



PostgreSQL Performance Tips & Techniques

15th November 2019

Hari Kiran
PostgreSQL Consultant, 2ndQuadrant



THP & Defrag

- Enabled by default
Disable! Disable! Disable!
- Consolidate Fragmented huge pages
- Defrags huge pages
- Application Wait



IO Scheduler

`noop` or `none` or `deadline`

What makes sense?

- When set to `noop`, host takes care of proper schedule
- I/O requests into a FIFO (First in, First Out) queue

Note: Use `pgbench` to test TPS to measure the impact of this I/O scheduler.



PG Background processes

Can't love them all... Can't leave them all

- Checkpointer
- Background Writer
- AutoVacuum



Checkpointing

- WALs & Recovery
- Not quite often; not too lazy
- Get away from defaults



Background Writer

- `maxwritten_clean` → `bgwriter_lru_maxpages`.
- Look at `buffers_clean` and `buffers_backend`
- `buffers_clean` > `buffers_backend`.
 - increase `bgwriter_lru_multiplier`
 - decrease `bgwriter_delay`

Signs of an insufficient shared buffers and hot part of your data don't fit into shared buffers and forced to travel between RAM and disks.

- `buffers_backend_fsync` should always show 0



AutoVacuum

- Don't disable autovacuum. Seriously!
- Busy databases & large ones - decrease the scale factor, more frequent clean-up (pg_stat_user_tables)
- Got good storage & multiple cores - increase the throttling parameters
- Increasing autovacuum_max_workers alone will not really help in most cases.
- Per table parameters using ALTER TABLE - makes the system more complex and more difficult to inspect.



Engineering Decisions





Connection Management - PgBouncer

Application -> PgBouncer

PgBouncer -> pool

PgBouncer -> PostgreSQL server connection

Disconnection -> PgBouncer pool



pg_dump... Dump faster

pg_dump Sections

- --section=pre-data
 - Most DDL statements
- --section=data
 - Table Data, Large Objects, Sequence Values
- --section=post-data
 - Indexes, Triggers, Rules
 - All Constraints except CHECK constraints

Tables are locked - cannot run DDL during dump

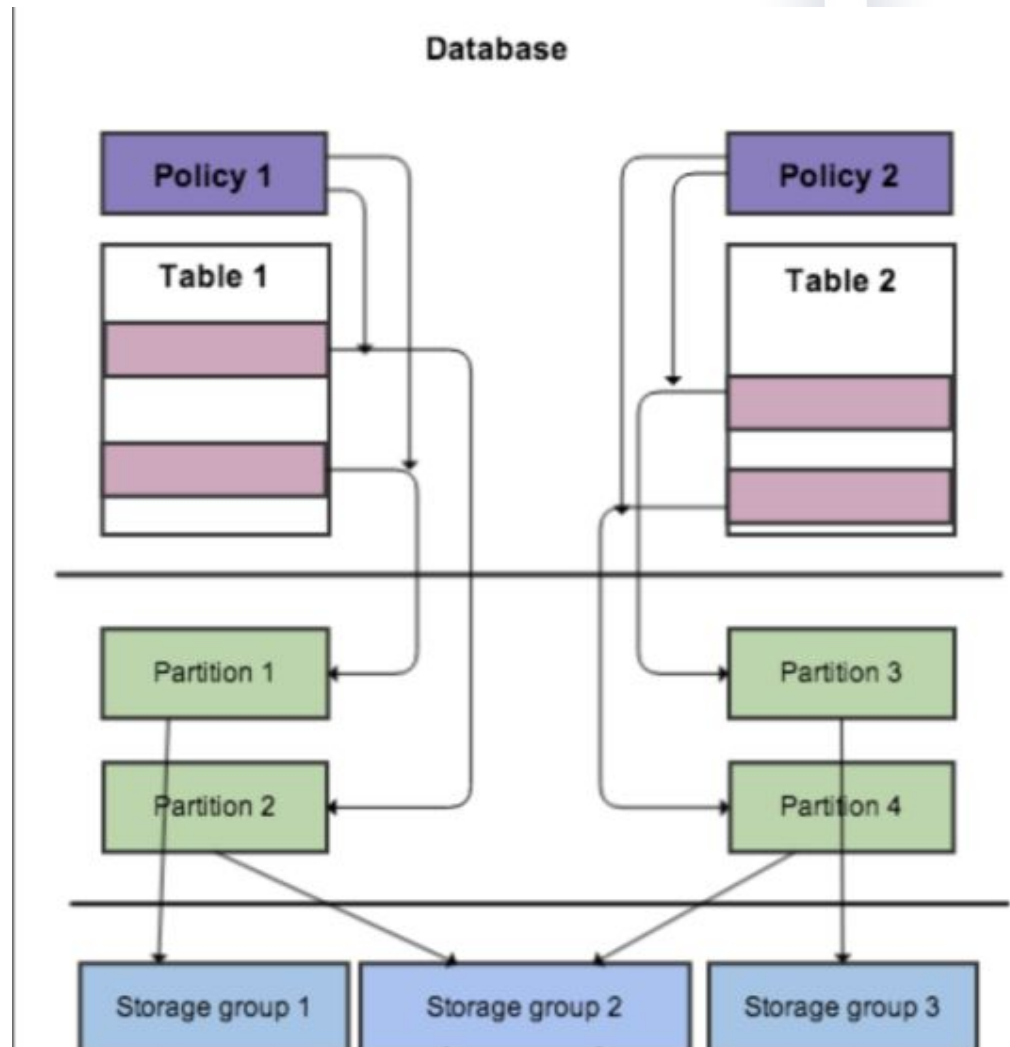
– Can specify --lock-wait-timeout=Nms

Note: Autovacuum won't start until after dump



Table Partitioning

Storage groups are
Database Tablespaces





Partition Pruning

Planner-time partition pruning

- A binary search quickly identifies matching LIST and RANGE partitions
- A hashing function finds the matching partitions for HASH partitioned tables

Execution-time partition pruning

- Phase-one: pruning is performed during executor initialization (not shown in EXPLAIN plans)
- Phase-two: remove partitions using parameters that are only known when the executor is actually running



High Availability

- `hot_standby_feedback`
- `max_standby_streaming_delay`
- `max_standby_archive_delay`



High Availability





Data Segregation

- Hard Separation via Different Databases in Same Cluster
- Soft Separation in the Same Database
 - Views
 - Row Level Security
 - Schemas



Pass by our booth!



Thank you!
Questions?

Hari Kiran
hari.kiran@2ndquadrant.com



2ndQuadrant PostgreSQL Solutions

Website <https://www.2ndquadrant.com/>
Blog <https://blog.2ndquadrant.com/>
Email info@2ndquadrant.com