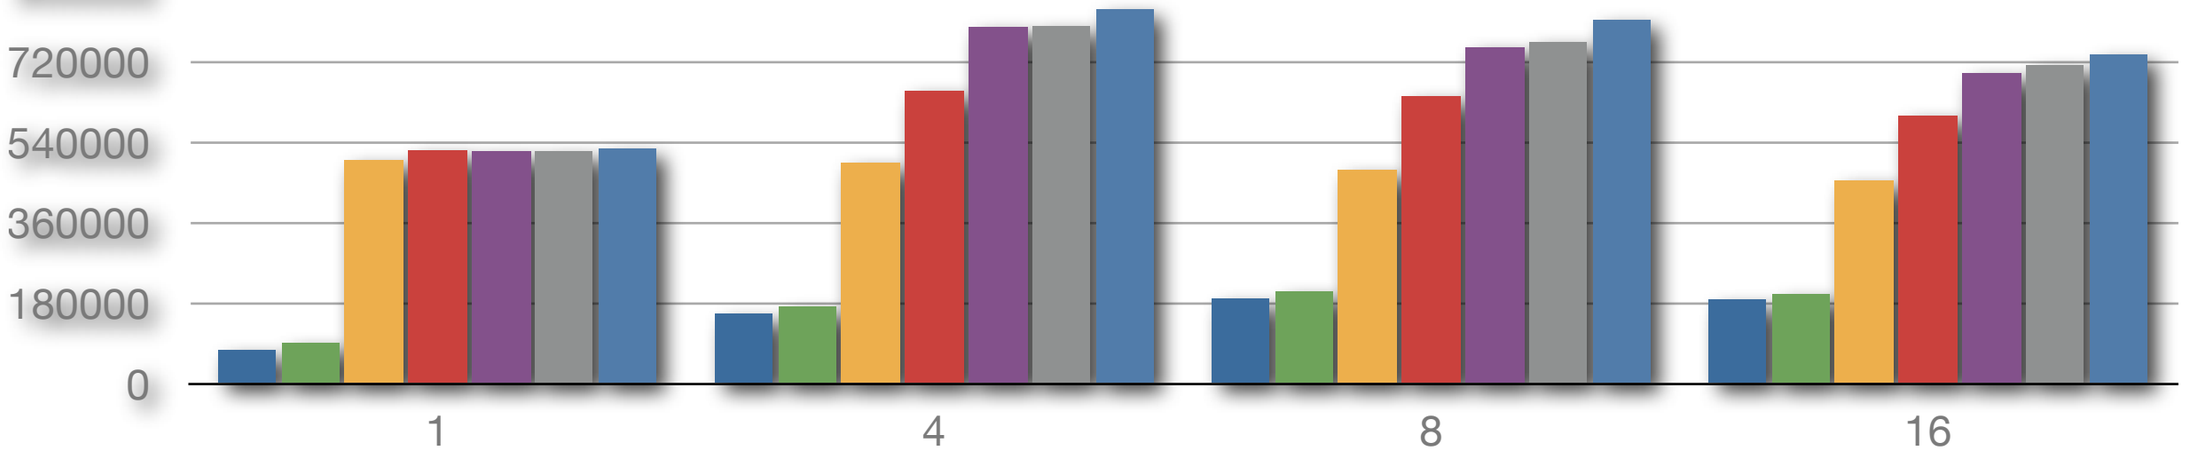

raidtest test -r -d /dev/gzero -n <NPROC>

Difference at 95.0% confidence
681139 +/- 13138.2
436.055% +/- 8.41088%
(Student's t, pooled s = 5796.45)

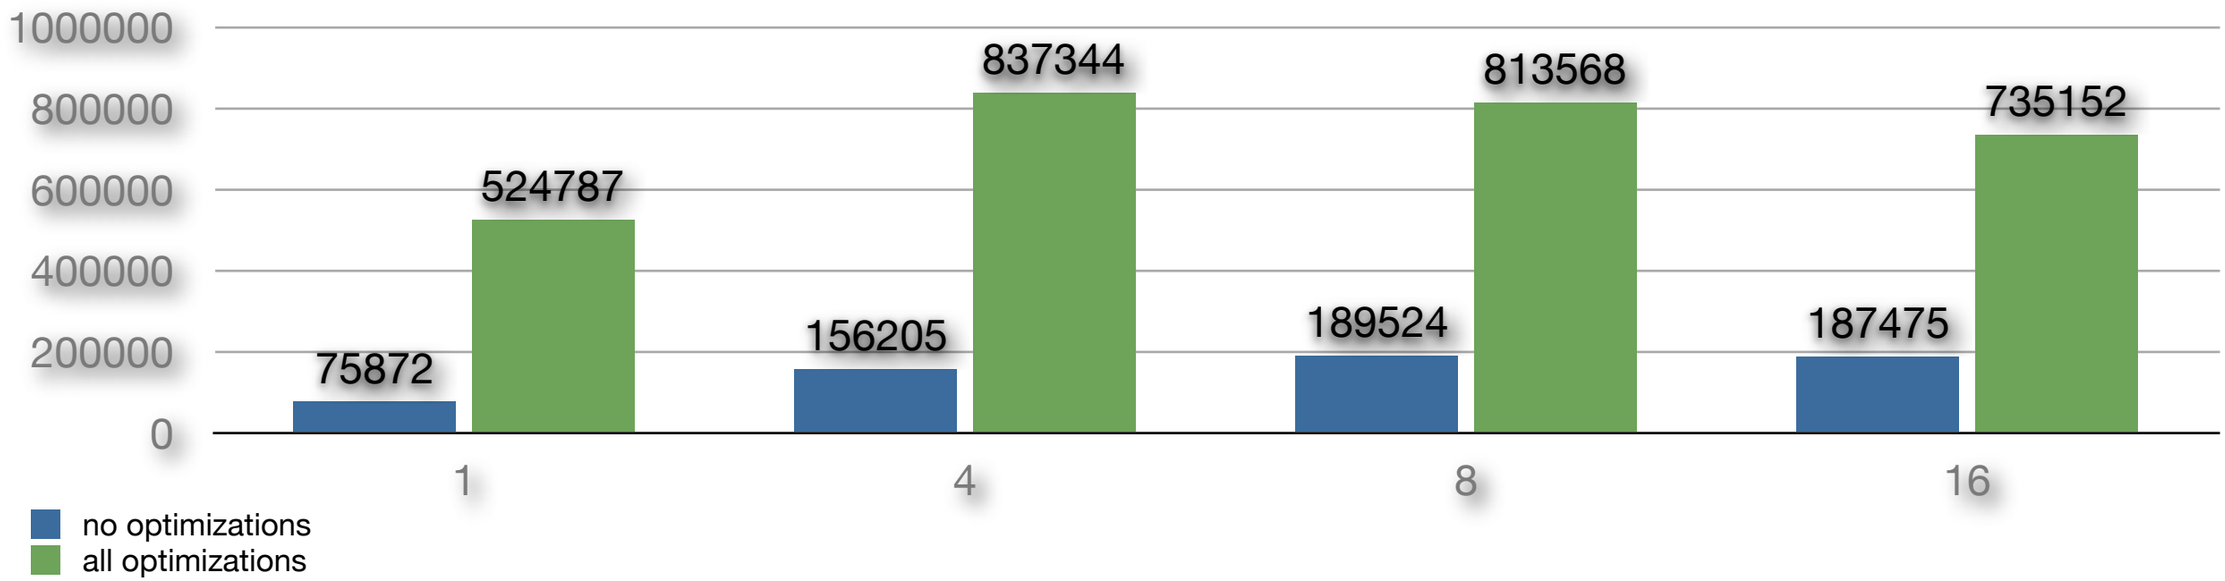

	1	4	8	16
kern.geom.collectstats=1	75872	156205	189524	187475
kern.geom.collectstats=0	91720	172085	206593	200498
debug.geomperf_direct_io=1	500066	493803	477353	454186
debug.geomperf_dev_ethernal=1	521437	654480	642448	599282
debug.geomperf_percpu_pbufs=1	518837	797164	752151	694321
MTX_POOL_SLEEP_SIZE=1024	519512	799757	764694	712428
pa_index(pa) ((pa) >> 17)	524787	837344	813568	735152
kern.geom.collectstats=1	75872	156205	189524	187475
pa_index(pa) ((pa) >> 17)	524787	837344	813568	735152