Using vc-pgsql-sniffer

Baron Schwartz • VividCortex

PGConfSV • 2015

About

- VividCortex is the best way to see what your databases are doing in production
- Baron Schwartz, CEO/Founder
  - @xaprb
  - baron@vividcortex.com

Background

- VividCortex has an agent-based architecture
- Agents are self-contained and written in Go
- Sniffing/decoding TCP is a key competency
- There is a lot of very smart code in them
- Performance is critical

Our TCP Sniffer

- Our sniffing libraries are:
- Based on libpcap with modifications
- Higher performance than alternatives
- Able to handle a lot of noise without choking
- The result of several years of constant improvement

Our Gift

- We believe in and support open source
- We decided to give away our sniffer
- It's not open source, but it's free as in beer

What's Performance?

- People have funny ideas about performance and how to improve it
- See my Percona white paper, Goal Driven Performance Optimization
- Fundamentally:
  - Performance is response time (credit to Cary Millsap)
  - You can't optimize what you can't measure

Measure What Matters

- I have solved many impossible performance problems as follows:
  - Measure what matters (hint: query performance)
  - Look for the tallest tent pole & use Amdahl's Law
  - Rinse and repeat

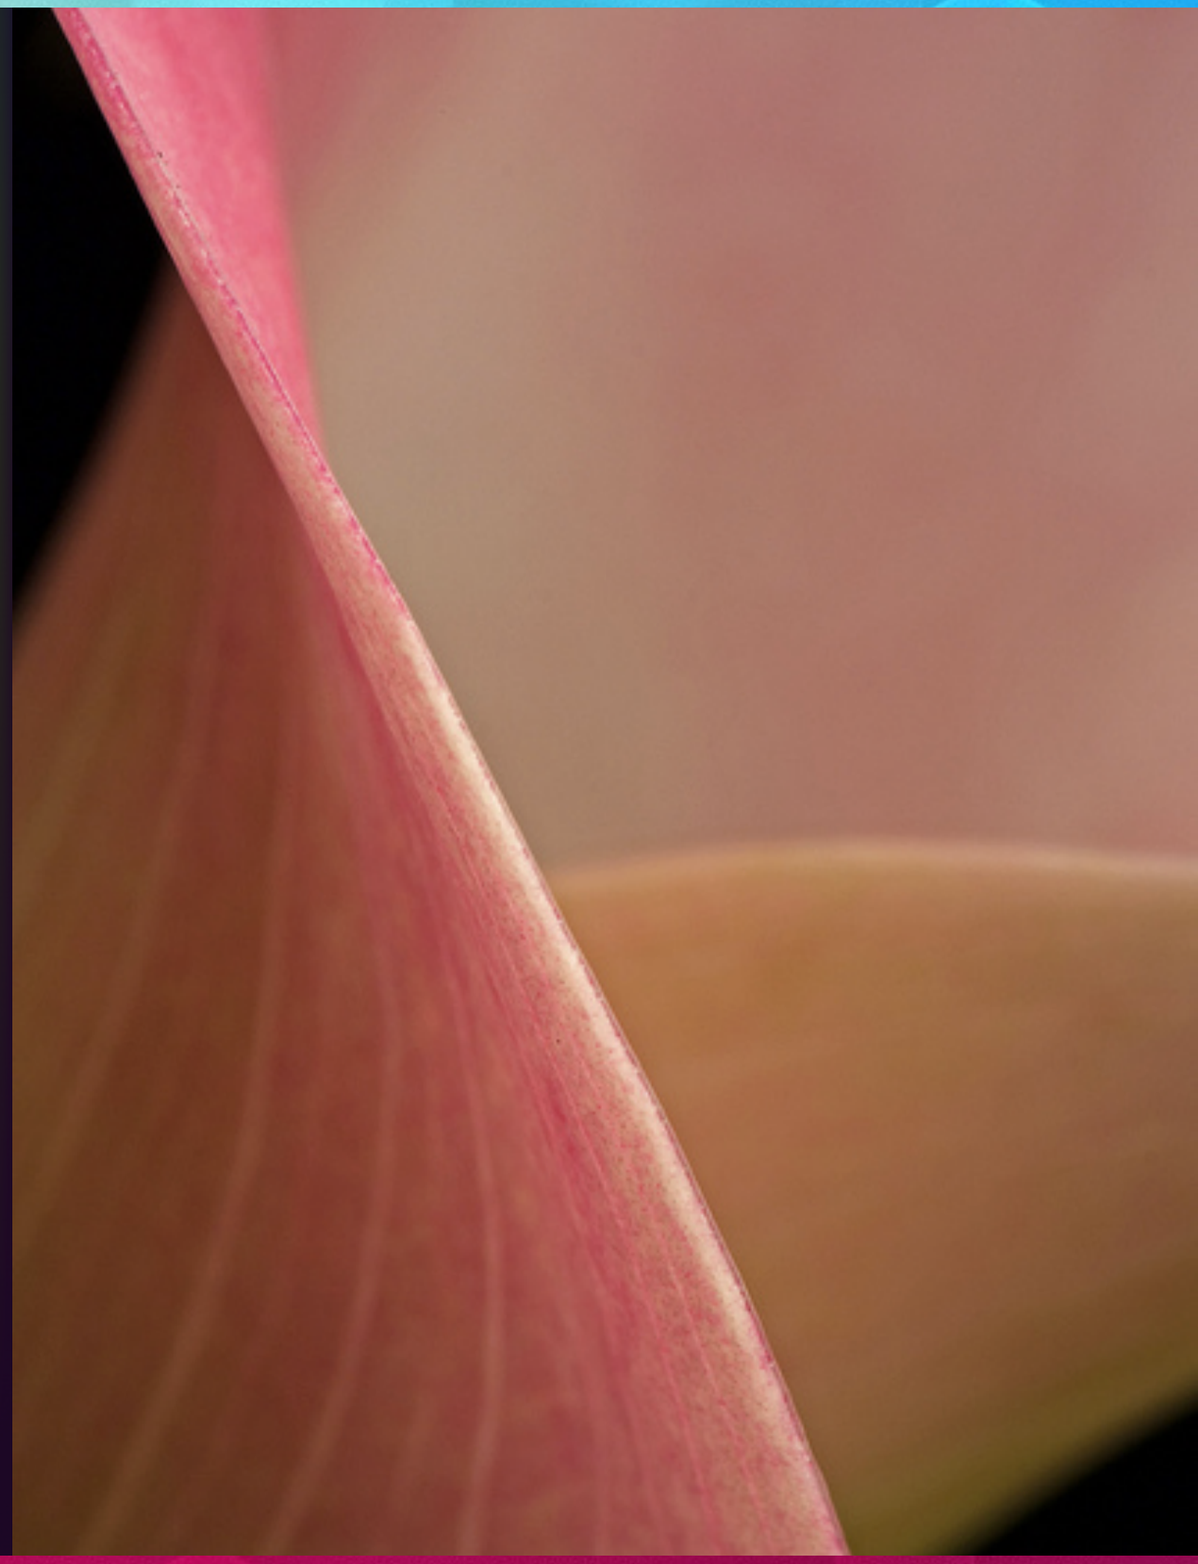

Heavy Hitters

If you focus on the heavy hitters, the rest will take care of itself.

How To Measure

- Logging?
- pg_stat_statements?
- TCP decoding?

I've Done Log Analysis

- I've built powerful tools for PG log analysis
- But logs in general have a lot of problems
- Performance overhead
- Operational risk
- Difficulty of enabling, collecting, and analyzing

PG_STAT_STATEMENTS

- It's pretty good, but has limitations
- Doesn't merge bind vars in IN()
- Has finite number of rows
- Only measure what was thought of to measure

Traffic Capture

- TCP capture+decode is a good option!
- Low overhead
- No reconfiguration
- No operational hazards
- Cons: no SSL, no Unix socket capture, packet loss, unknown prepared statements…

Free Download

- vividcortex.com/resources/

Demonstration

- Download and unzip
- Run as root/sudo
- Pipe output to a file
- Use pt-query-digest to analyze the results

Questions?

- baron@vividcortex.com
- @xaprb
- linkedin.com/in/xaprb

- … and again, that download URL: vividcortex.com/resources/

Photo Credits

- Chocolates: skrb - https://www.flickr.com/photos/skrb/5984342555
- Dew: taufuuu - https://www.flickr.com/photos/ghailon/1156522176
- Silhouette: https://www.flickr.com/photos/28481088@N00/2925783507
- Bus stop: Robert Couse-Baker - https://www.flickr.com/photos/29233640@N07/14033204315
- calla edge: mclcbooks - https://www.flickr.com/photos/39877441@N05/5455416496/
- Windmills: omarparada - https://www.flickr.com/photos/omarparada/9776594294
- Airplanes: presidioofmonterey - https://www.flickr.com/photos/presidioofmonterey/10710648865
- Droplet collision: https://www.flickr.com/photos/6929418@N07/8682467843
- 1000 layers: doug88888 - https://www.flickr.com/photos/doug88888/3139395660
- Balancing Rocks: light_seeker - https://www.flickr.com/photos/light_seeker/7780857224
- Capilano Dam: barabanov - https://www.flickr.com/photos/barabanov/4733415724
- Survival Bias: hjl - https://www.flickr.com/photos/hjl/15942299782