

<https://www.flickr.com/photos/hoyvinmayvin/4906678960>

Theo Schlossnagle, Founder @Circonus, @postwait on Twitter

Operational Software

Ignorance
Buys
Pain




Tenet #1

Designing for the unknown should never burden today

Compromise when forced.

Design with tomorrow's expected volume,
but today's functional requirements.



Tenet # 2

Never be left curious as to
what your software *is* doing

Observability is the only fast-path to success.



Tenet #3


Understand what your software *was* doing

This is one of the trickiest tenets.

You can't log every instruction performed, every packet sent & received, every context switch, every I/O.

Log things that are infrequent,
measure things that are frequent.

What's frequent? I said this was tricky.



Tenet #4

Understand internal failures

The only thing worse than a malfunction in production, is one without sufficient actionable data.

Core dump, stack trace analysis, etc.

These can wake people up.




Tenet #5

Operator remediation of an external failure is a bug

A system failing, a disk failing, a network partition, etc.

Once the external failure is remediated, no human being should be involved in returning the (internal) system to correct operation.



Tenet #6 & #7

Tight coupling
reduces resiliency

Loose coupling
reduces debug-ability



Tenet # 8

Avoid difficult problems if possible

One does not “just add” replication, consensus and fault tolerance into systems.

Never solve a harder problem than you are presented.

The best solution to a problem is to remove the problem.





<https://www.flickr.com/photos/skynoir/8300610952>

A tail of two systems

Queueing

Storage



From RabbitMQ to Fq

RabbitMQ: Oh, the outages we've had.
violates #1, #2, #3, #5, #8

Fq:
(#1) build only semantics we need
(#2 / #3) DTrace probes
(#4) cores dumps and backtrace.io
(#5) handled by simplicity of design and use
(#6 / #7) push decoupling out, gain debug-ability
(#8) punt clustering downstream to clients



<https://www.flickr.com/photos/fotologic/2165675515>





Overview

Queues

Routes

Routes

noit.firehose

	Prefix	Program	Invocations	Approximate Latency
transient	check.	[no program / match all]	226.85M	256ns
transient	*	[no program / match all]	11.97M	260ns

ns level deep routing



From Postgres to Snowth

Postgres 8... hacked in a column store using arrays

Tenet #3: build pg_statsd

Tenet #5: ... build Snowth

Snowth: ring topology time-series store

(#1 / #8) only commutative operations,

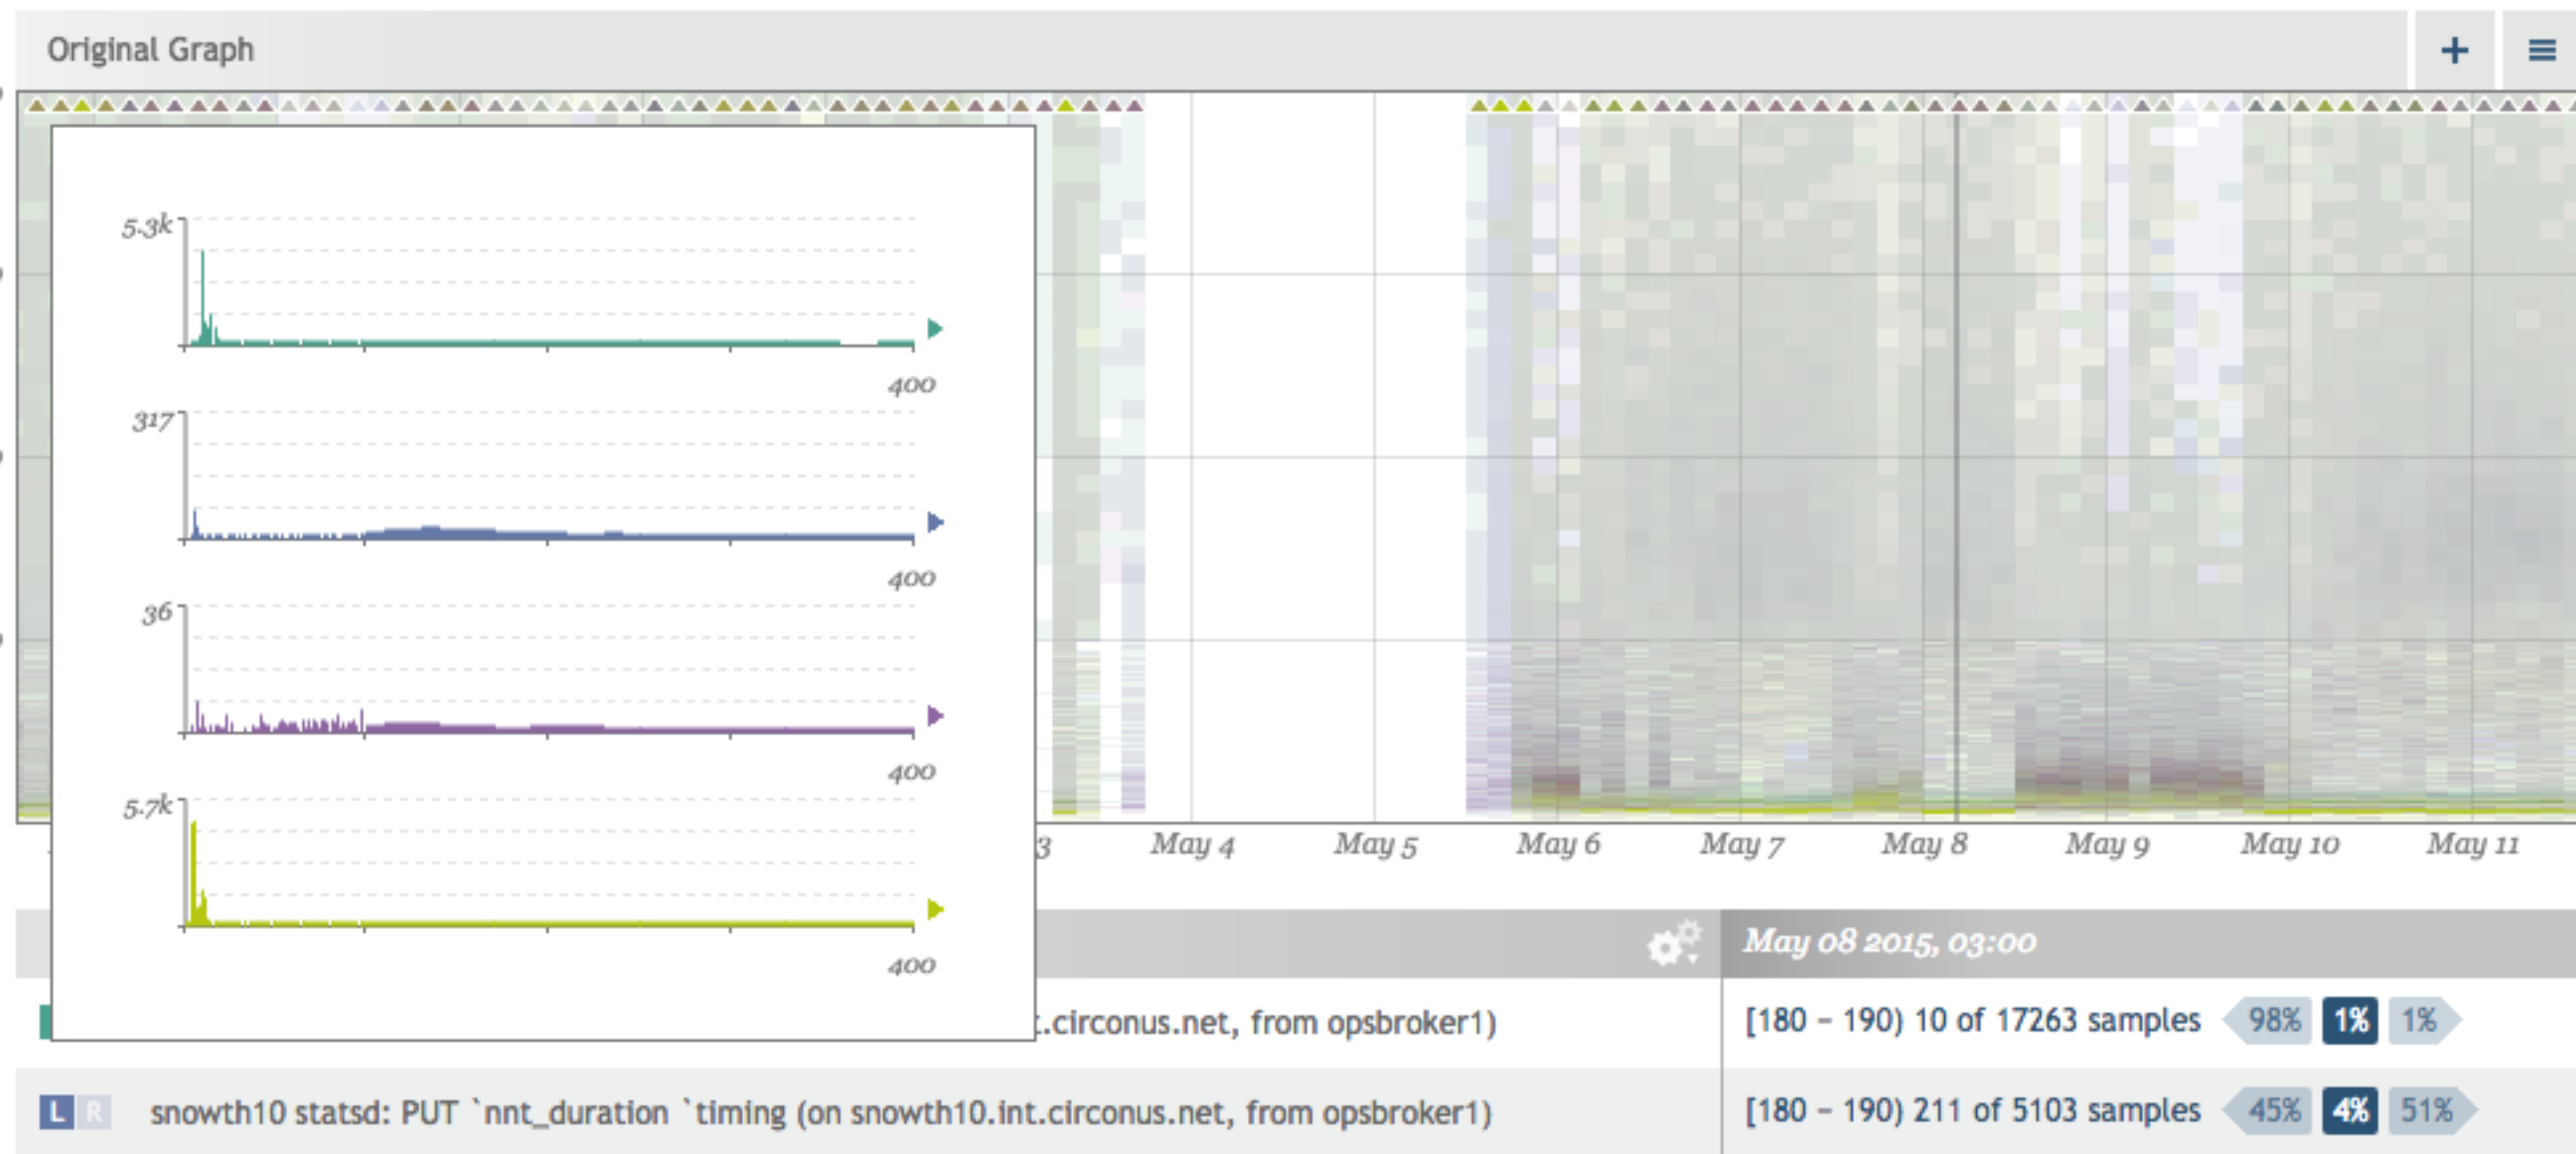

(#2) DTrace probes

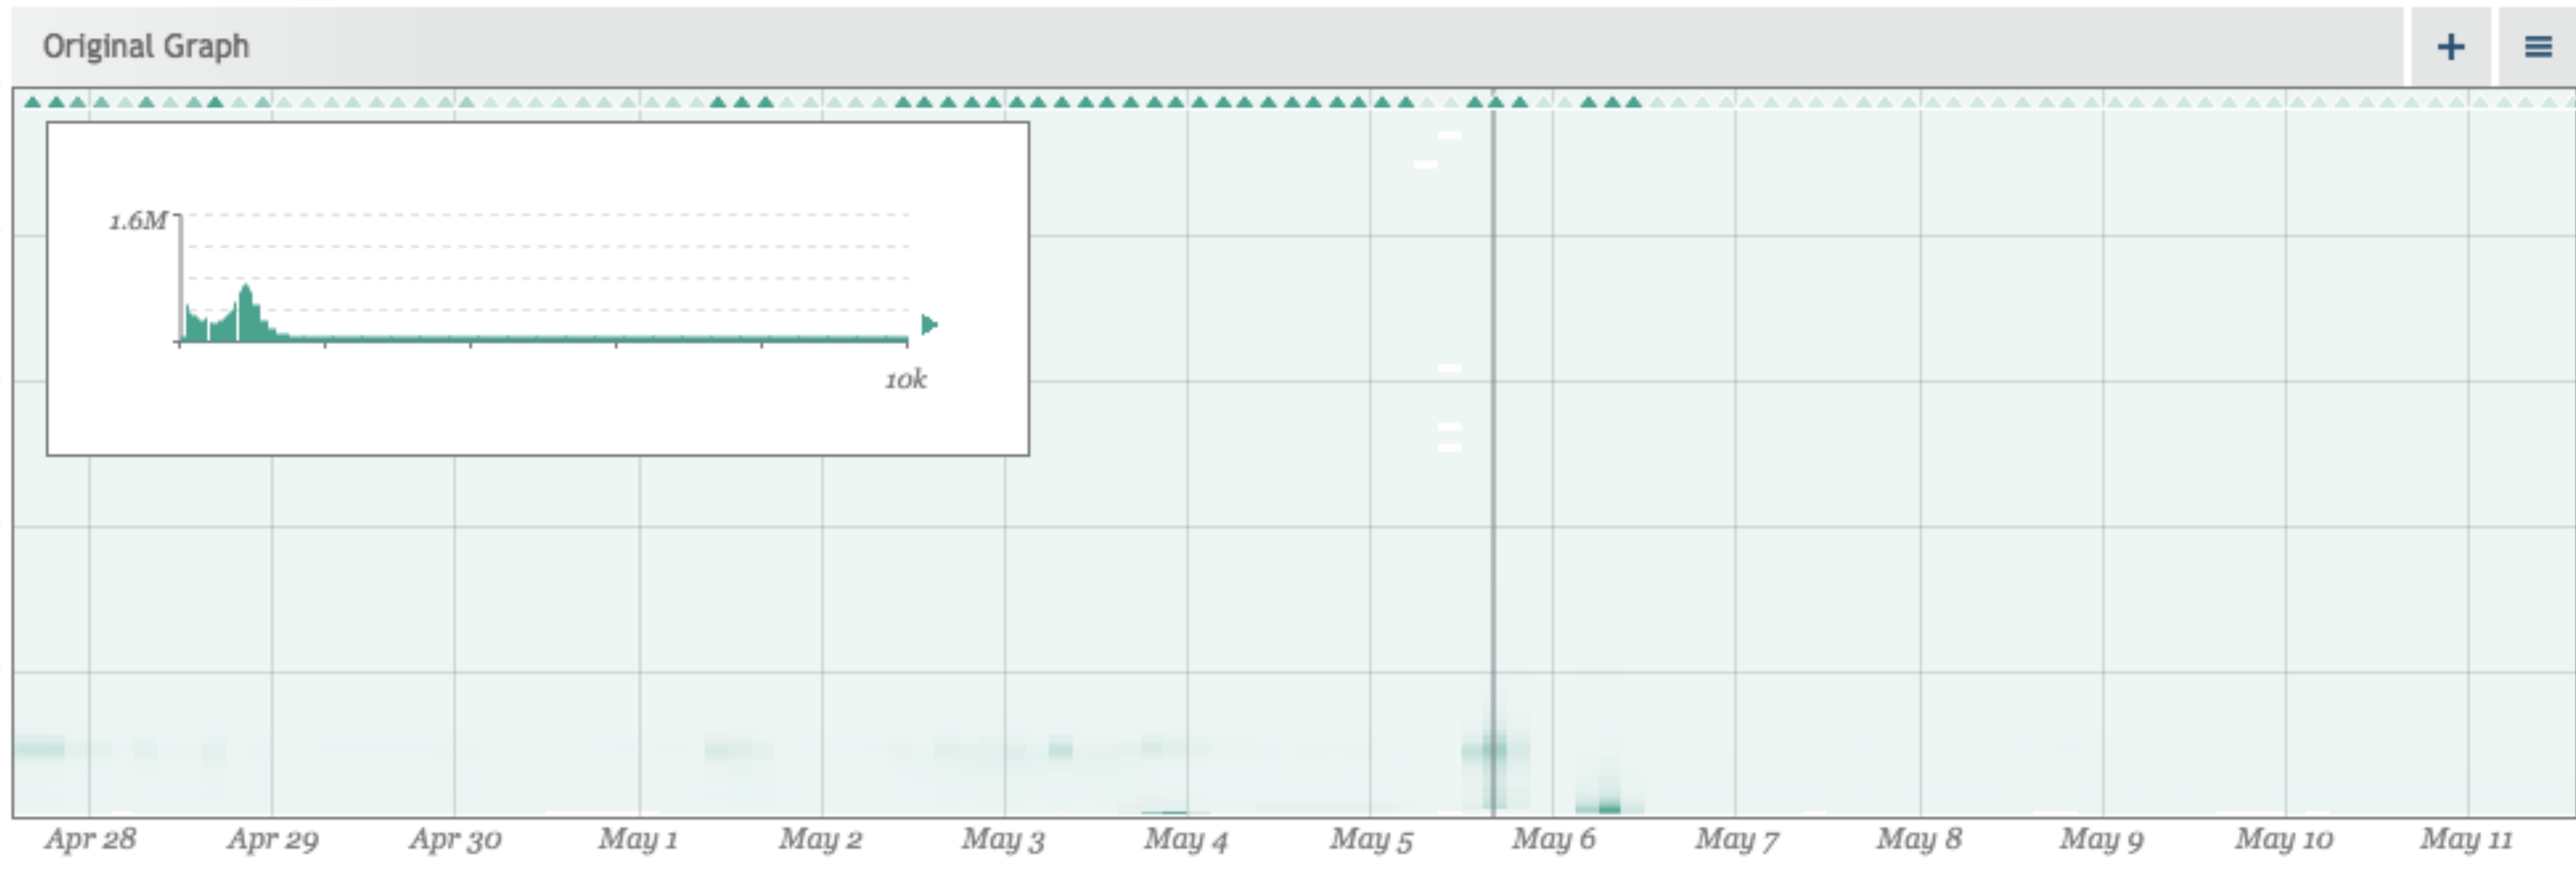

(#3) Circonus itself (statsd, histograms, etc.)

(#4) backtrace.io

(#7) implemented Zipkin (Dapper-like) tracing








⚙️ *May 05 2015, 15:00*

L R IO latency (µs) [0.94k - 0.95k] 716988 of 63274672 samples 53% **1%** 46%



quartile bands





“mvalue” best-guess modes

