Measuring shared_buffers

Greg Dostatni, DBA
COMMAND PROMPT, INC.
PGCONF US 2026

A simple idea:
Measure shared_buffers

COMMAND
PROMPT, INC.

Advice: 25% of system RAM (since 2008)

COMMAND PROMPT, INC.

Experiment Design

COMMAND PROMPT, INC.

We need to go bigger

- 128 GB of RAM
- ~250GB database simulating social media website
- 1 hour test runs with both read and write activity
- Work needs to be complex enough to require a lot of RAM

COMMAND PROMPT, INC.

Test Database

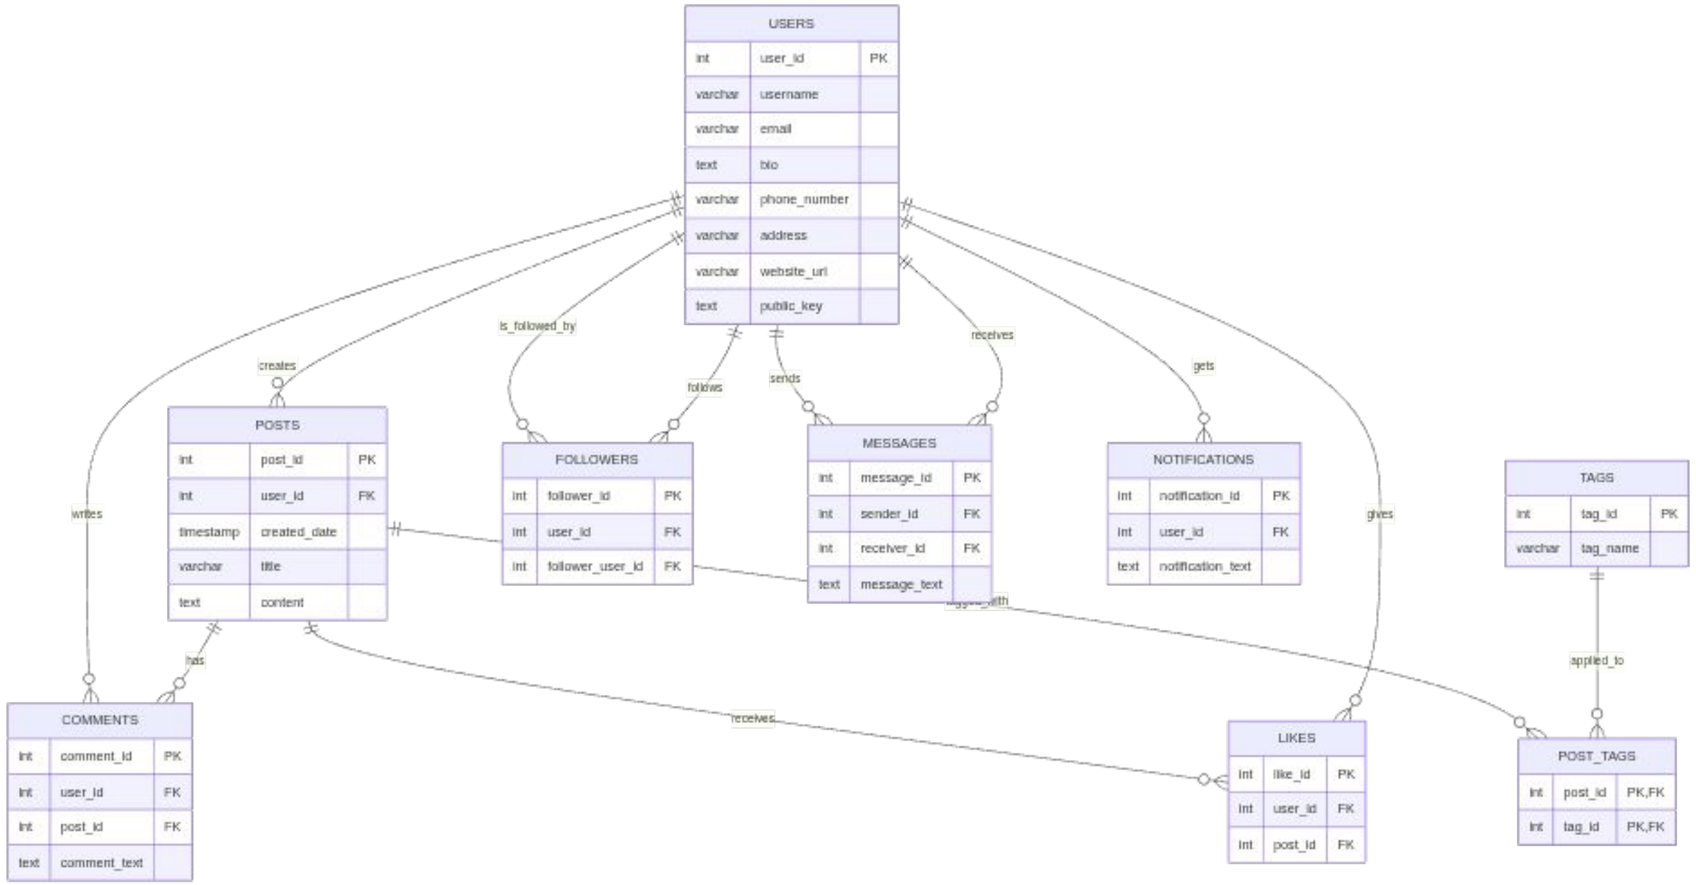

COMMAND PROMPT, INC.

Queries TPC-B vs. Analytics

```
1. BEGIN;
2. UPDATE pgbench_accounts SET abalance =
   abalance + :delta WHERE aid = :aid;
3. SELECT abalance FROM pgbench_accounts
   WHERE aid = :aid;
4. UPDATE pgbench_tellers SET tbalance =
   tbalance + :delta WHERE tid = :tid;
5. UPDATE pgbench_branches SET bbalance =
   bbalance + :delta WHERE bid = :bid;
6. INSERT INTO pgbench_history (tid, bid,
   aid, delta, mtime) VALUES (:tid, :bid,
   :aid, :delta, CURRENT_TIMESTAMP);
7. END;
```

```
1. with recent_posts AS (
2.     select posts.post_id as post_id
3.     from posts
4.     where posts.created_date >(select max(created_date)-
   interval '5 days'  from posts )
5.     and posts.user_id <> :user_id
6.     limit 100
7. )
8. select recent_posts.post_id, count(likes.like_id)
9. from recent_posts
10. join likes on likes.post_id=recent_posts.post_id
11. group by 1
12. order by 2 desc;
```

COMMAND PROMPT, INC.

Results:
TPS Reports

COMMAND
PROMPT, INC.

TPC-B (pgbench like)

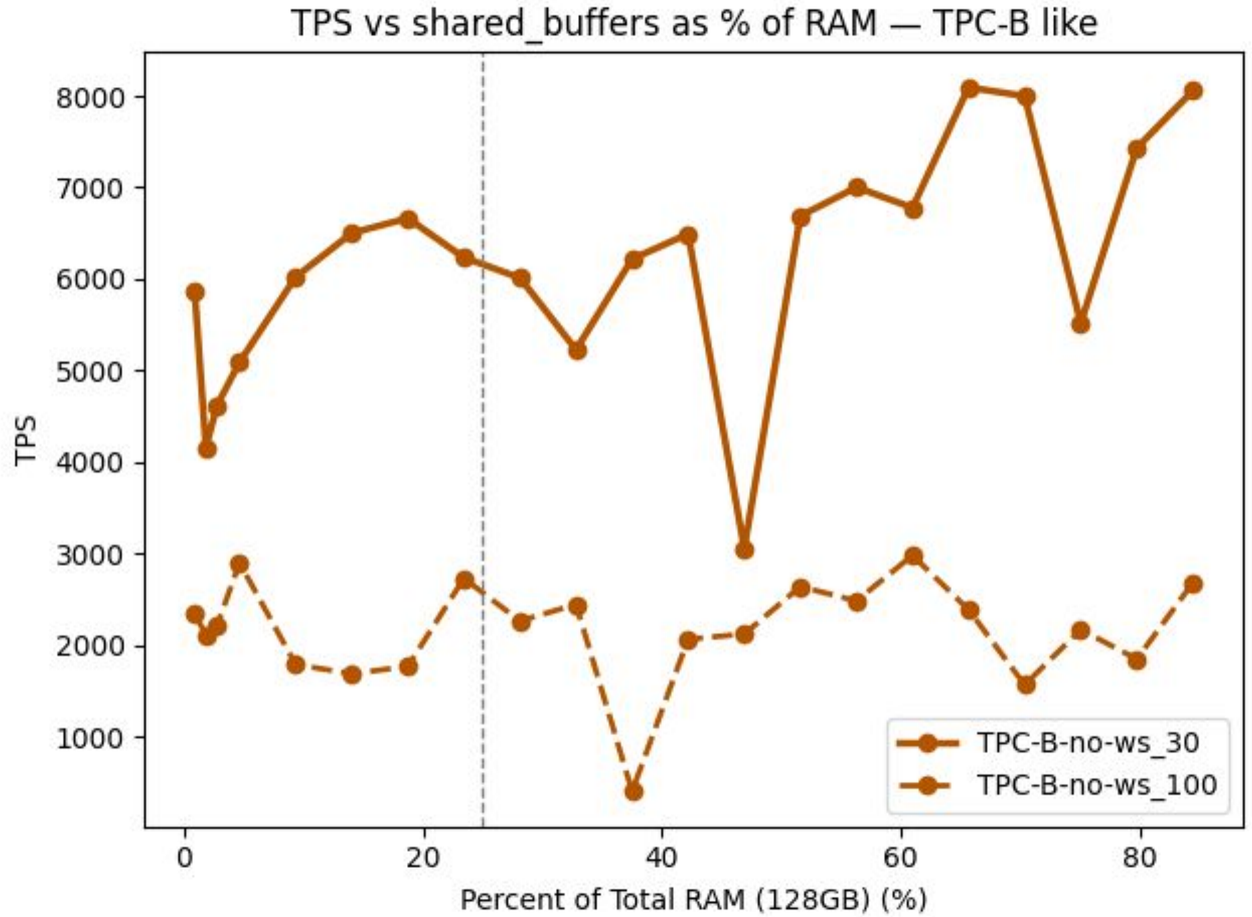

COMMAND PROMPT, INC.

Analytics

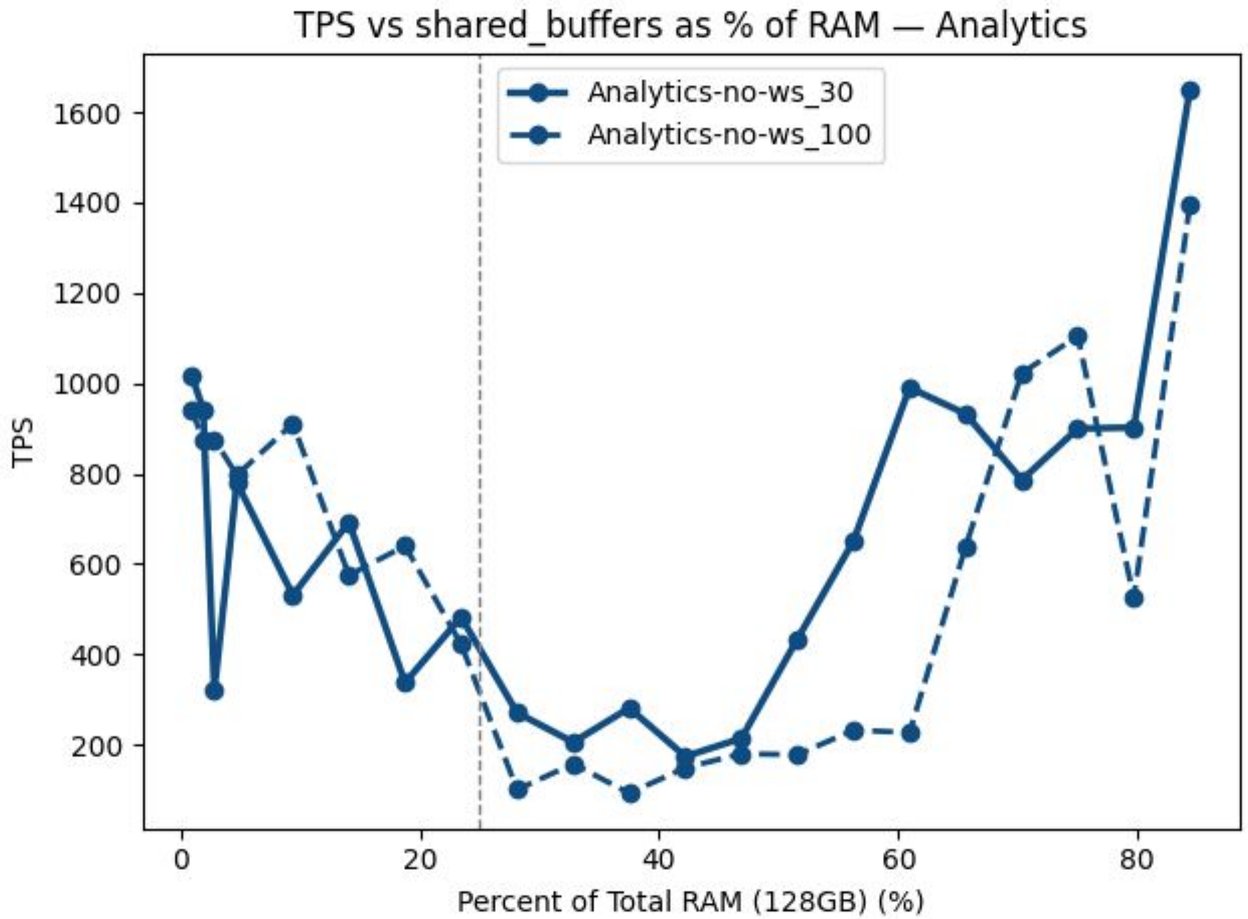

COMMAND PROMPT, INC.

Results:
TPS Reports (huge_pages)

COMMAND PROMPT, INC.

COMMAND
PROMPT, INC.

Huge Pages

Every process uses a lookup table (Translation Lookaside Buffer) to access system memory.

Huge pages	L1 Data entries	L2 Data	Max space L1 + L2
4K	64	2048	256KB + 8MB
2M	64	2048	128MB + 4GB
1G	64	64	64GB + 64GB

TPC-B (pgbench like)

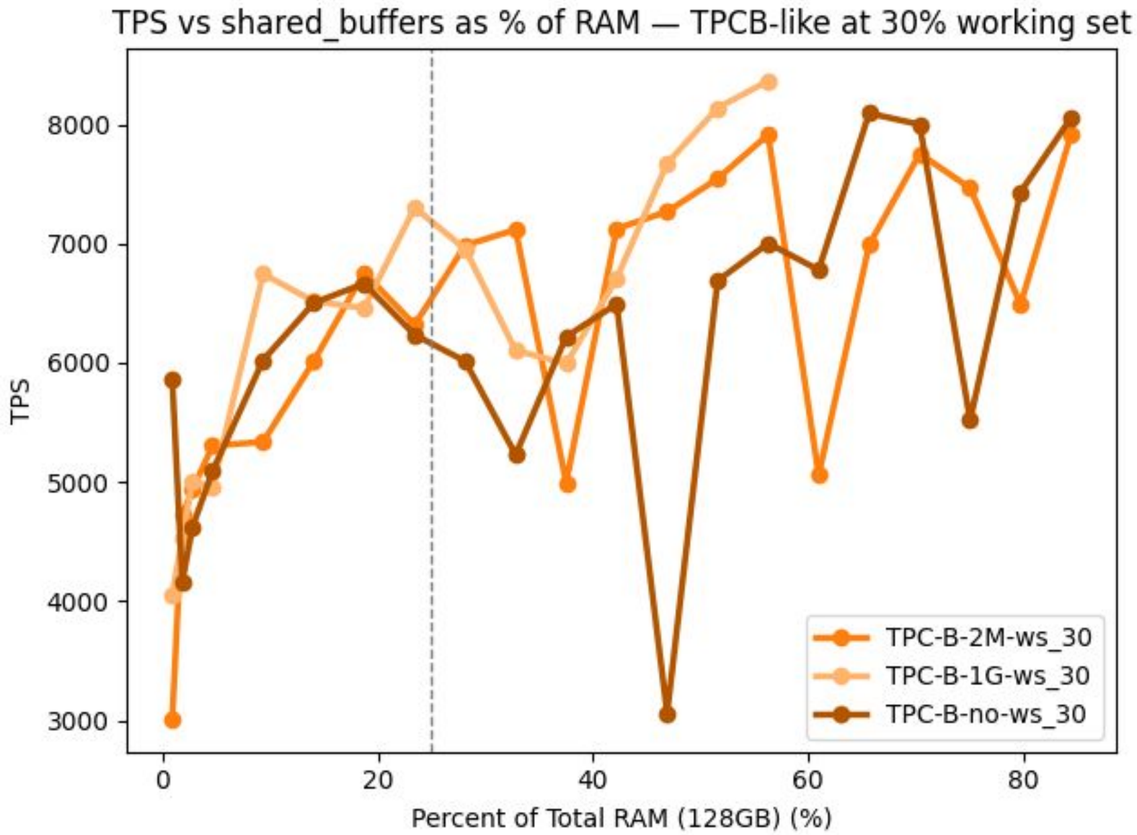

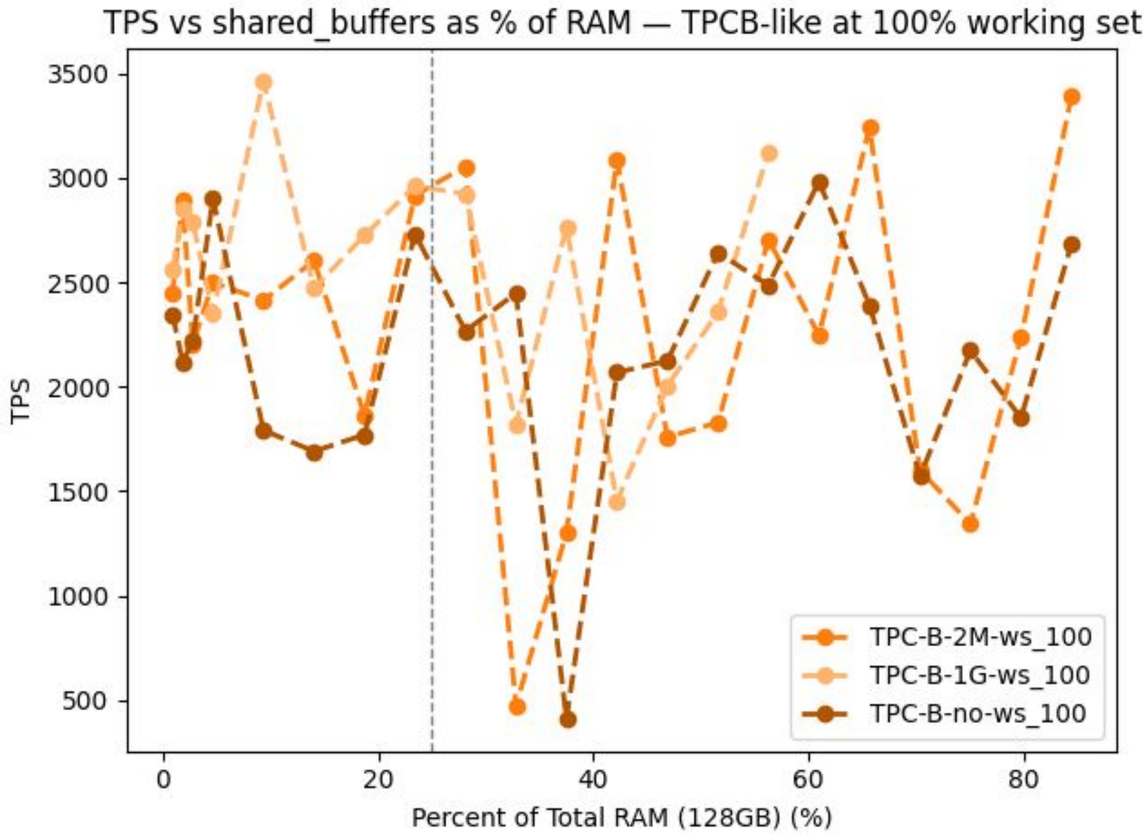

COMMAND PROMPT, INC.

Analytics

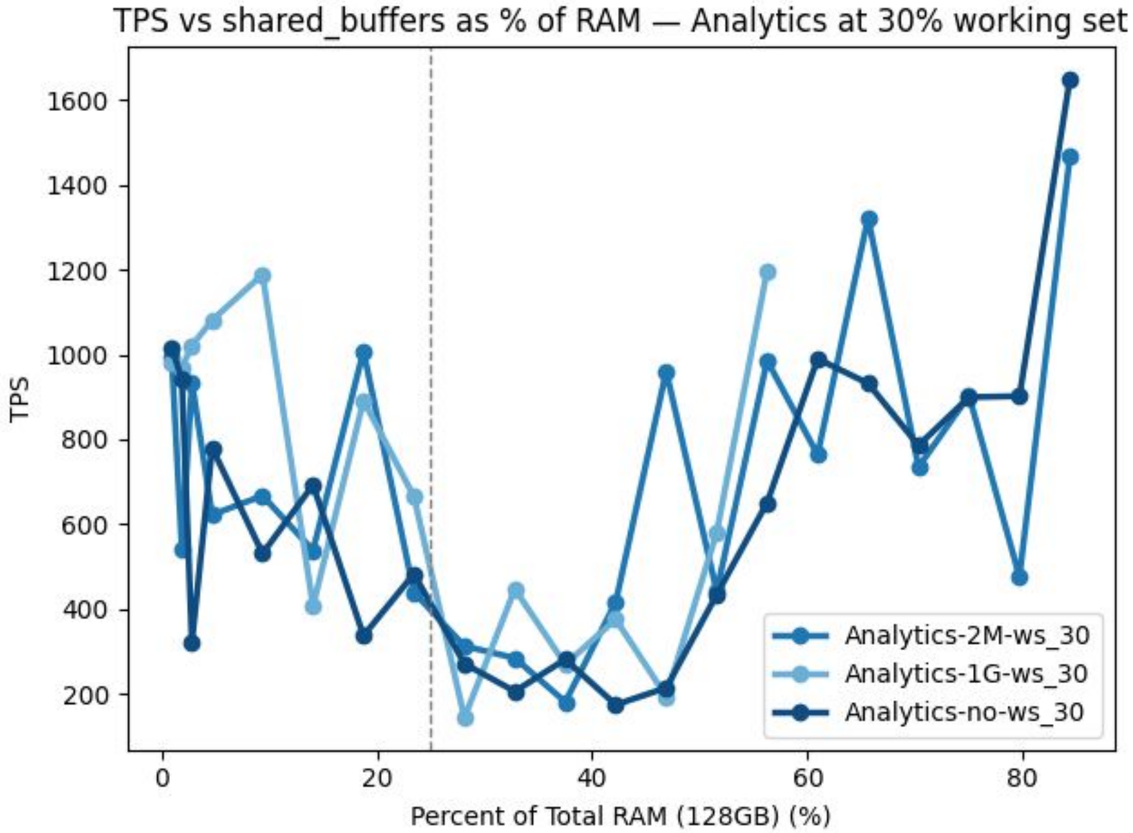

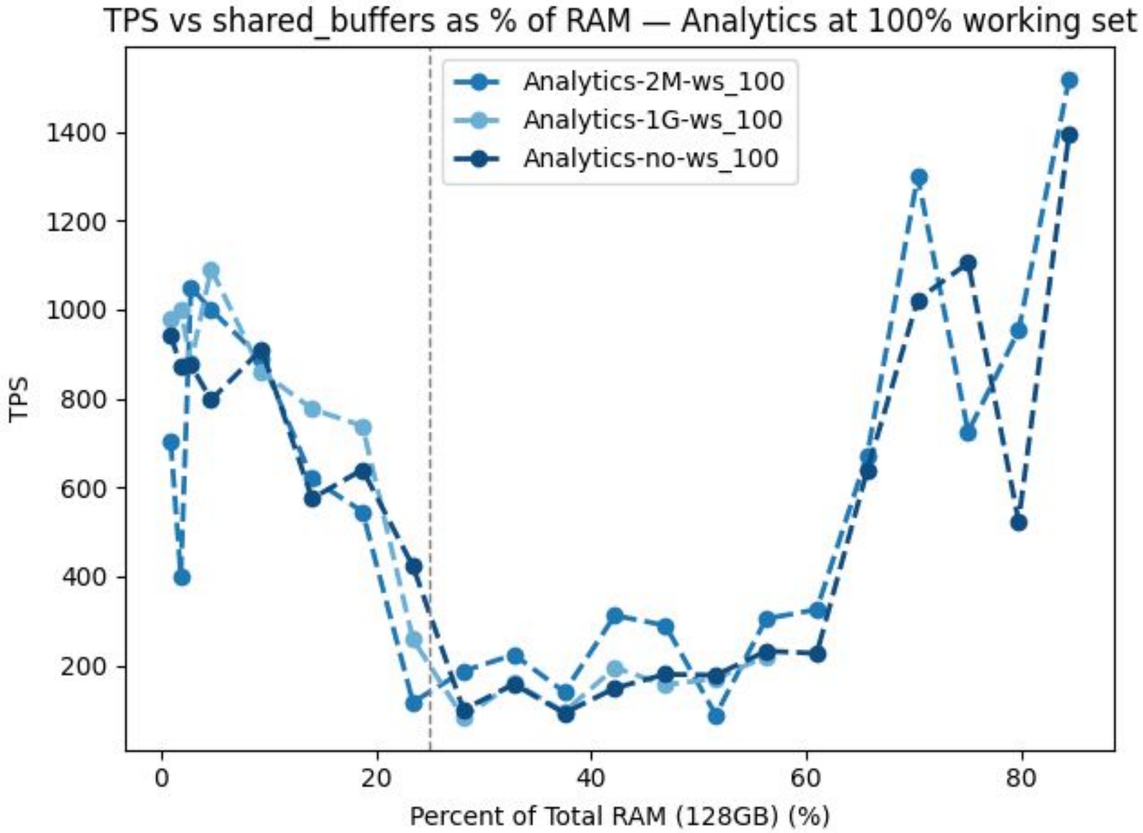

COMMAND PROMPT, INC.

Memory lost to access memory

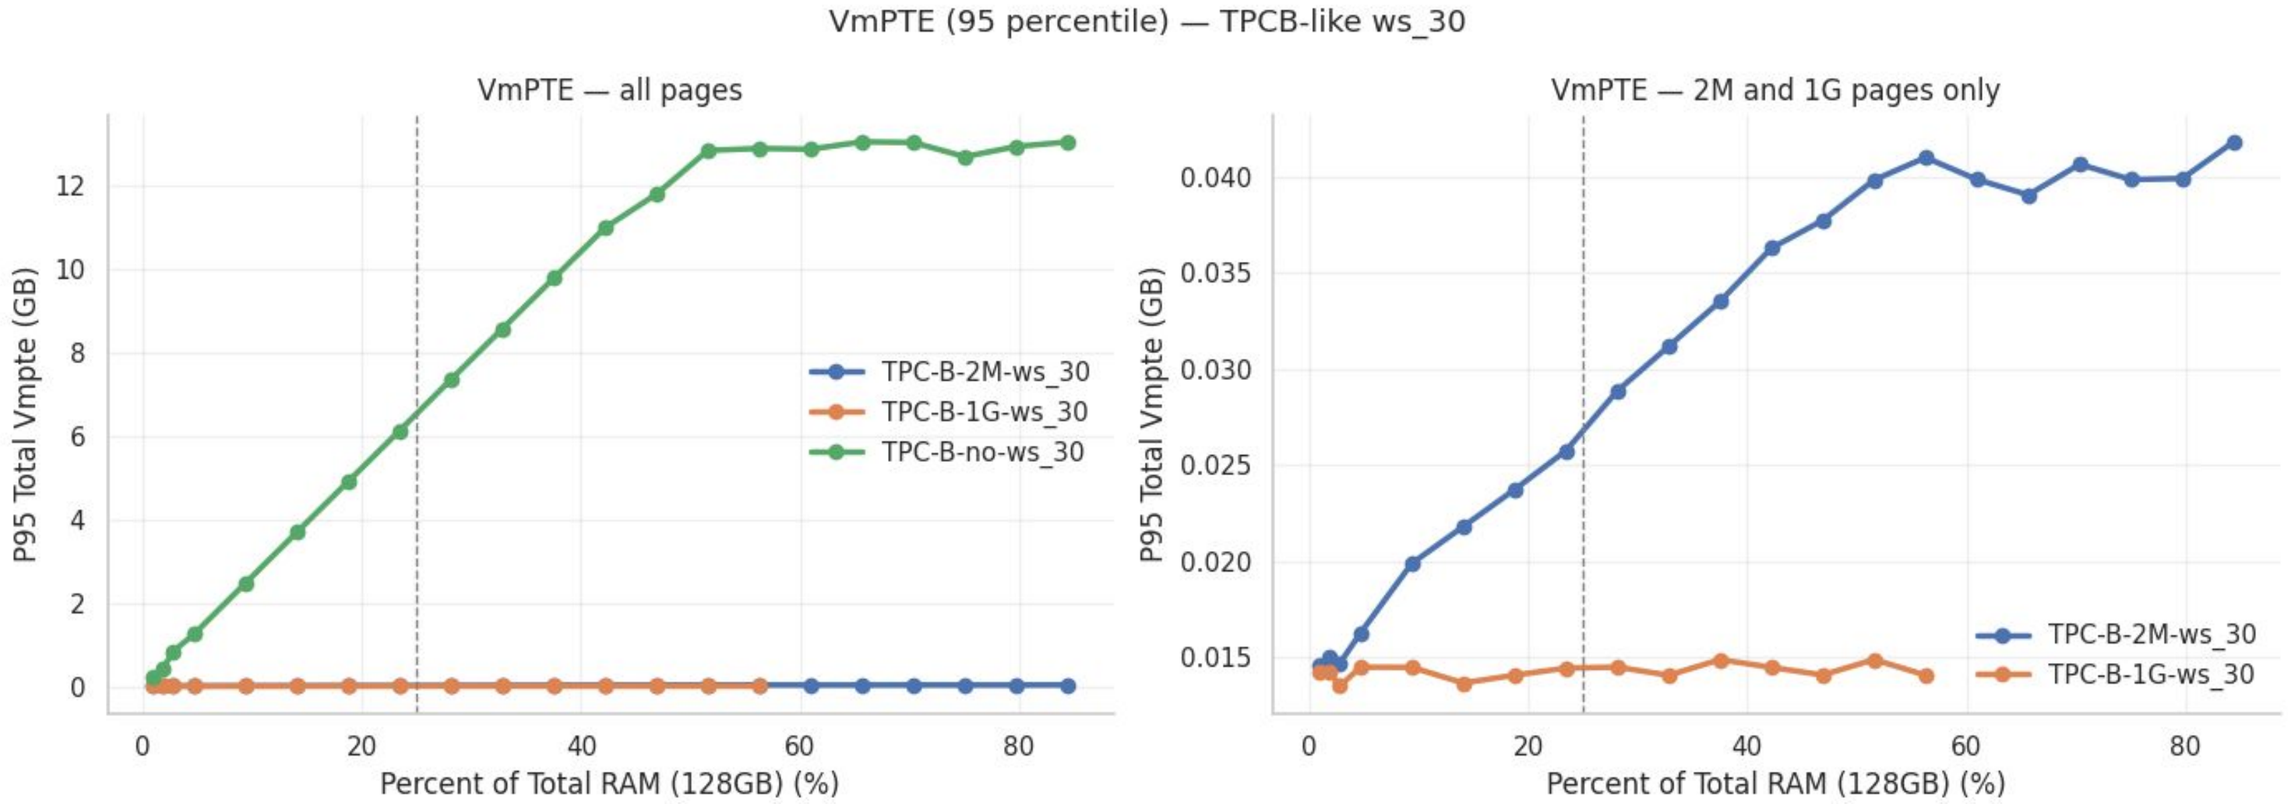

COMMAND PROMPT, INC.

Memory lost (pt 2)

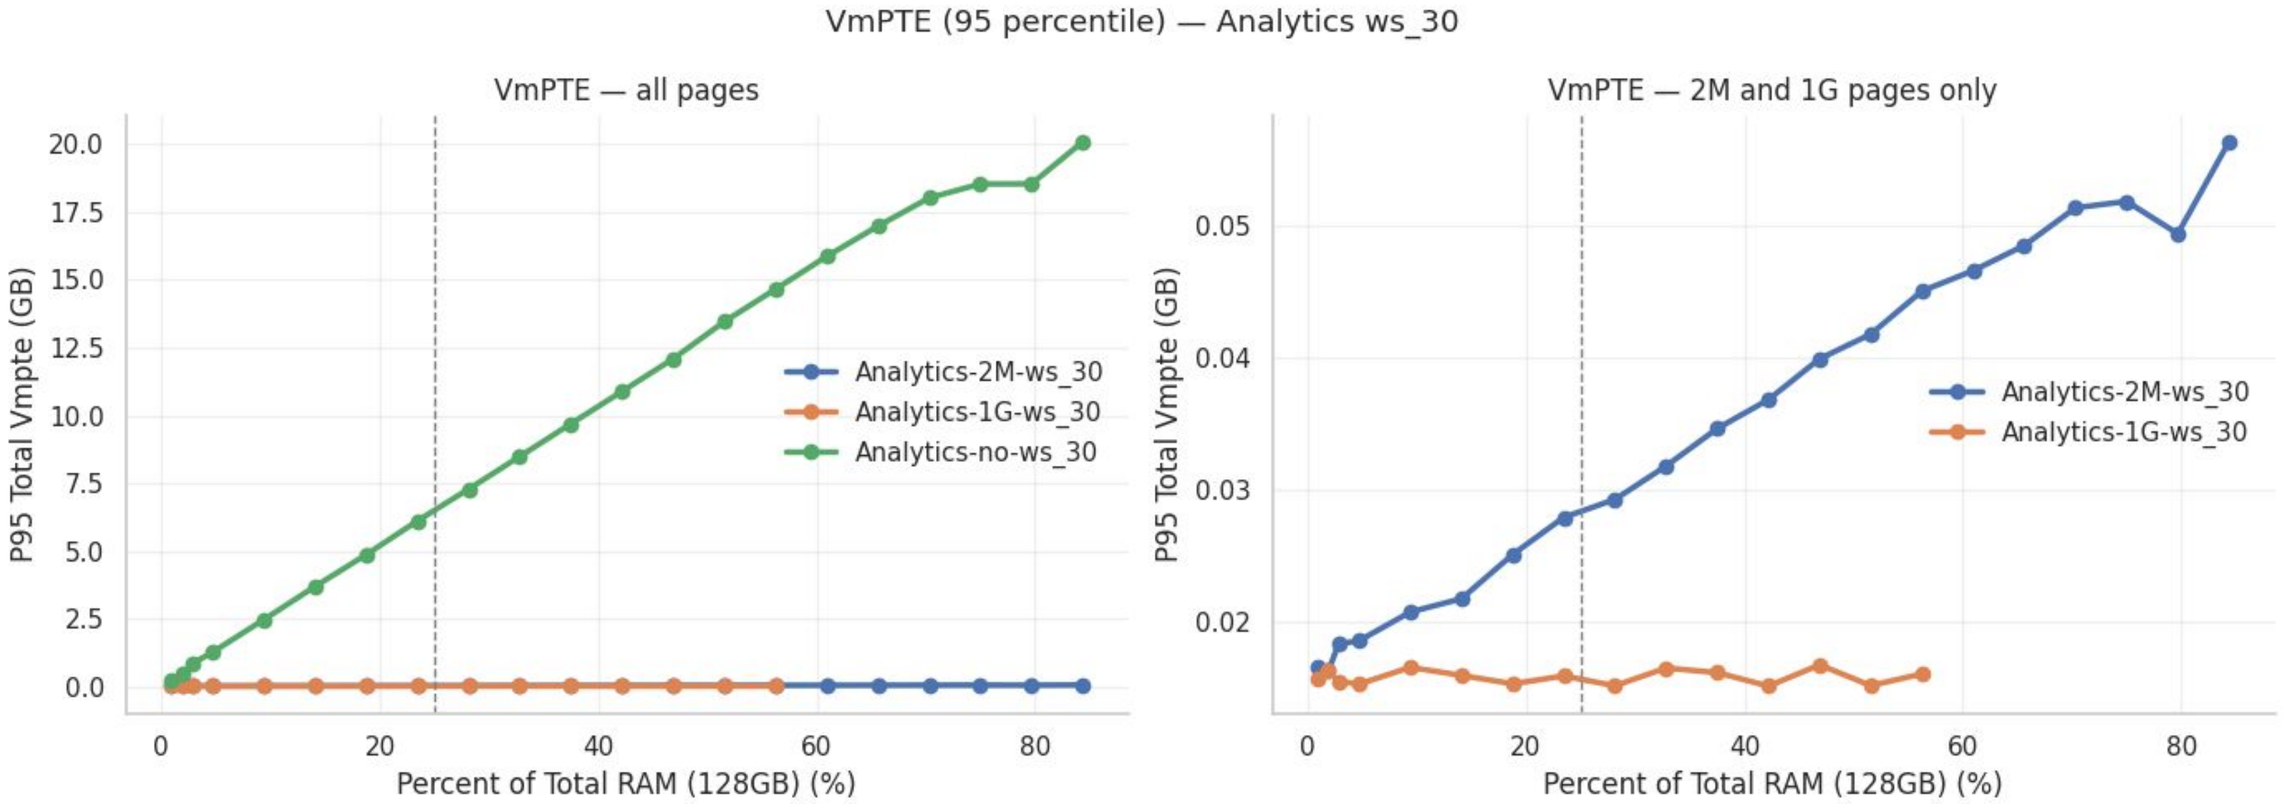

COMMAND PROMPT, INC.

Size of File Cache

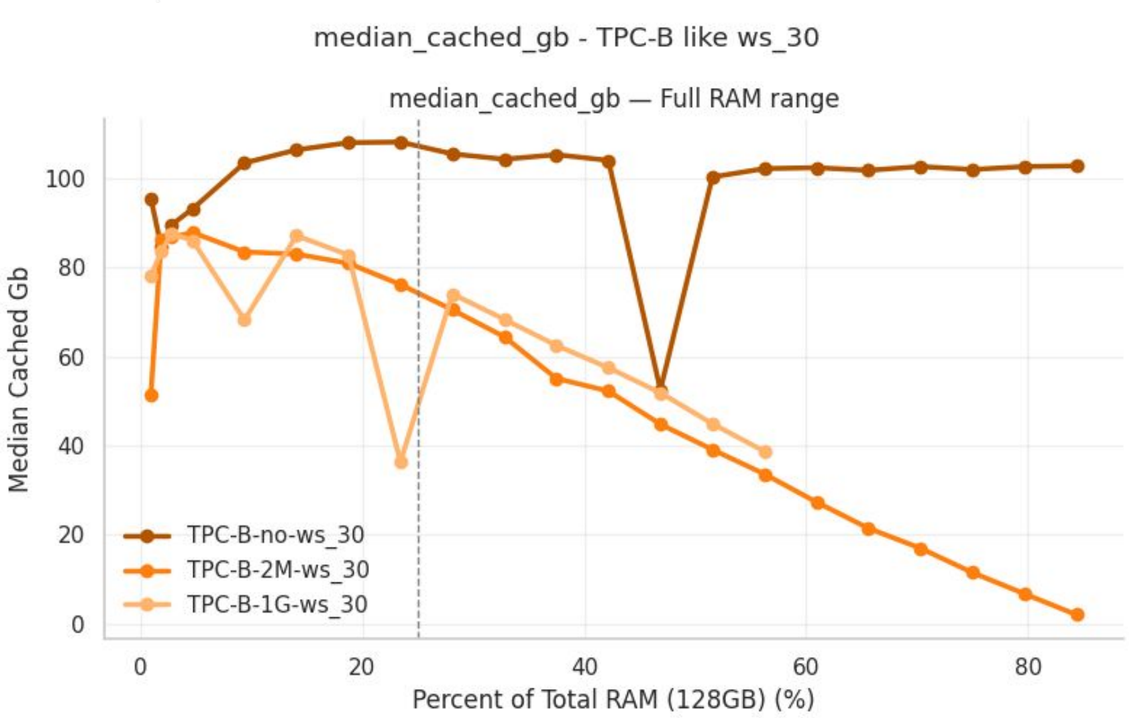

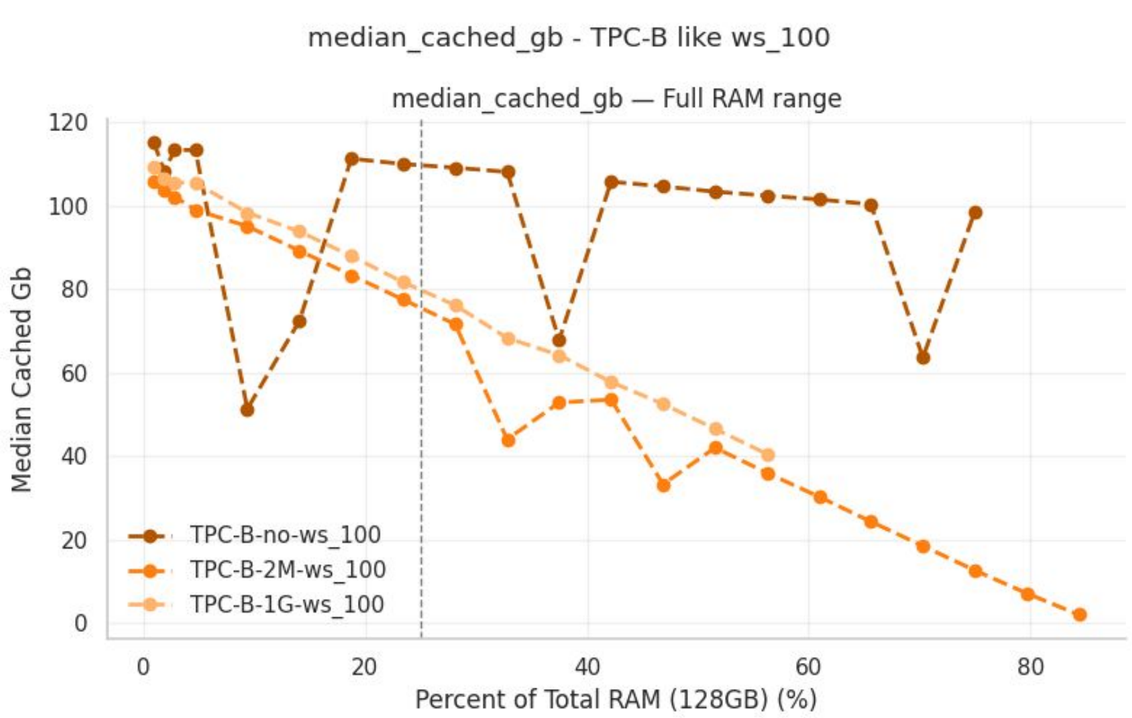

COMMAND PROMPT, INC.

Size of File Cache

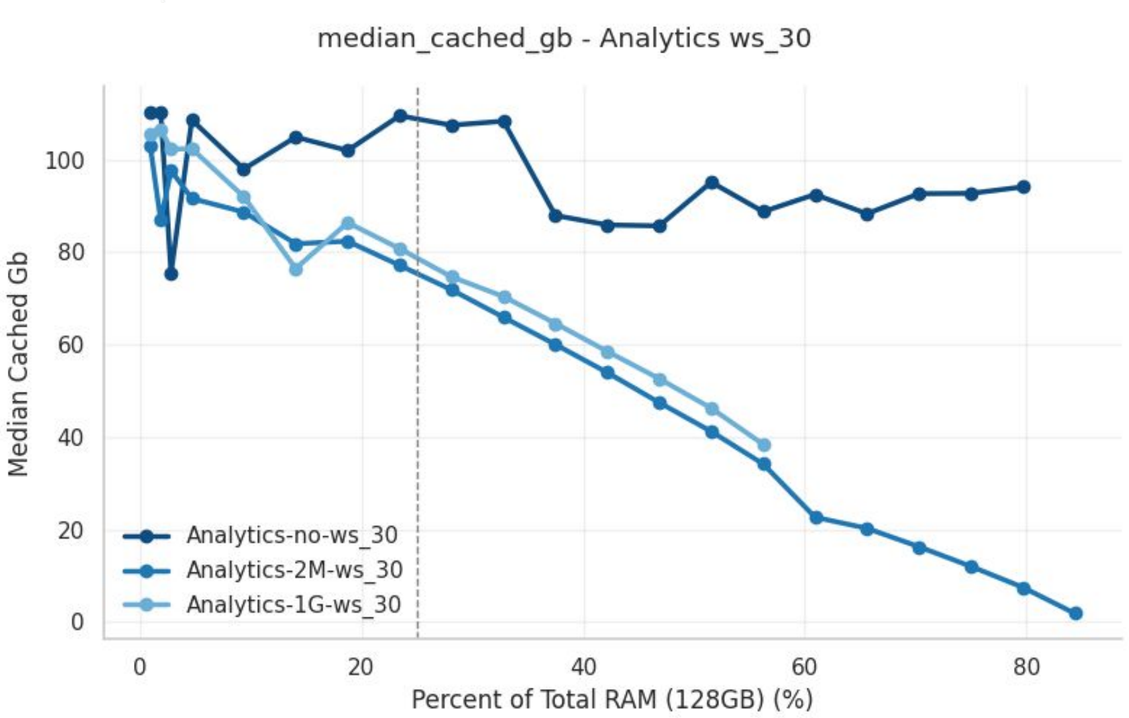

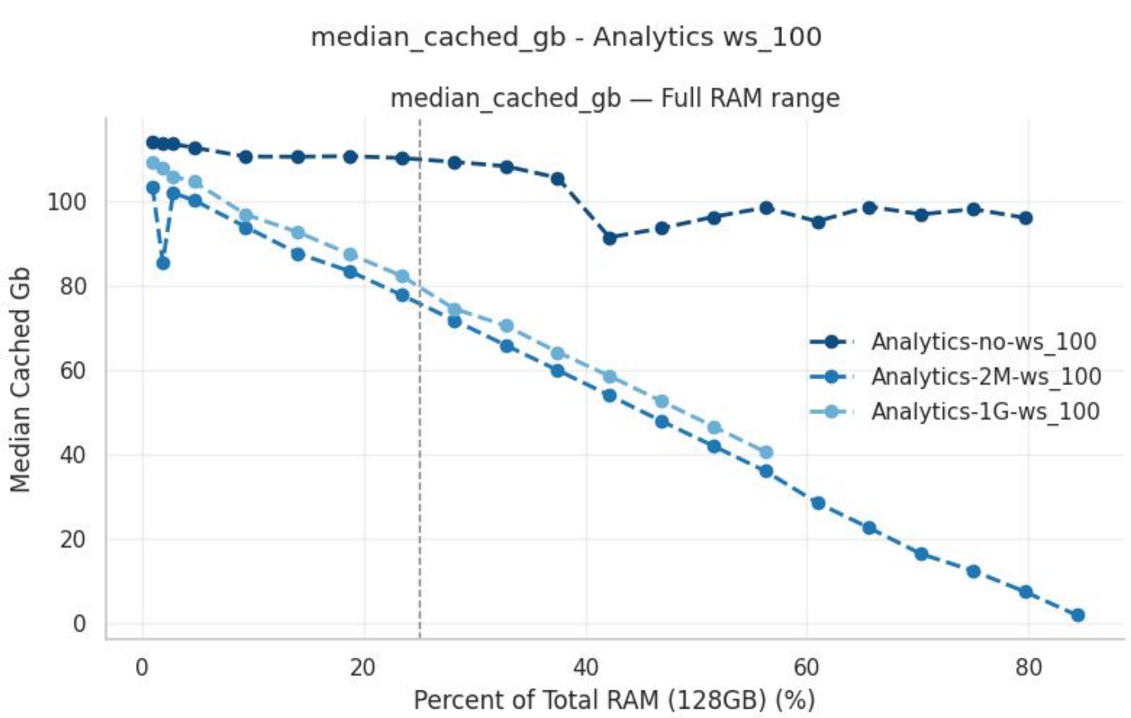

COMMAND PROMPT, INC.

Results:
Efficiency reports

COMMAND
PROMPT, INC.

COMMAND
PROMPT, INC.

Why?

- In cloud environments we pay per usage
- Moving to higher tier is expensive
- Optimize cost of providing service

COMMAND PROMPT, INC.

Disk IOPS per Million transactions

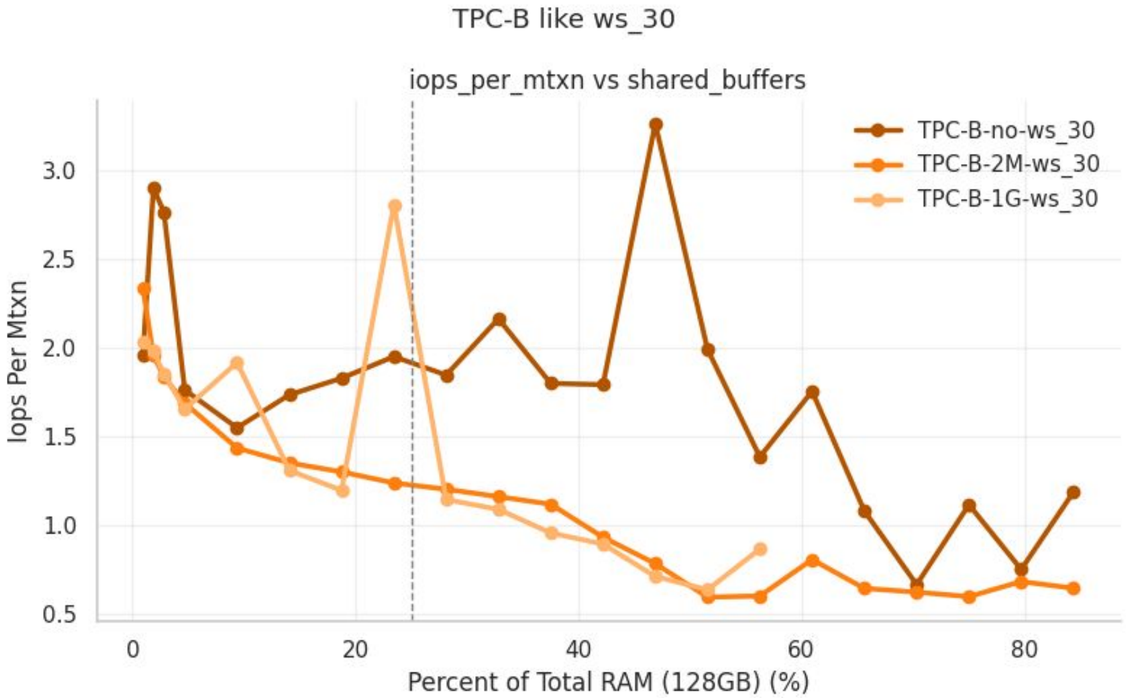

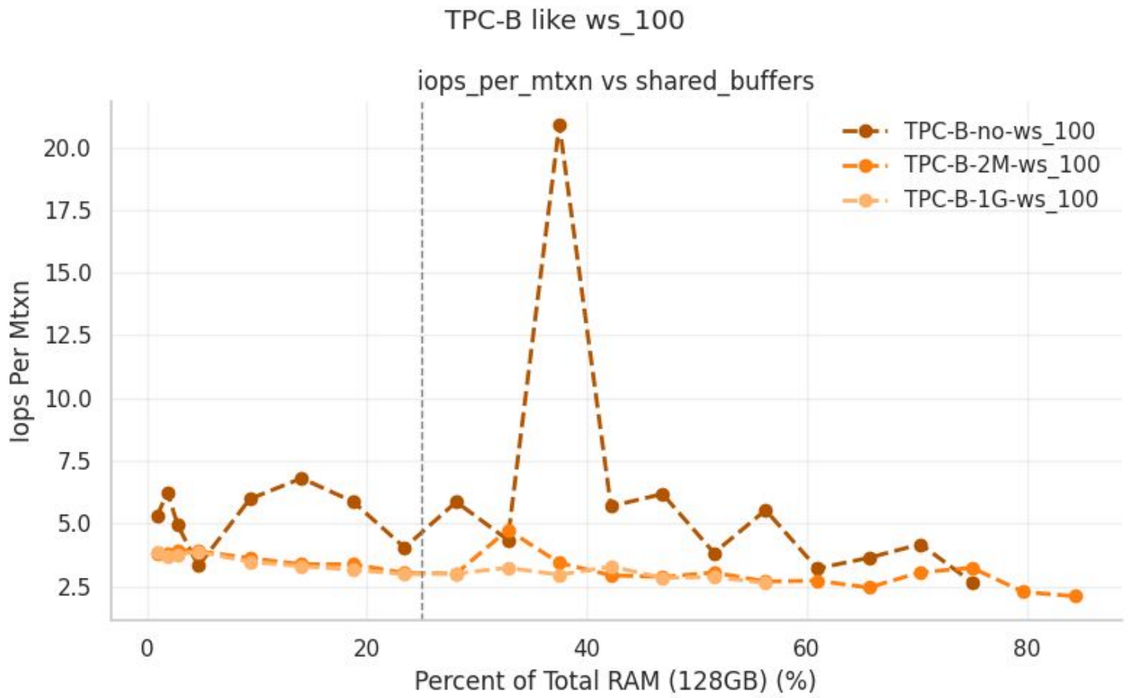

*lower is generally better

COMMAND PROMPT, INC.

Disk IOPS per Million Transactions

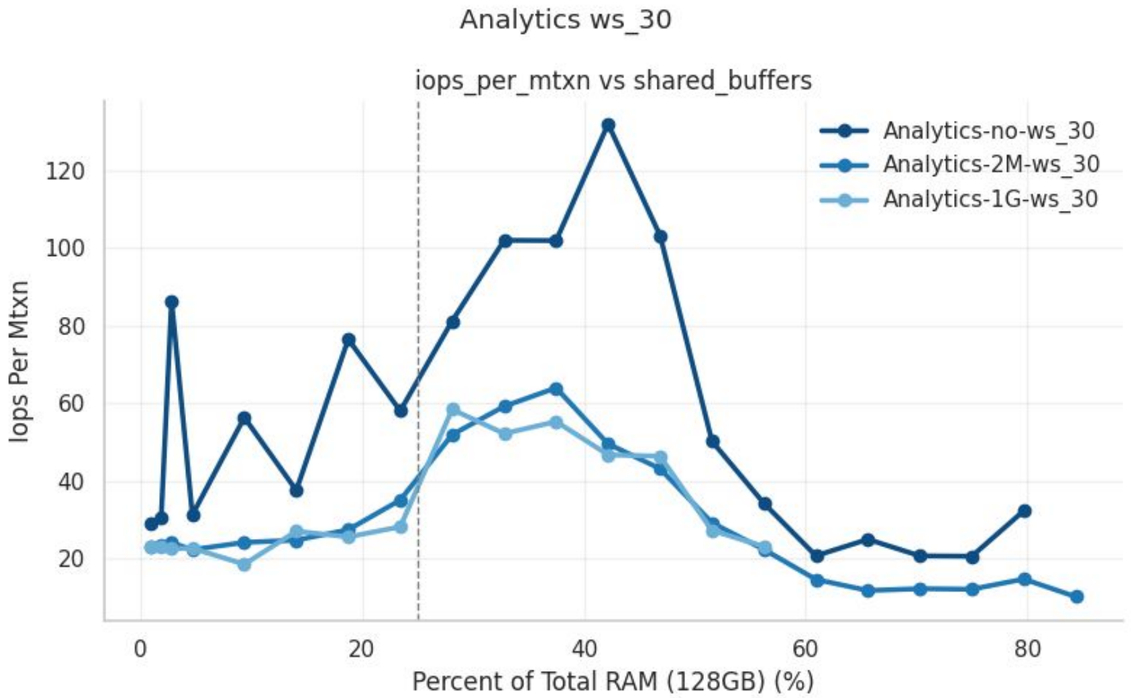

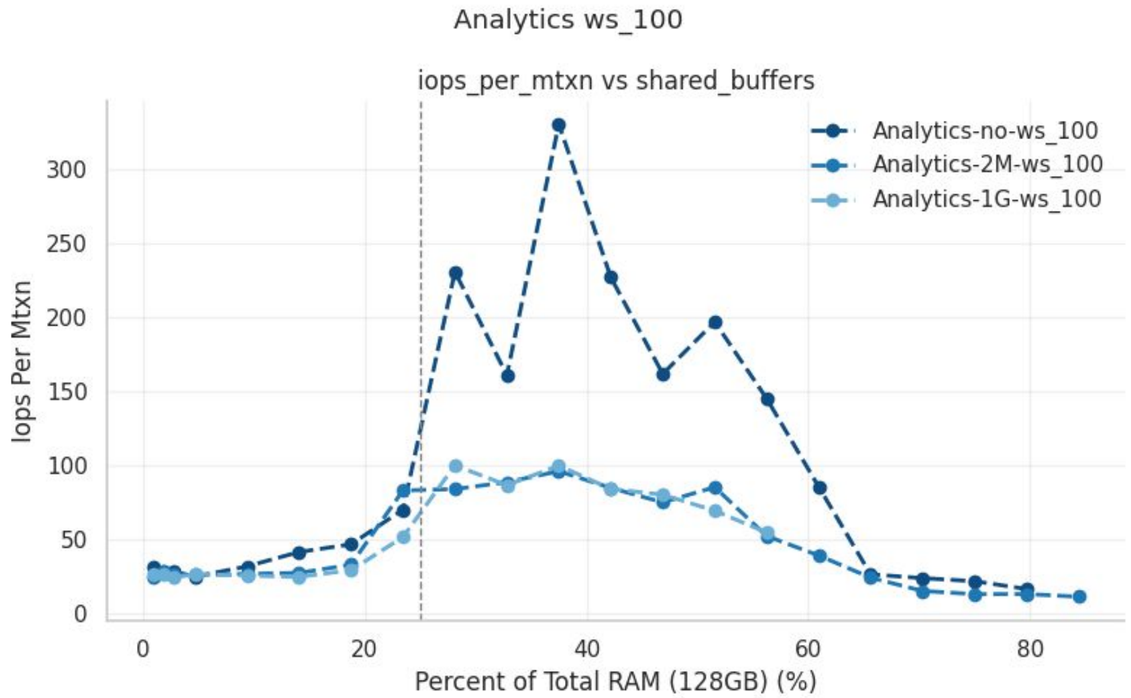

*lower is generally better

COMMAND PROMPT, INC.

CPU seconds

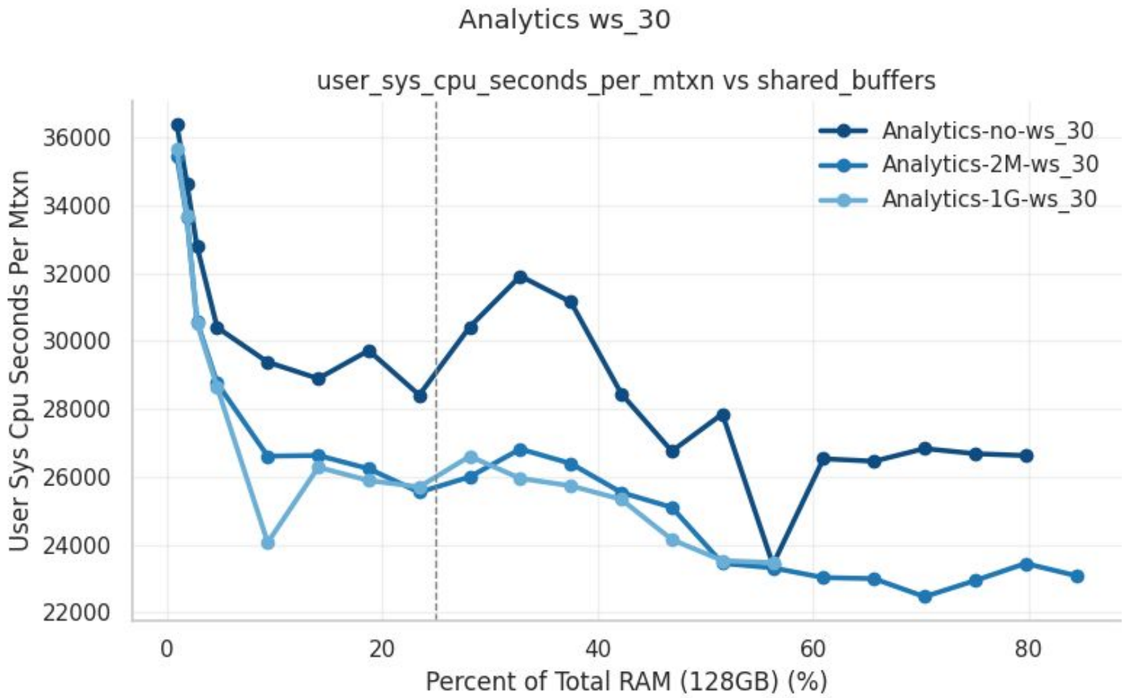

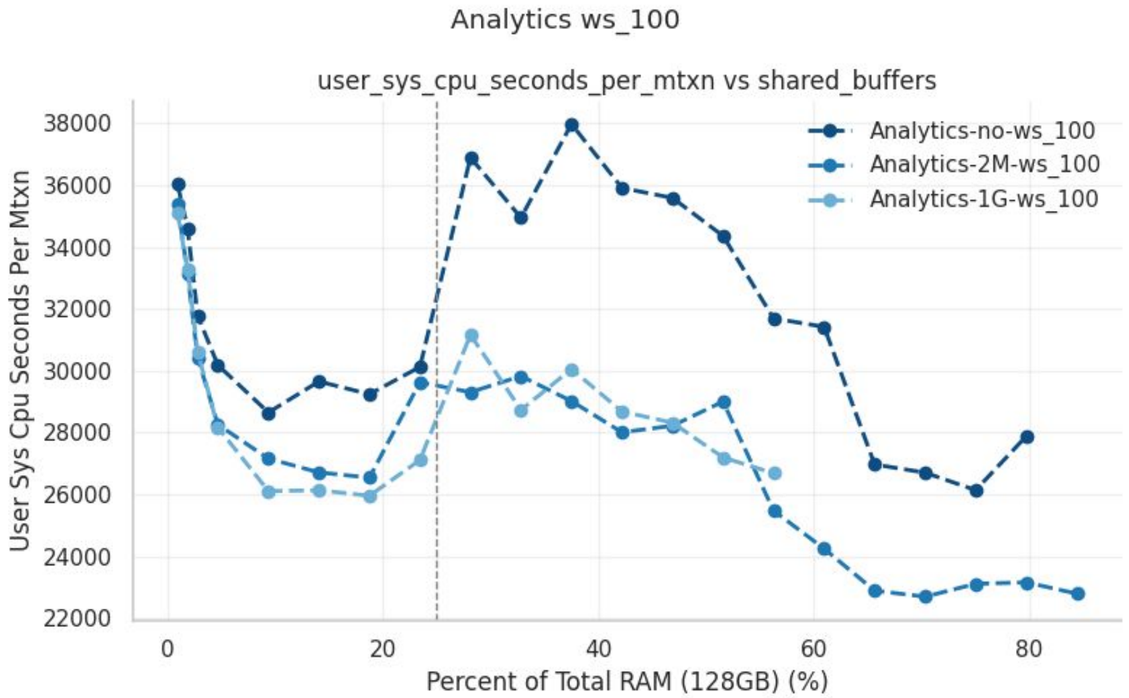

*lower is generally better

COMMAND PROMPT, INC.

Involuntary Context Switch

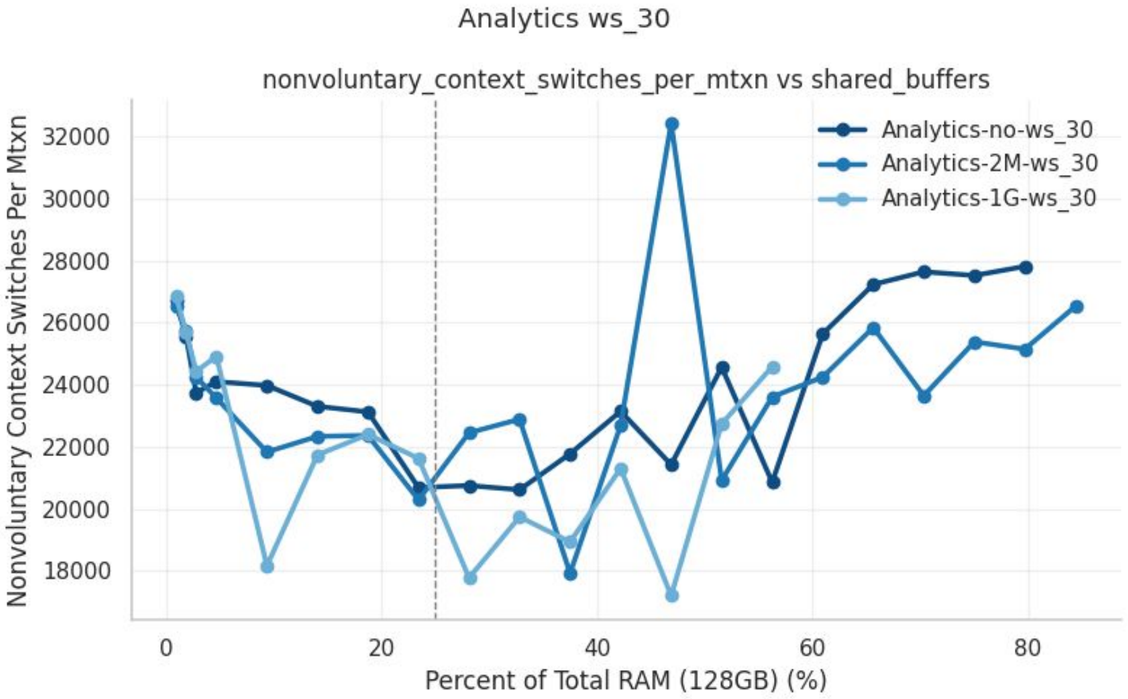

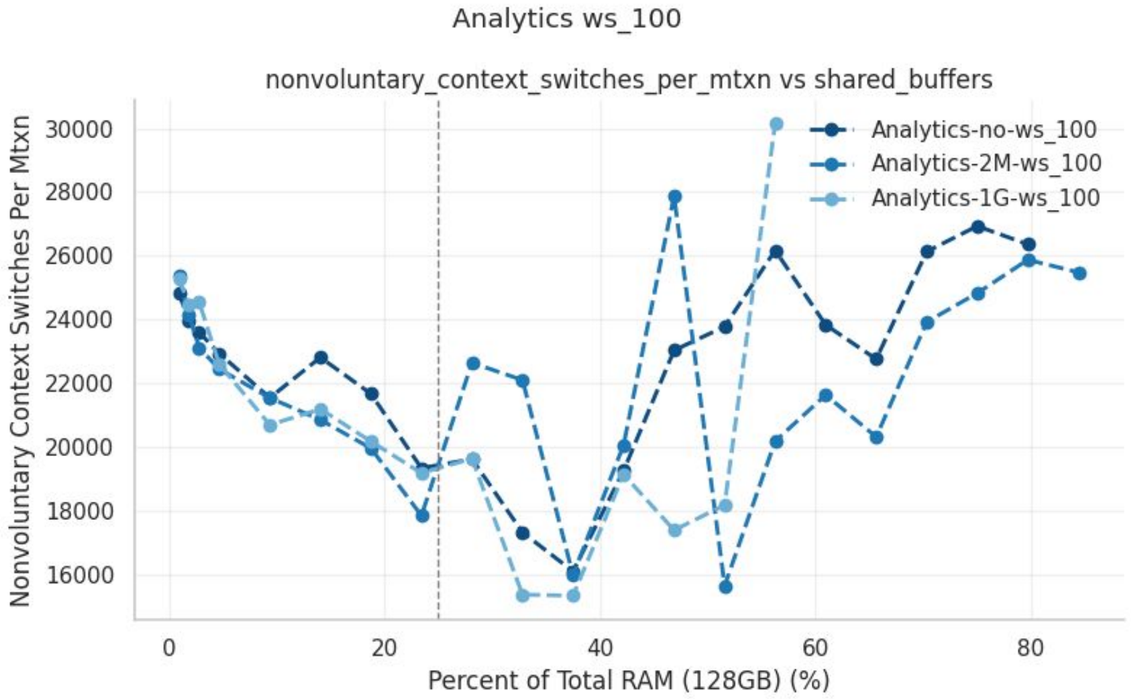

*lower is generally better

COMMAND PROMPT, INC.

Questions?

COMMAND
PROMPT, INC.

COMMAND PROMPT, INC.

+1 503 667 4564

www.commandprompt.com/contact-us

EXPERTS IN POSTGRES AND OPEN SOURCE INFRASTRUCTURE