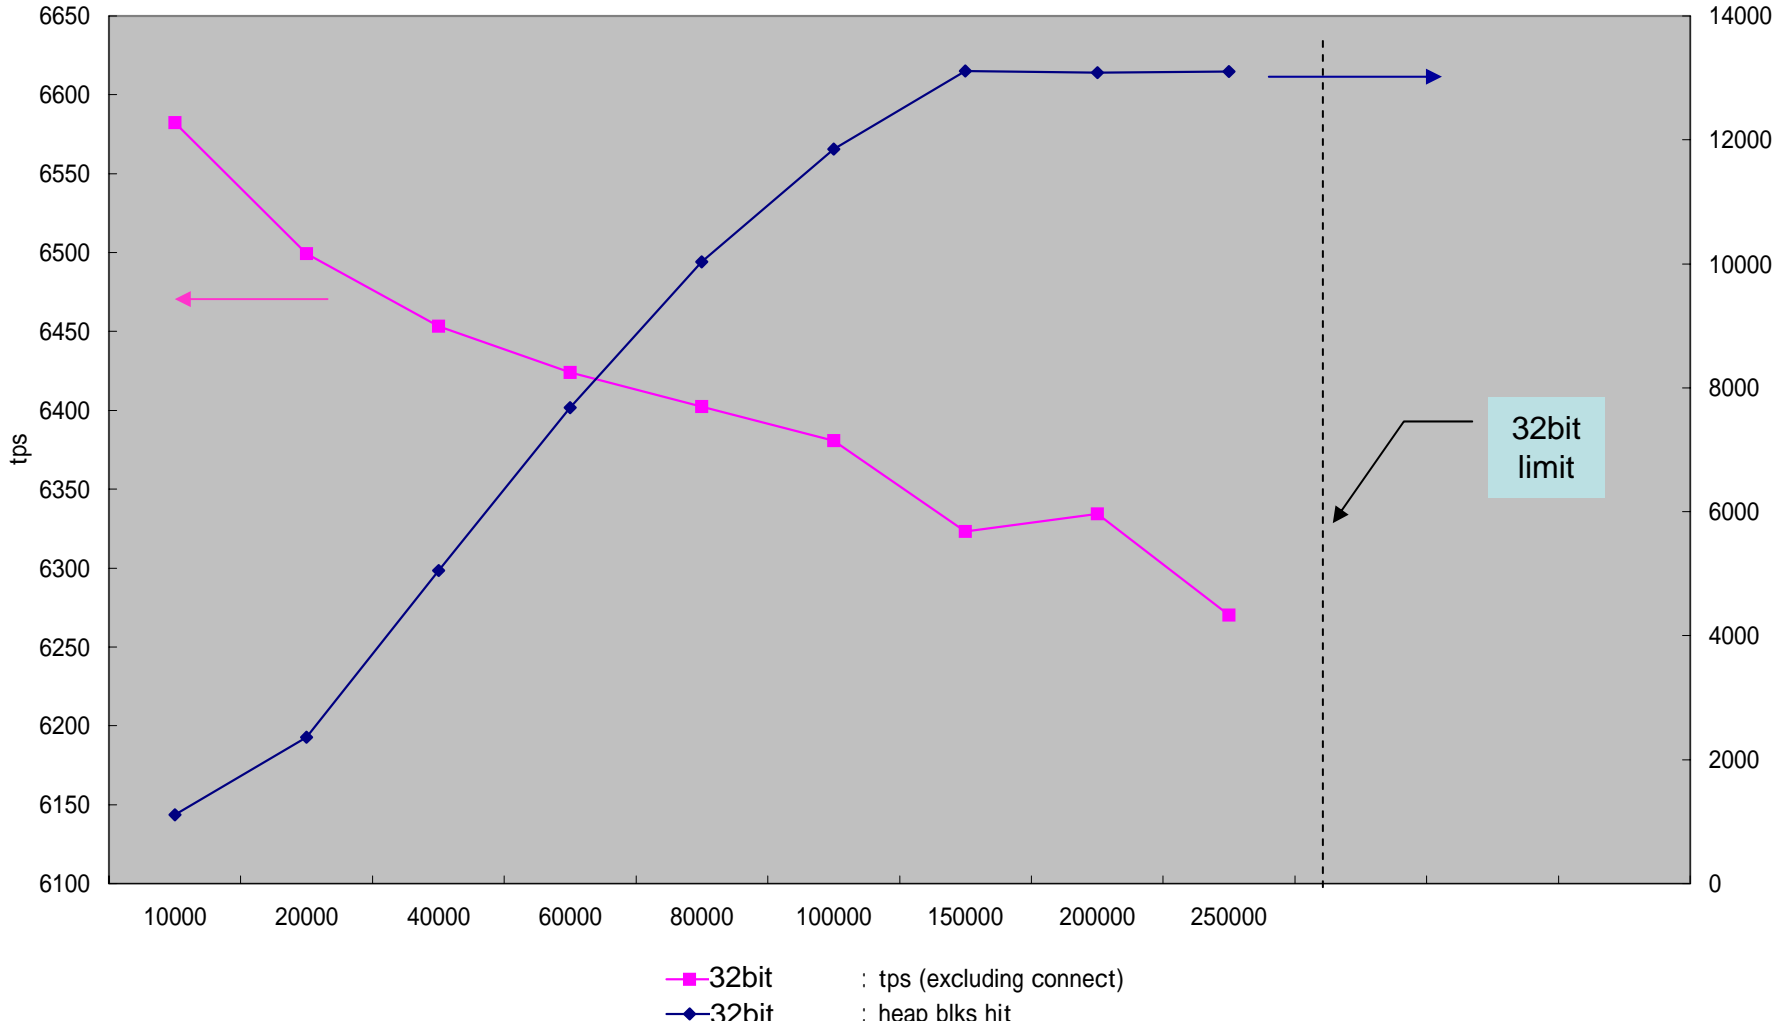# Benchmark Machine Spec

- CPU: Opteron 1.4GHz x 2, cache 1MB each
- MEM: 8GB
- HDD: IDE 120GB

Pgbench-1

shmem and XID: 32bit, PostgreSQL 8.0.1, cold start tps

pgbench (SELECT only, scaling factor = 200, client = 10, transaction = 10000)

32bit
limit

-■-32bit   tps (excluding connect)
-◆-32bit   heap blks hit

Pgbench-2

Shmem and XID: 32bit (original) PostgreSQL 8.0.1, Hot start (issued select (count *) from accounts; in advance)
Throuput and heap_blks_hit improves as we have more shmem.   It seems that more tuples reside in memory (not simple though).
In cold start, more memory doesn't improve the throuput. Maybe maitaining shmem will cost more than its benefit in this case.

pgbench (SELECT only, scaling factor = 200, client = 10, transaction = 10000,

| | 32bit | tps (excluding connect) |
| | 32bit | heap blks hit |
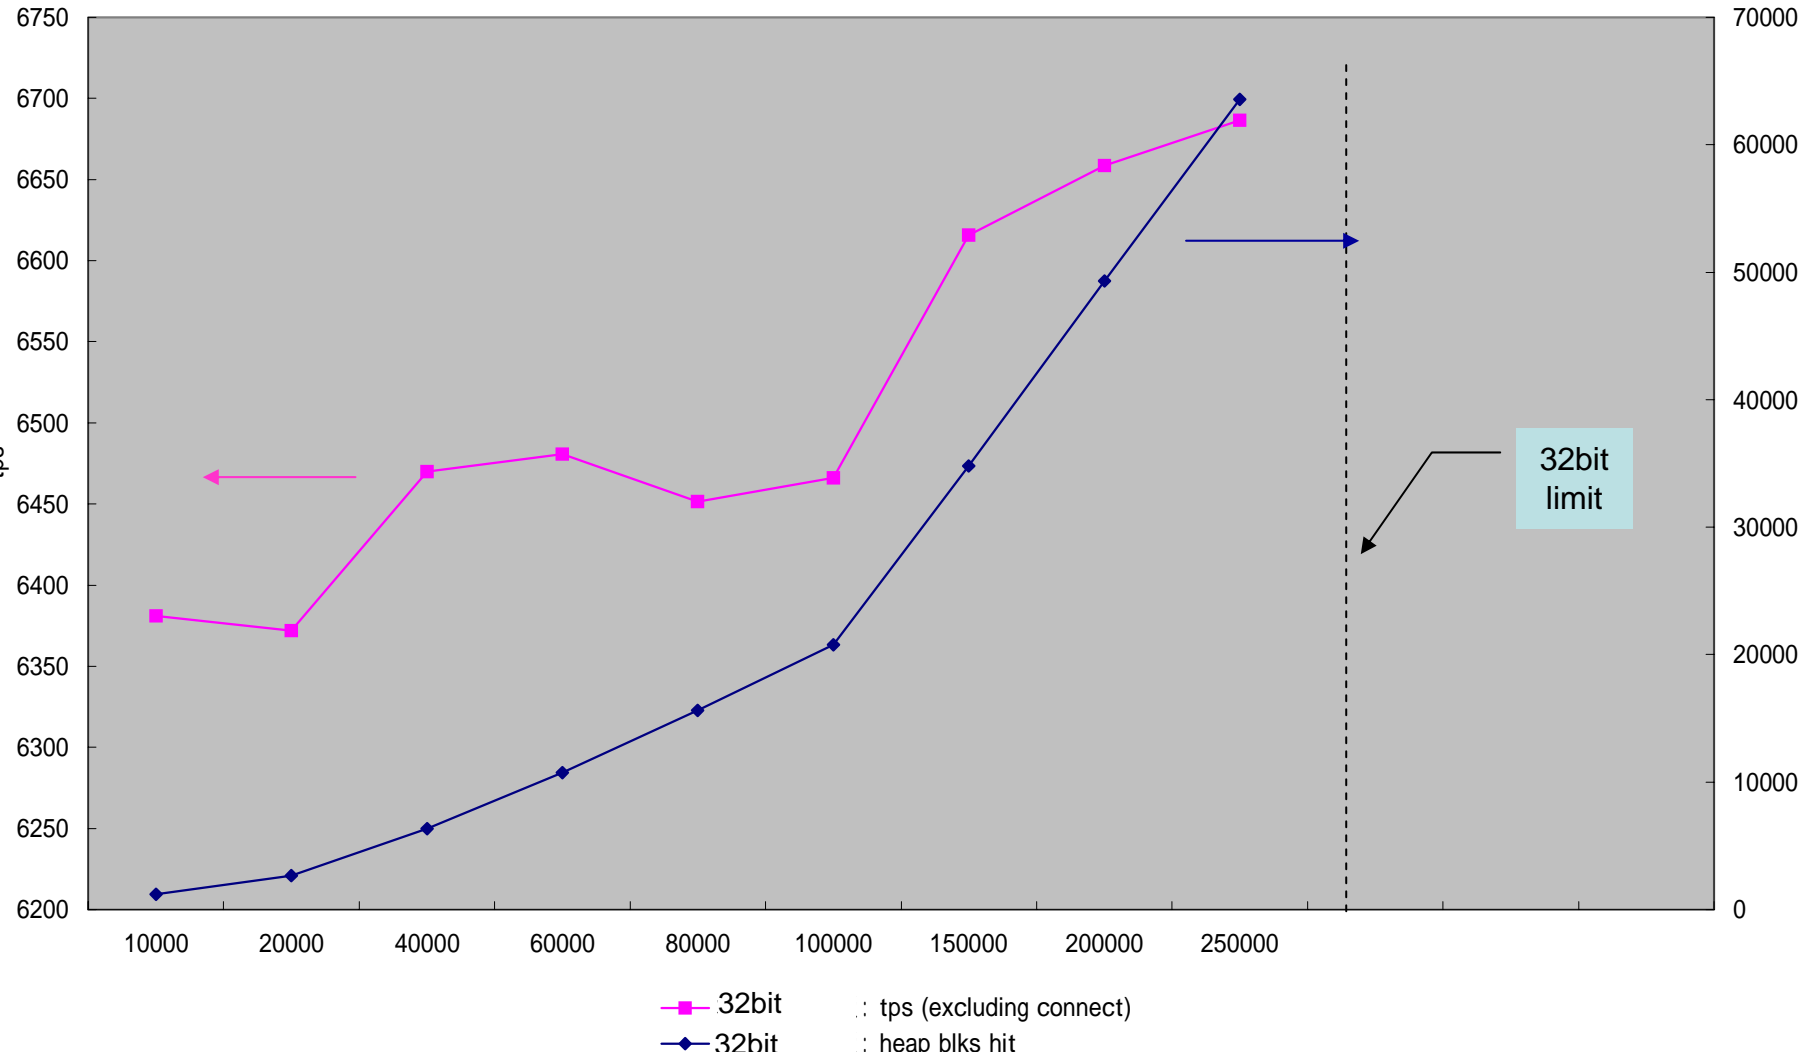
Pgbench-3

Shmem and XID 64bit (based on 8.0.1), hot start (issued select (count *) from accounts; in advance).

More shmem gets better throuput and heap_blks_hit beyond the 2GB limit, as seen in pgbench-2.

On the other hand, within 2GB limit, throuput and heap_blks_hit is slightly worth than 32bit mode, may be caused by 64bit overhead of XID.

pgbench (SELECT only, scaling factor = 200, client = 10, transaction = 10000,

32bit
limit
(2GB)

64bit	tps (excluding connect)	
64bit	heap blks hit	