CIRCONUS

PERFORMANCE MONITORING – REVOLUTIONIZED

What's in a Number?

monitoring and measuring;

the spiritual and the carnal.

Theo Schlossnagle @postwait

- Entrepreneur: OmniTI, Message Systems, Circonus
- Author: Scalable Internet Architectures, Web Operations
- Speaker: over 75 speaking engagements around the world
- Engineer: Senior ACM Member, IEEE Member, ACM Queue
- Open Source:

  Apache Traffic Server, Varnish, Node.js, node-iptrie, node-amqp, Solaris HTTP accept filers, OpenSSH + SecurID, JLog, Mungo, Portable umem, Reconnoiter, mod_log_spread/spreadlogd, Spread, Wackamole, Zetaback, concurrencykit, Net::RabbitMQ, MCrypt, node-gmp, node-compress, and others.

Monitoring ~ Voyeurism

- Watching from the outside in.
- Observing systems and their behaviors.
- Ultimately, a single observation distills into two things:
  - a set of circumstances
  - a value

Circumstances: Dimensions

- What are you looking at?
  - outbound octets on
  - an Ethernet port (1/37) on a switch
  - in rack 12, in datacenter 6
  - for client ACME
- What time is it?

Circumstances: Dimensions

- What are you looking at?
  - length of time until DOM ready for
  - a user loading a web page (/store/foo/item/bar)
  - using Chrome 10
  - from Verizon FioS in Fulton, Maryland.
- What time is it?

Value

- Ultimately, there are two types of values we care about:
  - text
    - an version number, http code, ssh fingerprint
  - numeric
    - 0.000000000093123, 18446744064986053382, 1.1283e97, 6.33e-23, 0, -1, 1000, 42

Text Data

- Incredibly useful,
  because the *“when”* can be correlated
  with other events and conditions.

- Idea:
  - if the distinct set of text values is sufficiently small,
  - consider the text value as a dimension, and
  - the new value = 1

Text Metrics in Use

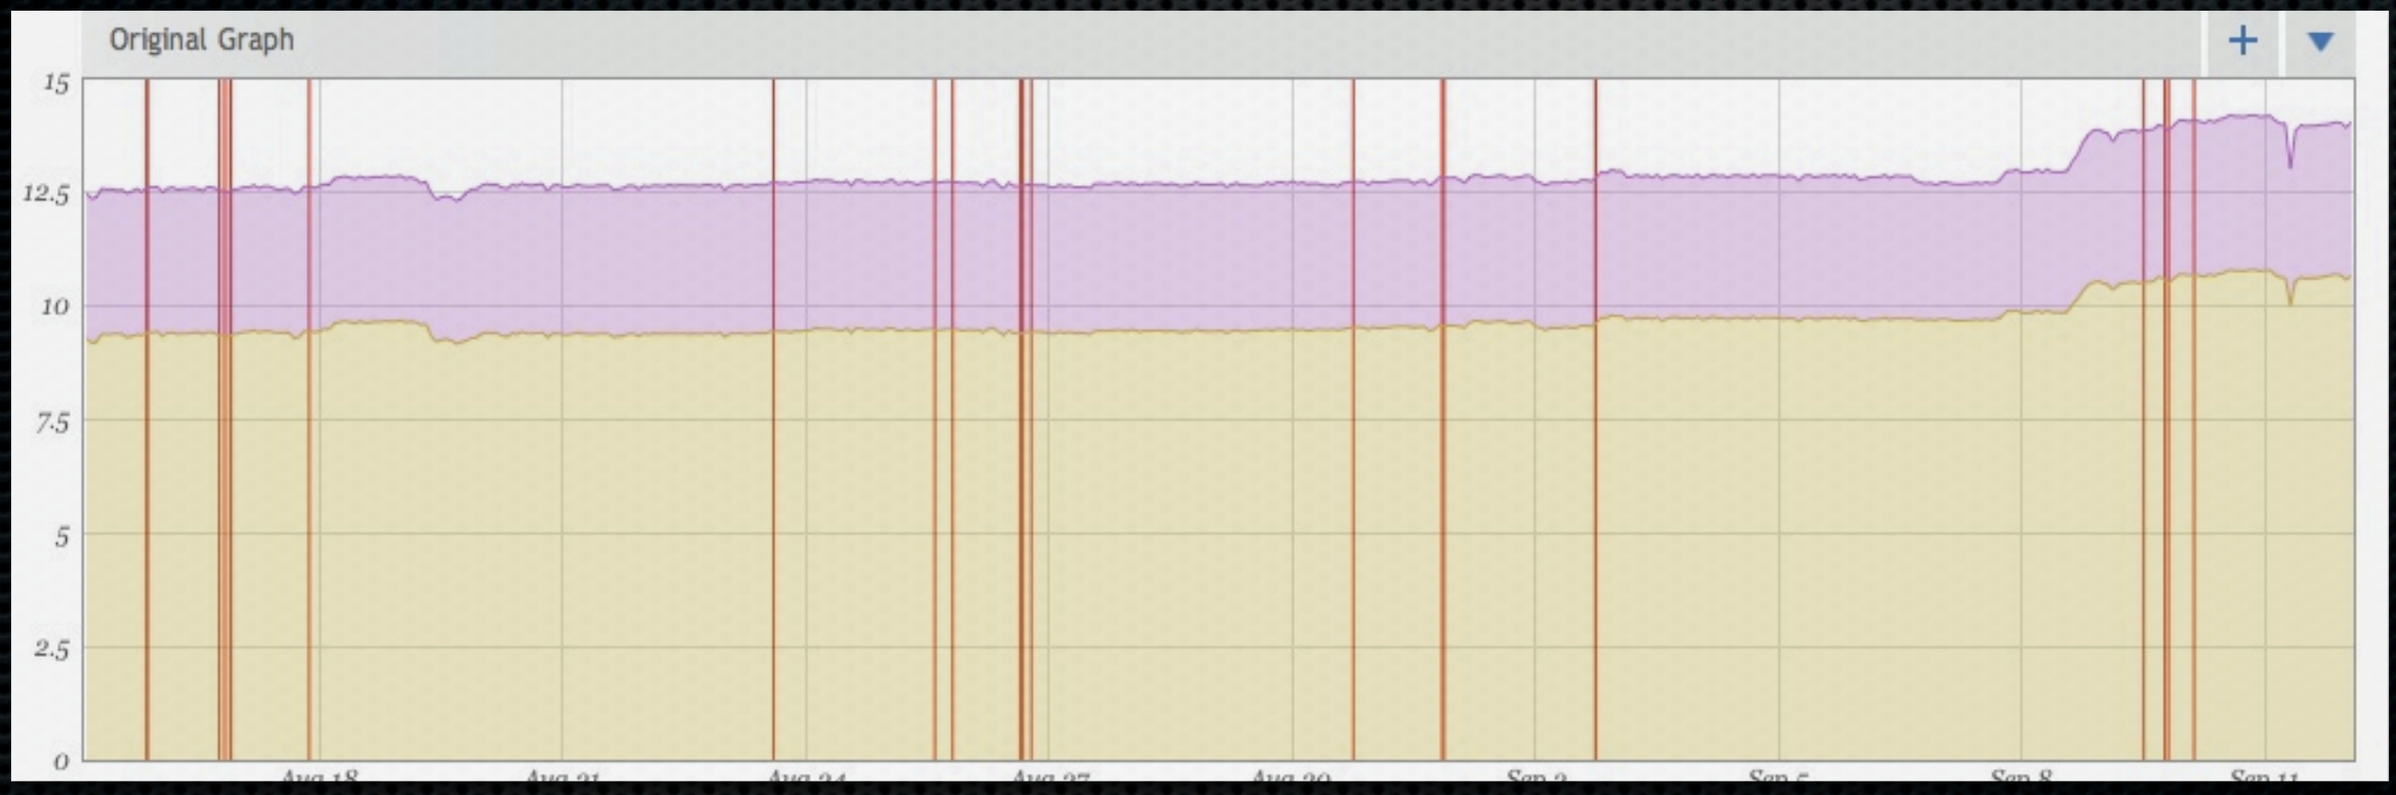

- Code deployments vs. process activity

Numeric Data

- There tends to be a lot of it.
- Example web hits:
  - one box, one metric, 1 billion / month.
  - that's only 400 events/second.
  - each event has about 20 dimensions.

Numeric Metrics in Use

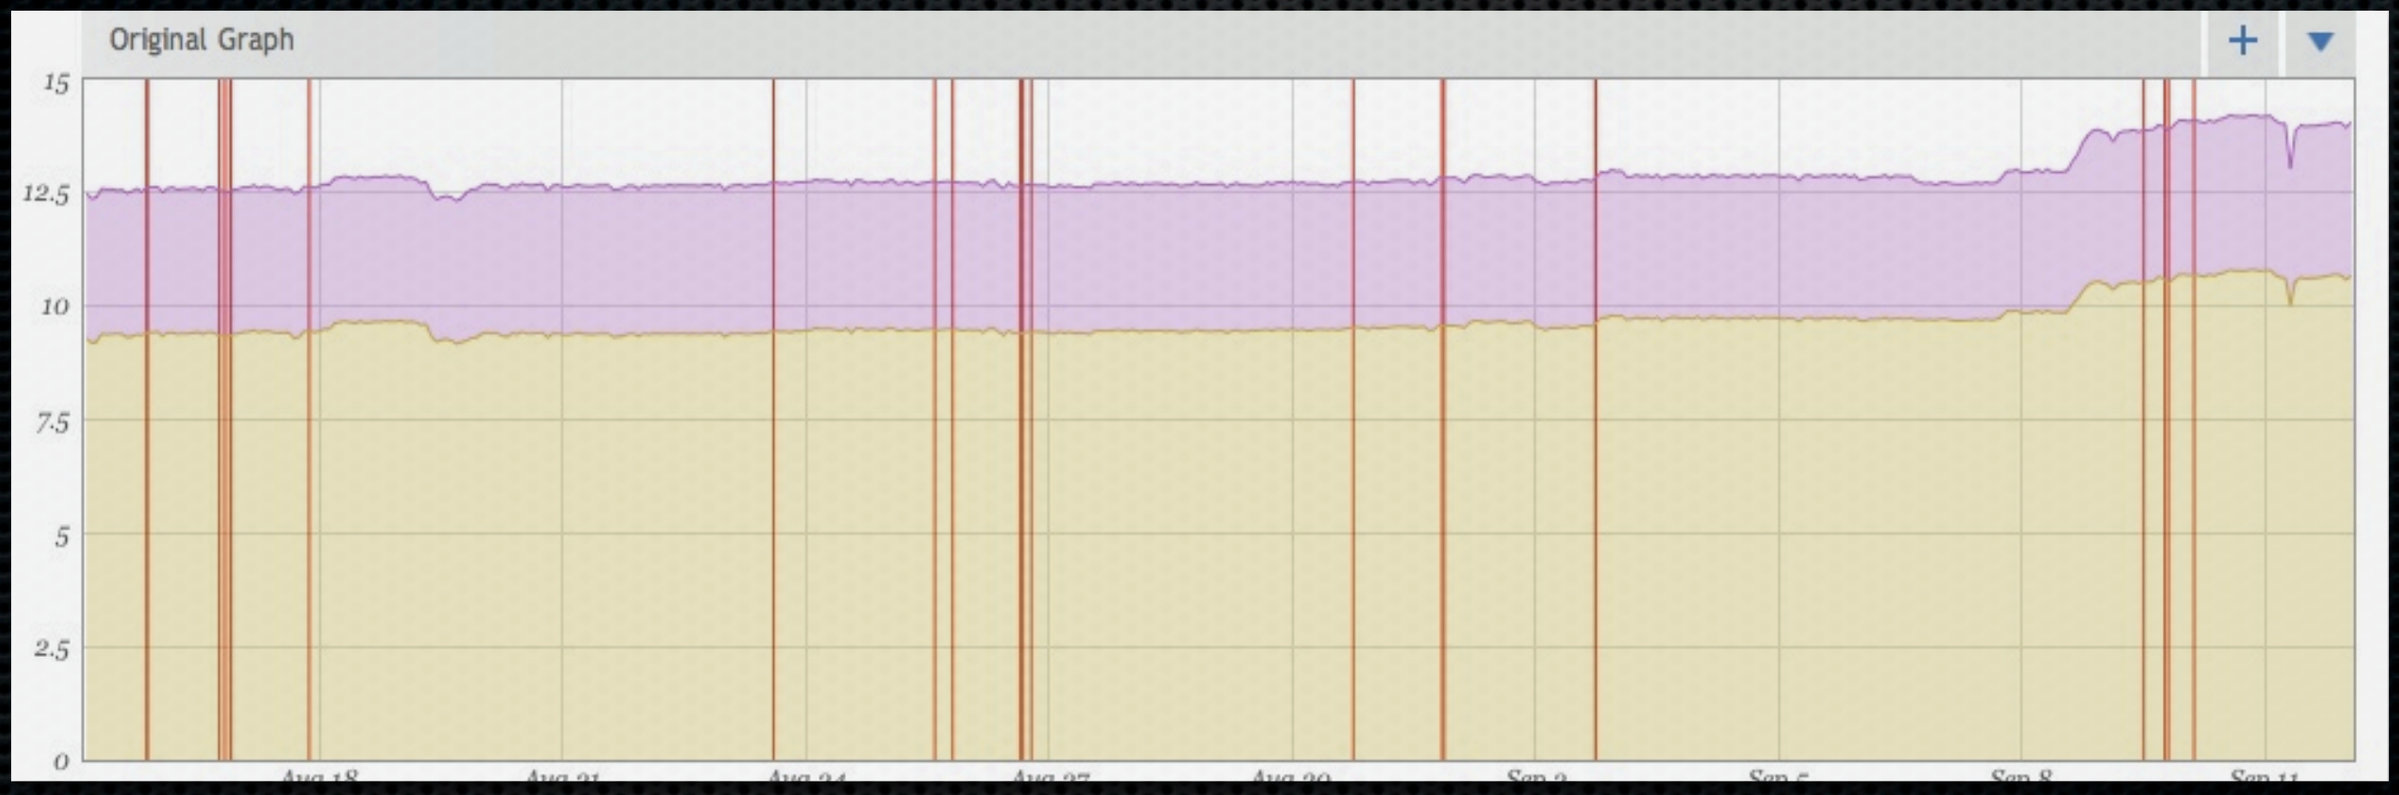

672px ≅ 4 weeks (2419200s) or 3600s ≅ 1px

Numeric Metrics in Use

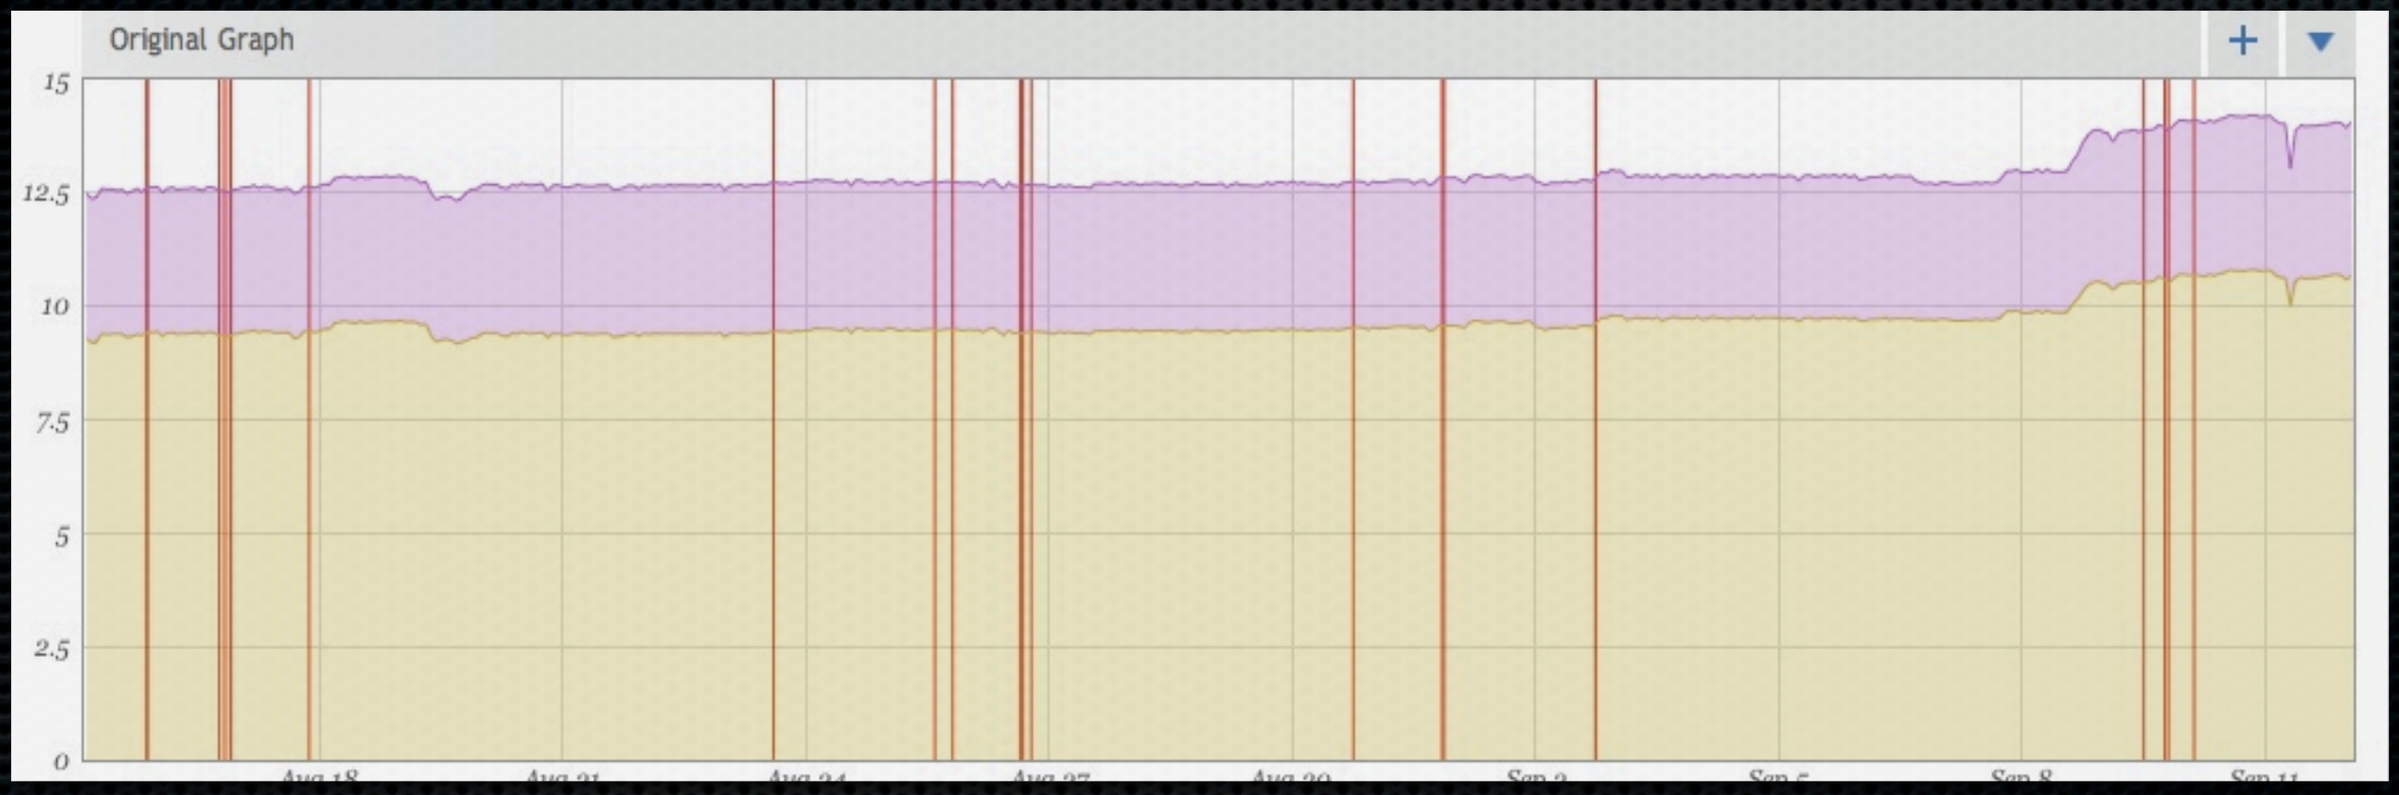

672px ≅ 4 weeks (2419200s) or 3600s ≅ 1px

- in previous example: 24000 samples per ‘point’

Numeric Metrics in Use

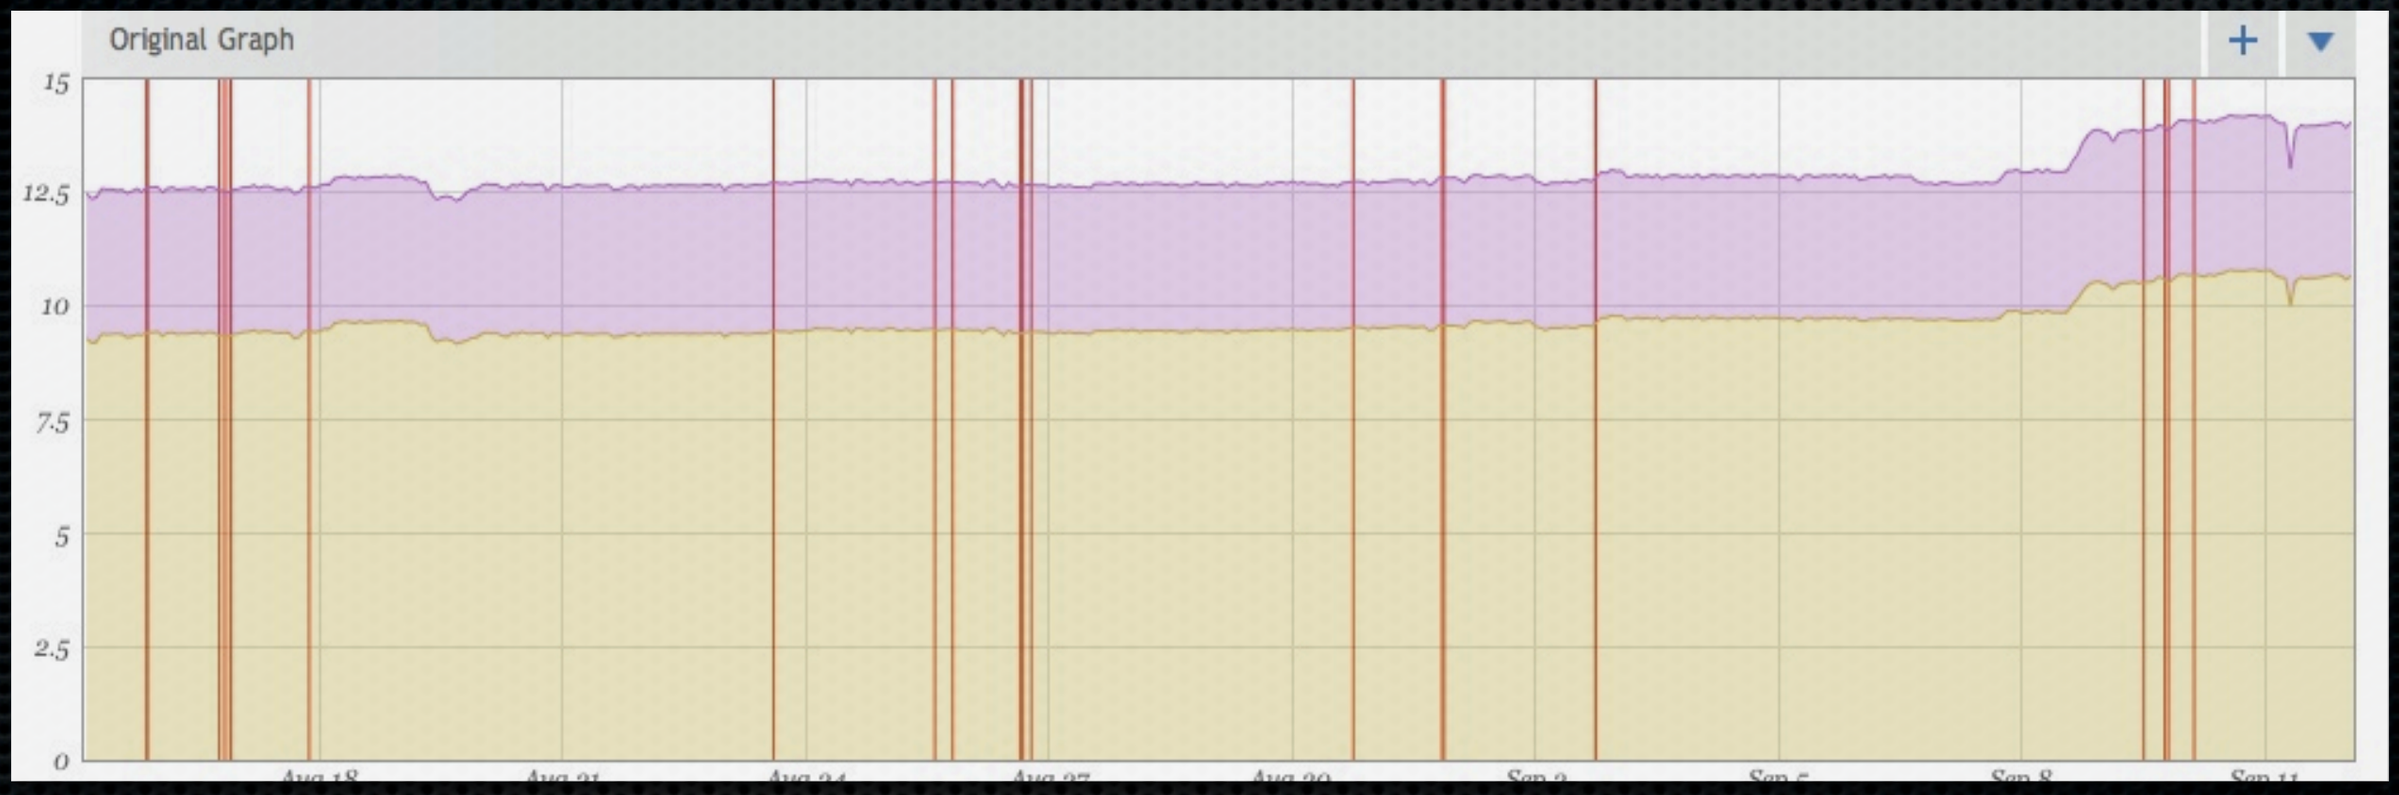

672px ≅ 4 weeks (2419200s) or 3600s ≅ 1px

- in previous example: 24000 samples per ‘point’
- even at one minute samples: 60 samples per ‘point’

An Event

each is a singularity:

individual;

unrepeatable;

unique.

mondolithic.com

Passive vs. Active

- Active is a synthesized event.
  - You perform some action an measure it.
- Passive is the collection of other events.
  - Stuff that really happened.
- Ask yourself which is better.

Conceding the Truth

- Passive events are ideal in high traffic environments.
- Active events are necessary in low traffic environments.

- “high and low” are relative, so...
  - you need both
  - always

Numbers: how they work.

- Think of a car's speedometer.
  - it shows you how fast you are going (sorta)
  - at the moment you read it (almost)
- the only reason this is acceptable is:
  - you have a rudimentary understanding of physics,
  - you are in the car, and
  - your body is exceptionally good at detecting acceleration.

Computers are different.

- computers are fast: billions of instructions per second.
- you aren't inside the computer.
- most "users" somehow forget rudimentary physics.
- interfaces and visual representations mislead you.
- a speedometer (or familiar readout) leads you to a false confidence in system stability between samples.

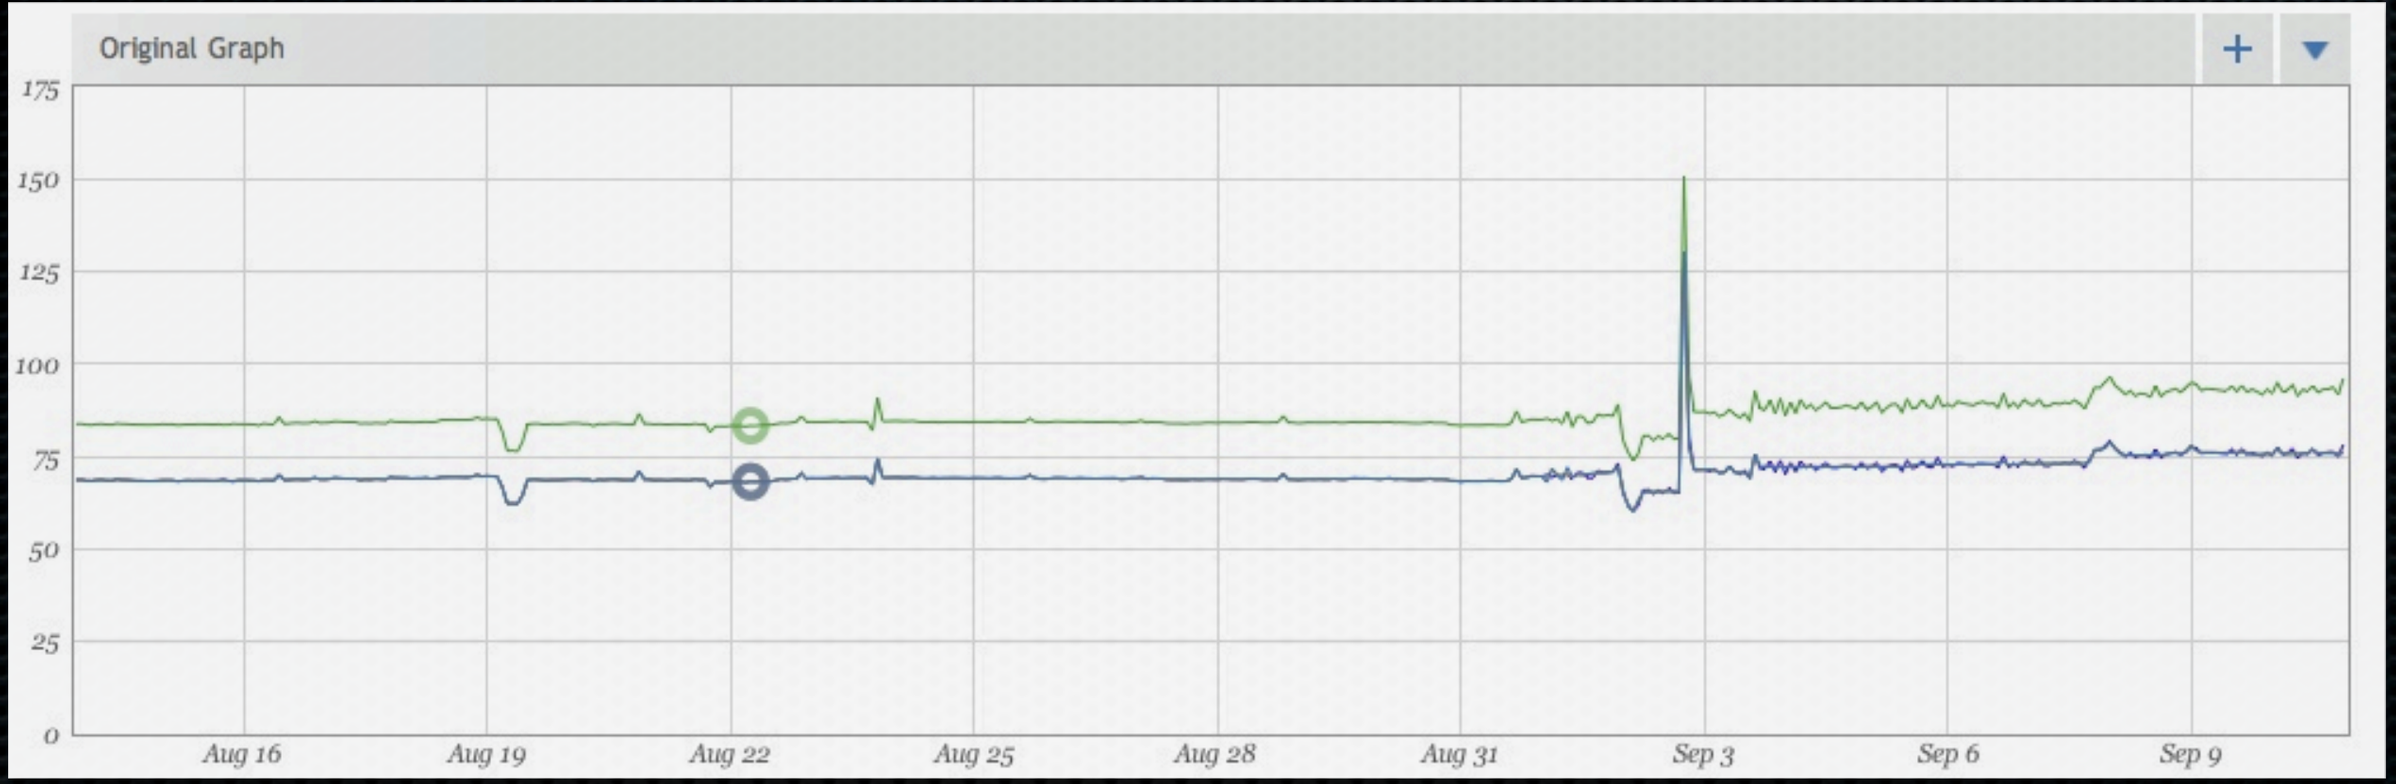

Write IOPS

Simple, relatively stable, easy to trend...
and wrong.

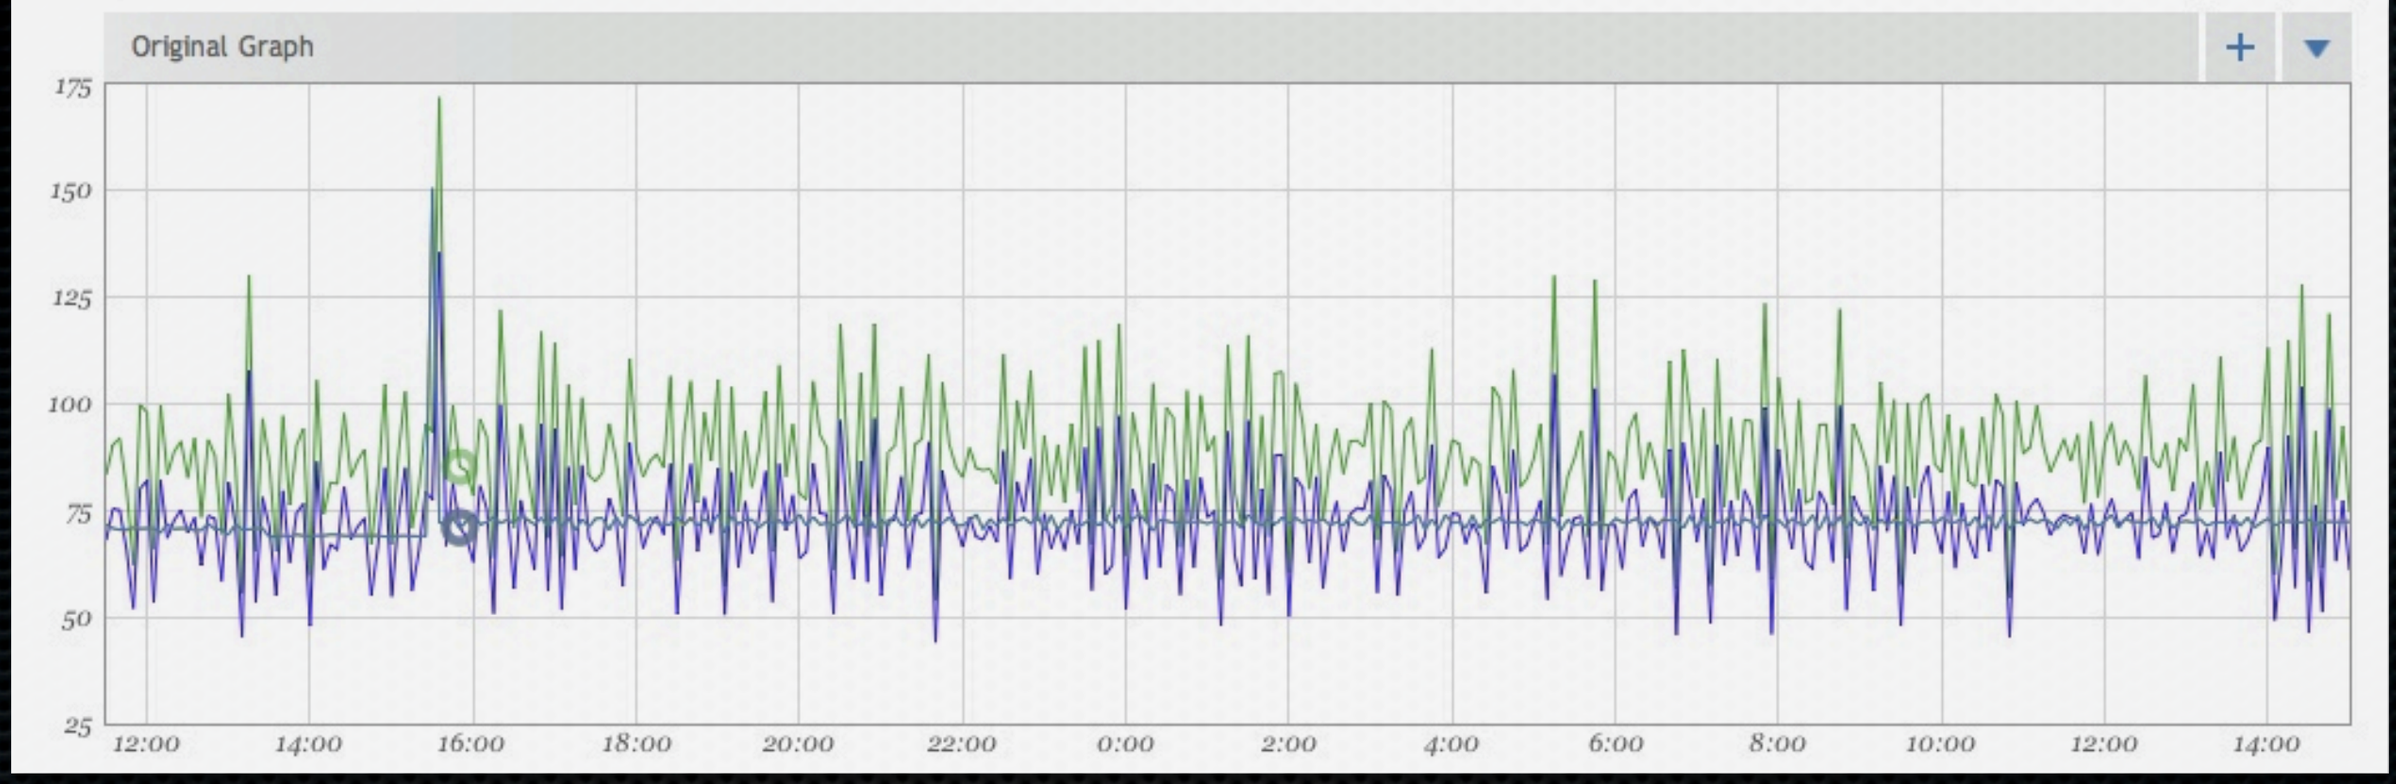

Zooming in shows insight.

It is not as smooth as we thought.

Good to know.

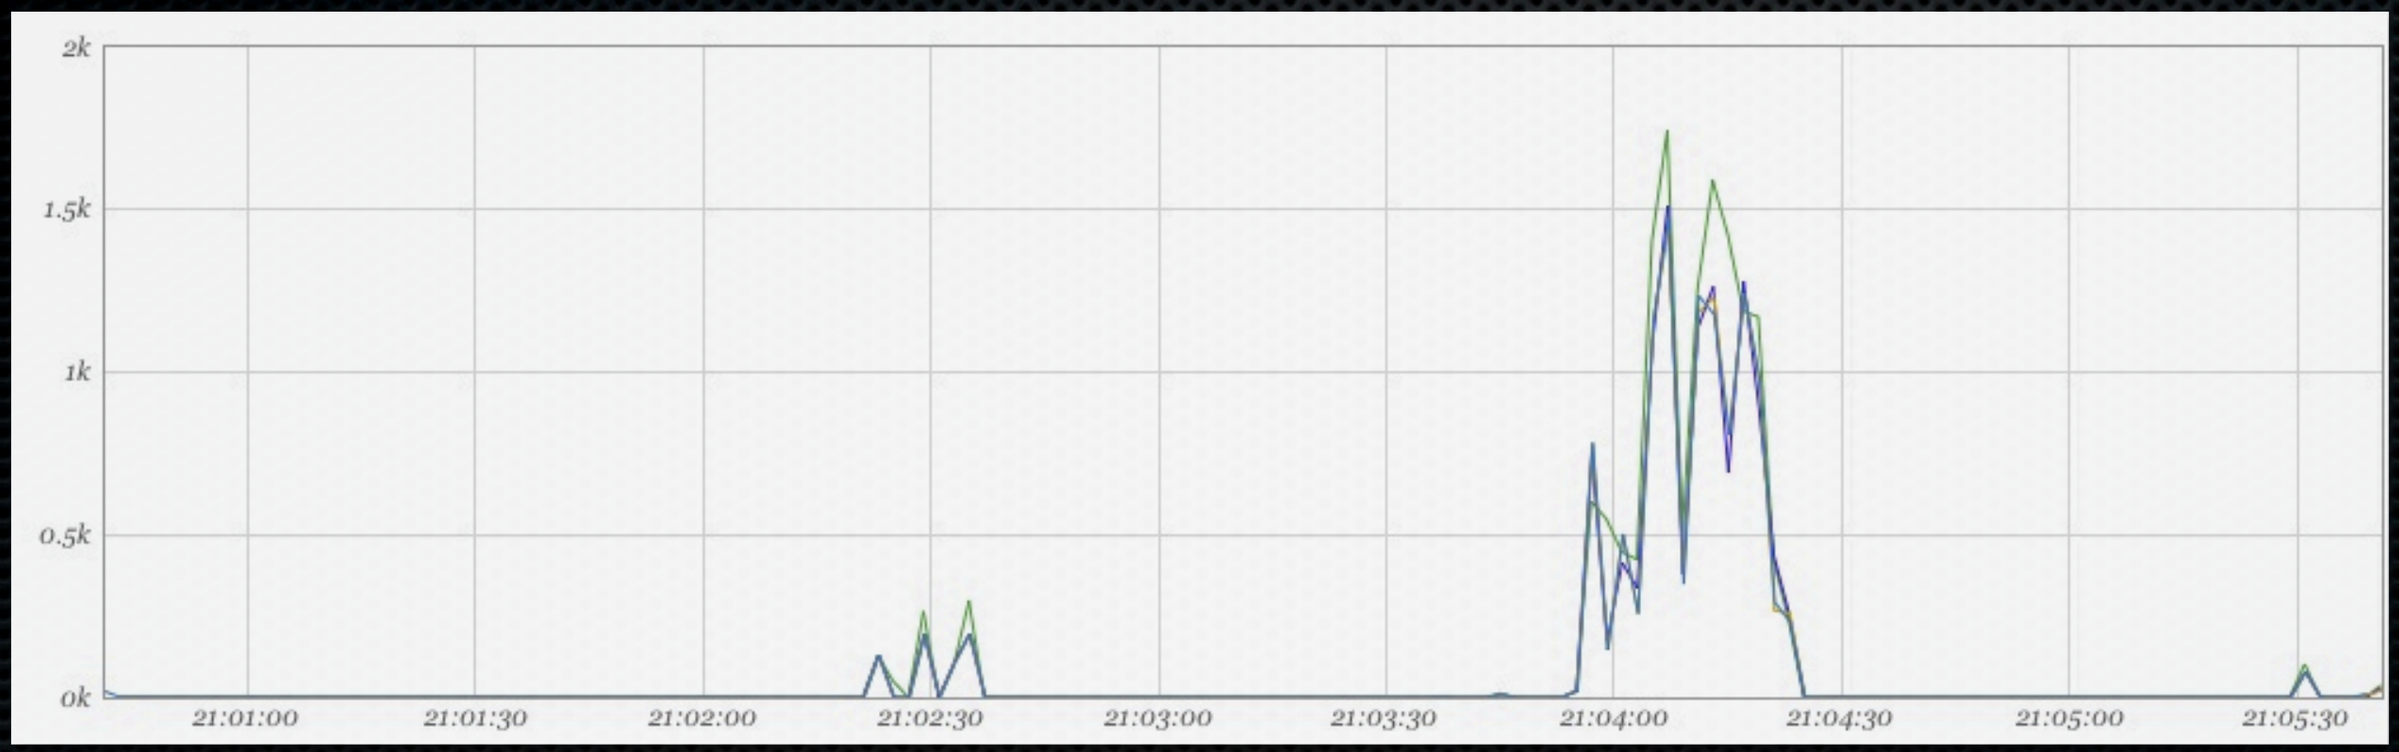

Dishonesty Revealed.

This is “inside” a 5 minute sample.

Note: we’re still sampling; now at 0.5 Hz.

You've Been Deceived.

1. Denial and Isolation
2. Anger
3. Bargaining
4. Depression
5. Acceptance

You've Been Deceived.

let's just skip this one, okay?

1. Denial and Isolation
2. Anger
3. Bargaining
4. Depression
5. Acceptance

You've Been Deceived.

let's just skip this one, okay?

1. Denial and Isolation

If you're not angry, get angry:
it's the next step in healing

2. Anger

3. Bargaining

4. Depression

5. Acceptance

You've Been Deceived.

1. Denial and Isolation
2. Anger
3. Bargaining
4. Depression
5. Acceptance

let's just skip this one, okay?

If you're not angry, get angry:
it's the next step in healing

This is where it gets interesting.

Bargaining.

- I can't reasonably store all the data:
  - likely true.
- Since I can't store all the data,
  I have to store a summarization
  - yes; they're called "aggregates."

Bargaining: you're losing.

- You are sampling things and losing data.
  - the speedometer example.

or

- You are taking an average of a large set of samples.

“Average” is a bad bargain.

- The average of a set of numbers is:
  - the single most useful statistic
- By itself, it provides very little insight into the nature of the original set.
- Hello statistics!

Rounding out your statistics.

- You should store, in order of importance:
  - average
  - cardinality
  - standard deviation
  - covariance
  - min, max
  - derivatives
  - derivative stddev
  - derivative covariance
  - derivative min, max

Rounding out your statistics.

- You should store, in order of importance:
  - average
  - cardinality
  - standard deviation
  - ~~covariance~~
  - ~~min, max~~
  - derivatives
  - derivative stddev
  - ~~derivative covariance~~
  - ~~derivative min, max~~

FWIW: we don't store these at Circonus (today)

What statistics give you.

- With the cardinality, average and stddev
  - you can understand if a “reasonable” number of the data points are “near” the average.
- However, this is often misleading itself.
- We expect normal distributions or (more likely) gamma distributions

So, we can make due.

- No.
- Let’s revisit the five stages of grieving.

You've Been Deceived.

1. Denial and Isolation
2. Anger
3. Bargaining
4. Depression
5. Acceptance

let's just skip this one, okay?

If you're not angry, get angry:
it's the next step in healing

This is where it gets interesting.

You've Been Deceived.

1. Denial and Isolation
2. Anger
3. Bargaining
4. Depression
5. Acceptance

let's just skip this one, okay?

If you're not angry, get angry:
it's the next step in healing

This is where it gets interesting.

Here we are.

Gamma Gamma Gamma

average: 0.72

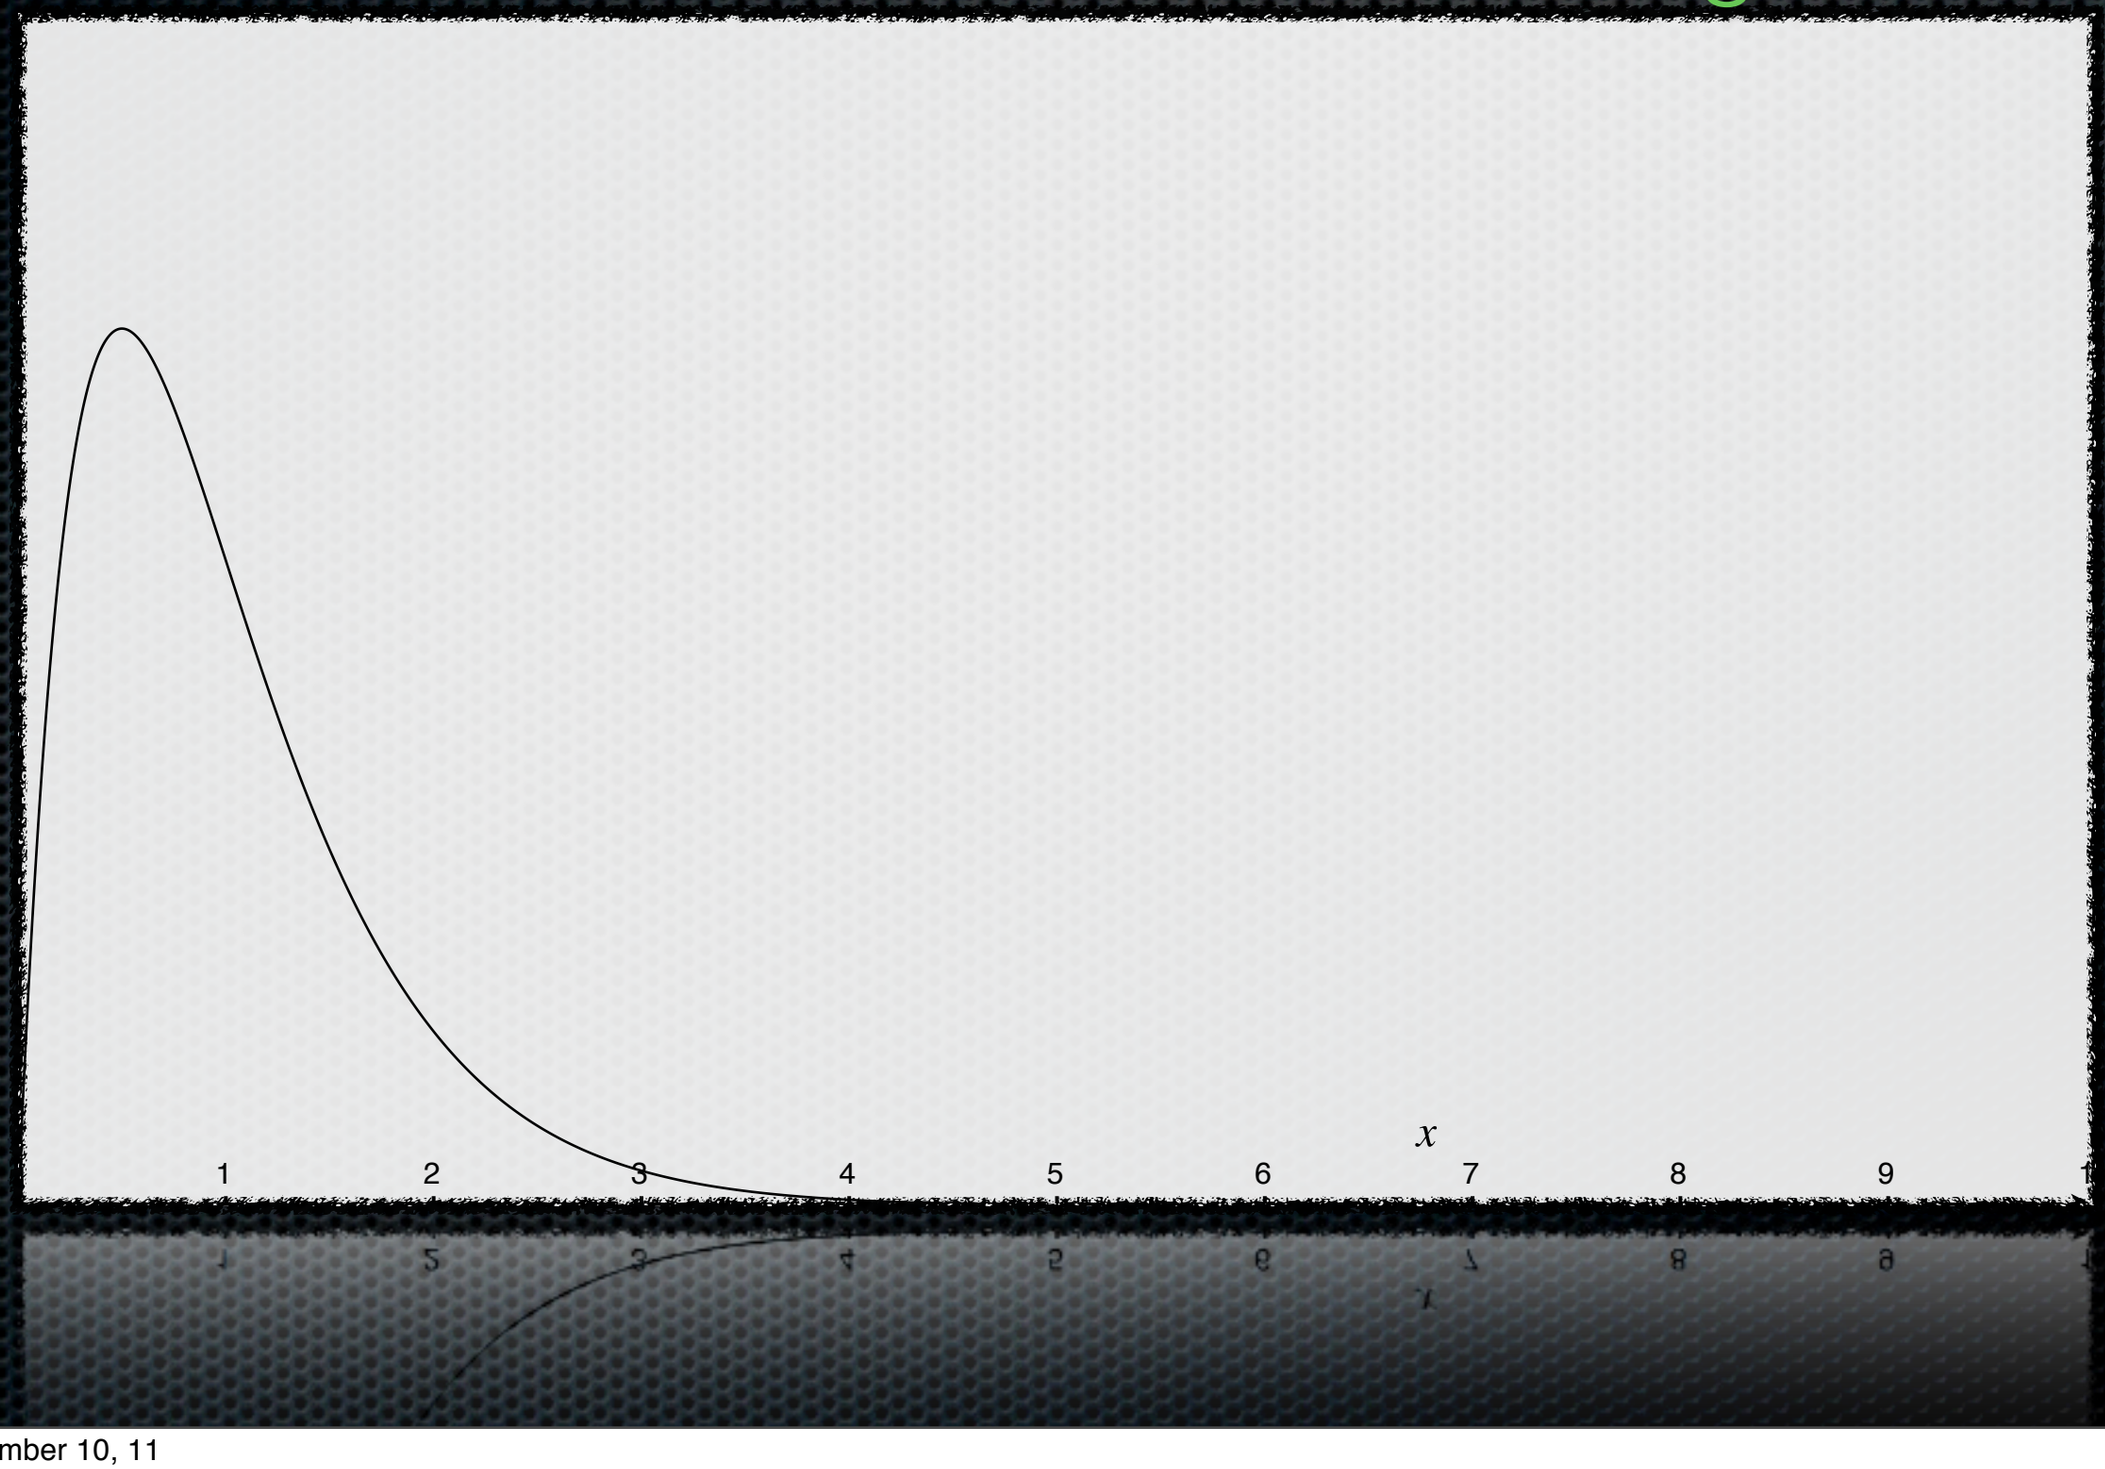

Gamma Gamma Gamma

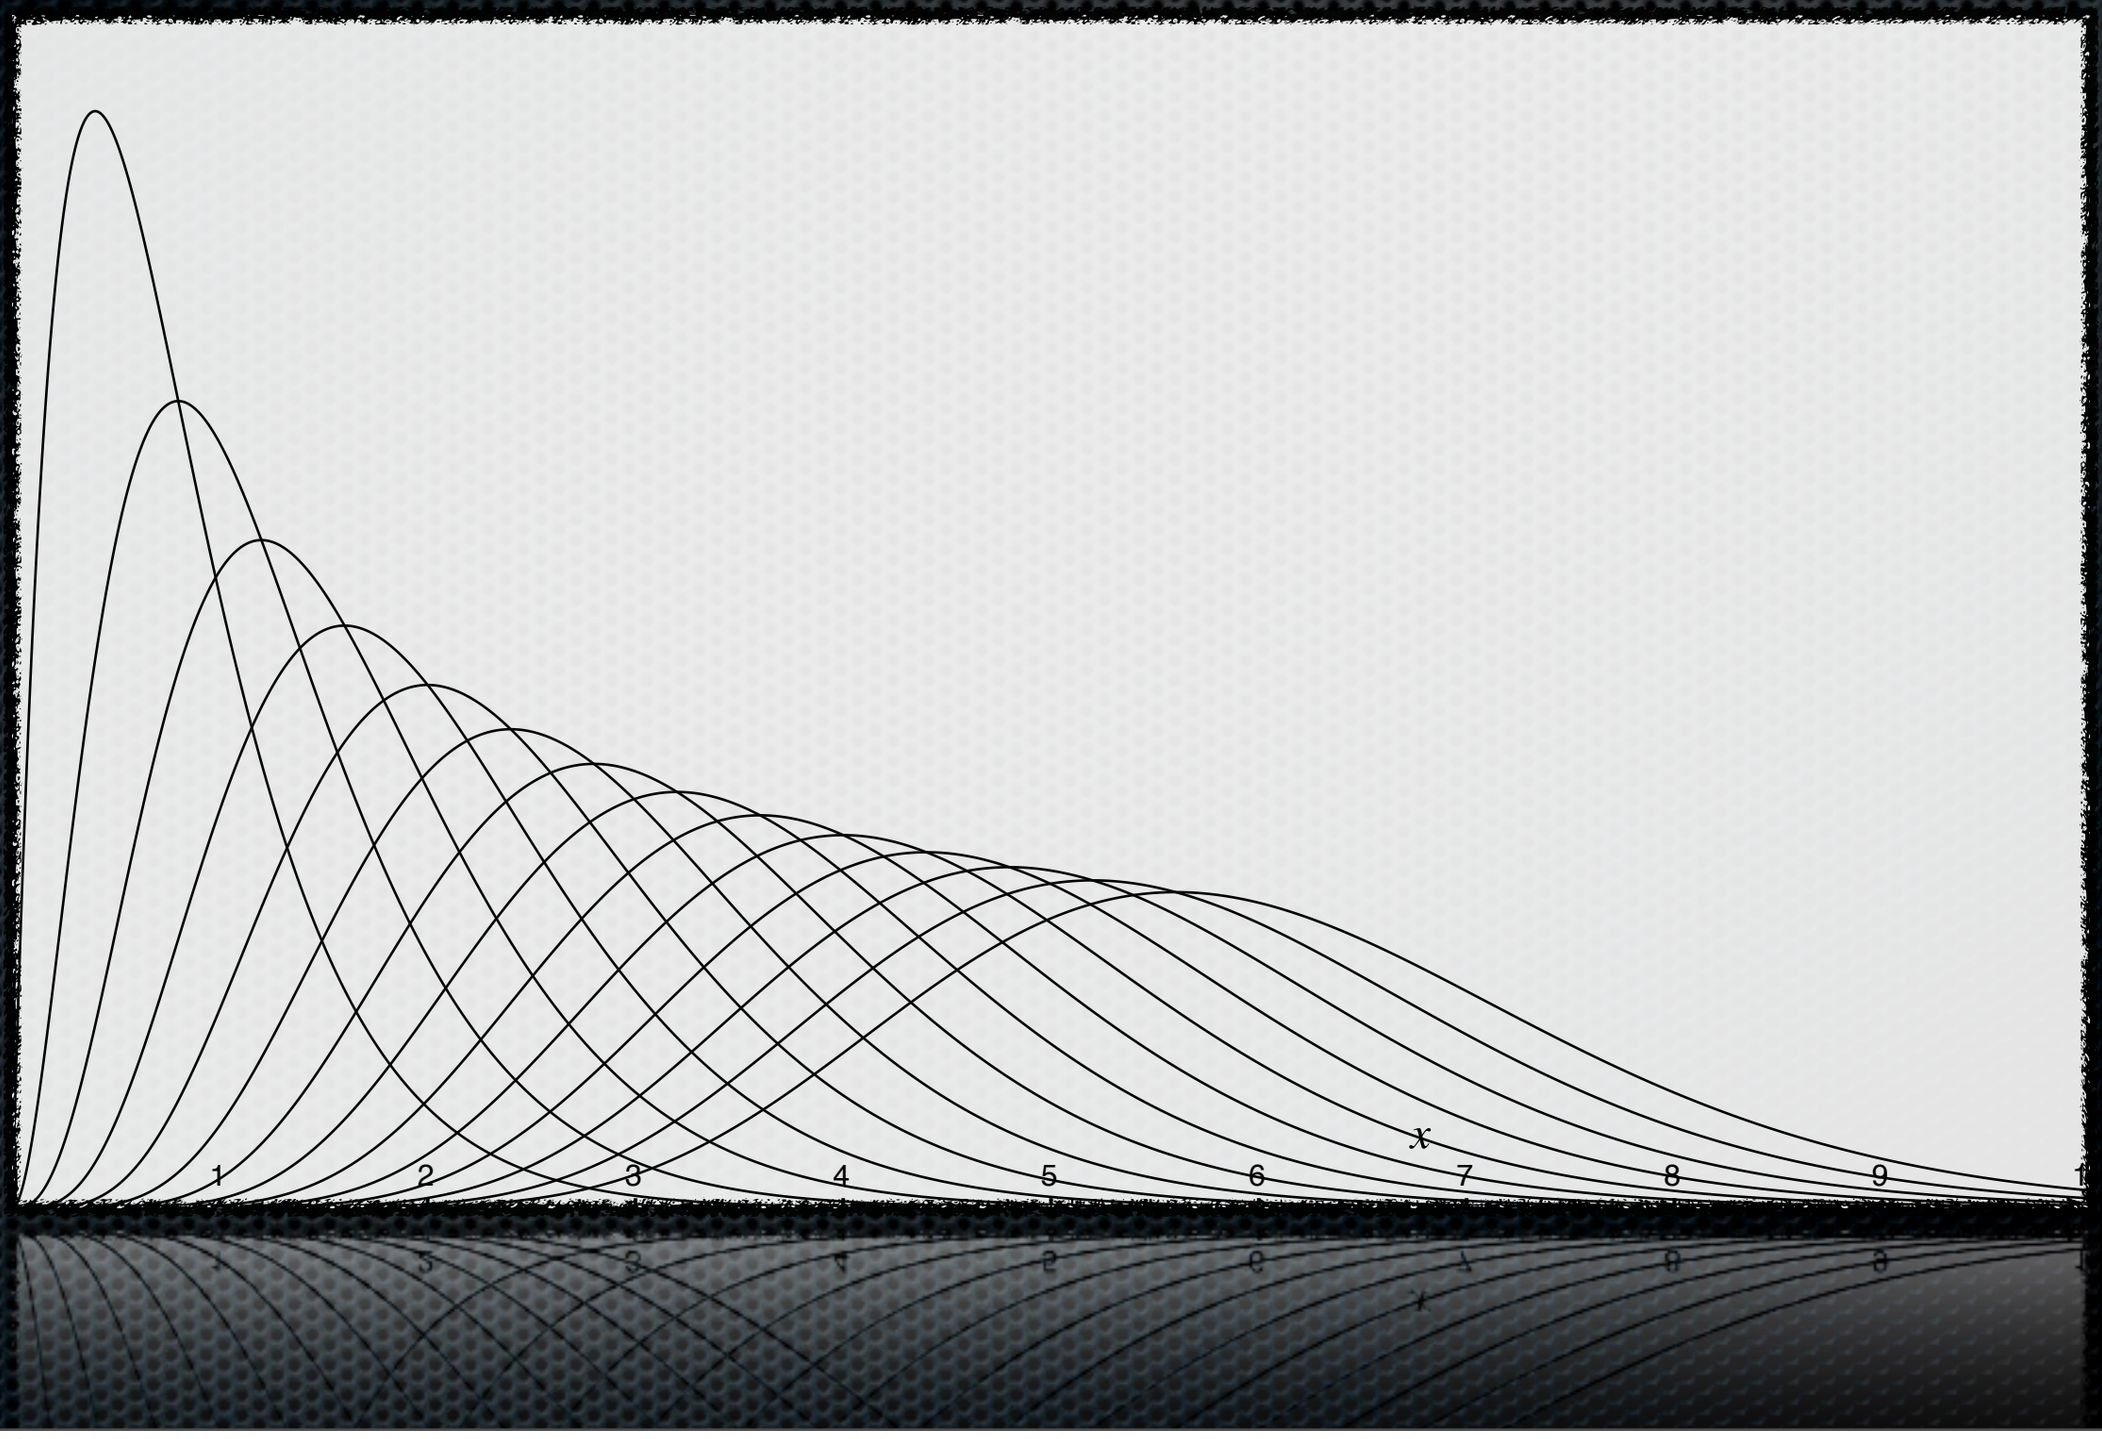

Gamma Gamma Gamma

average: 1.16

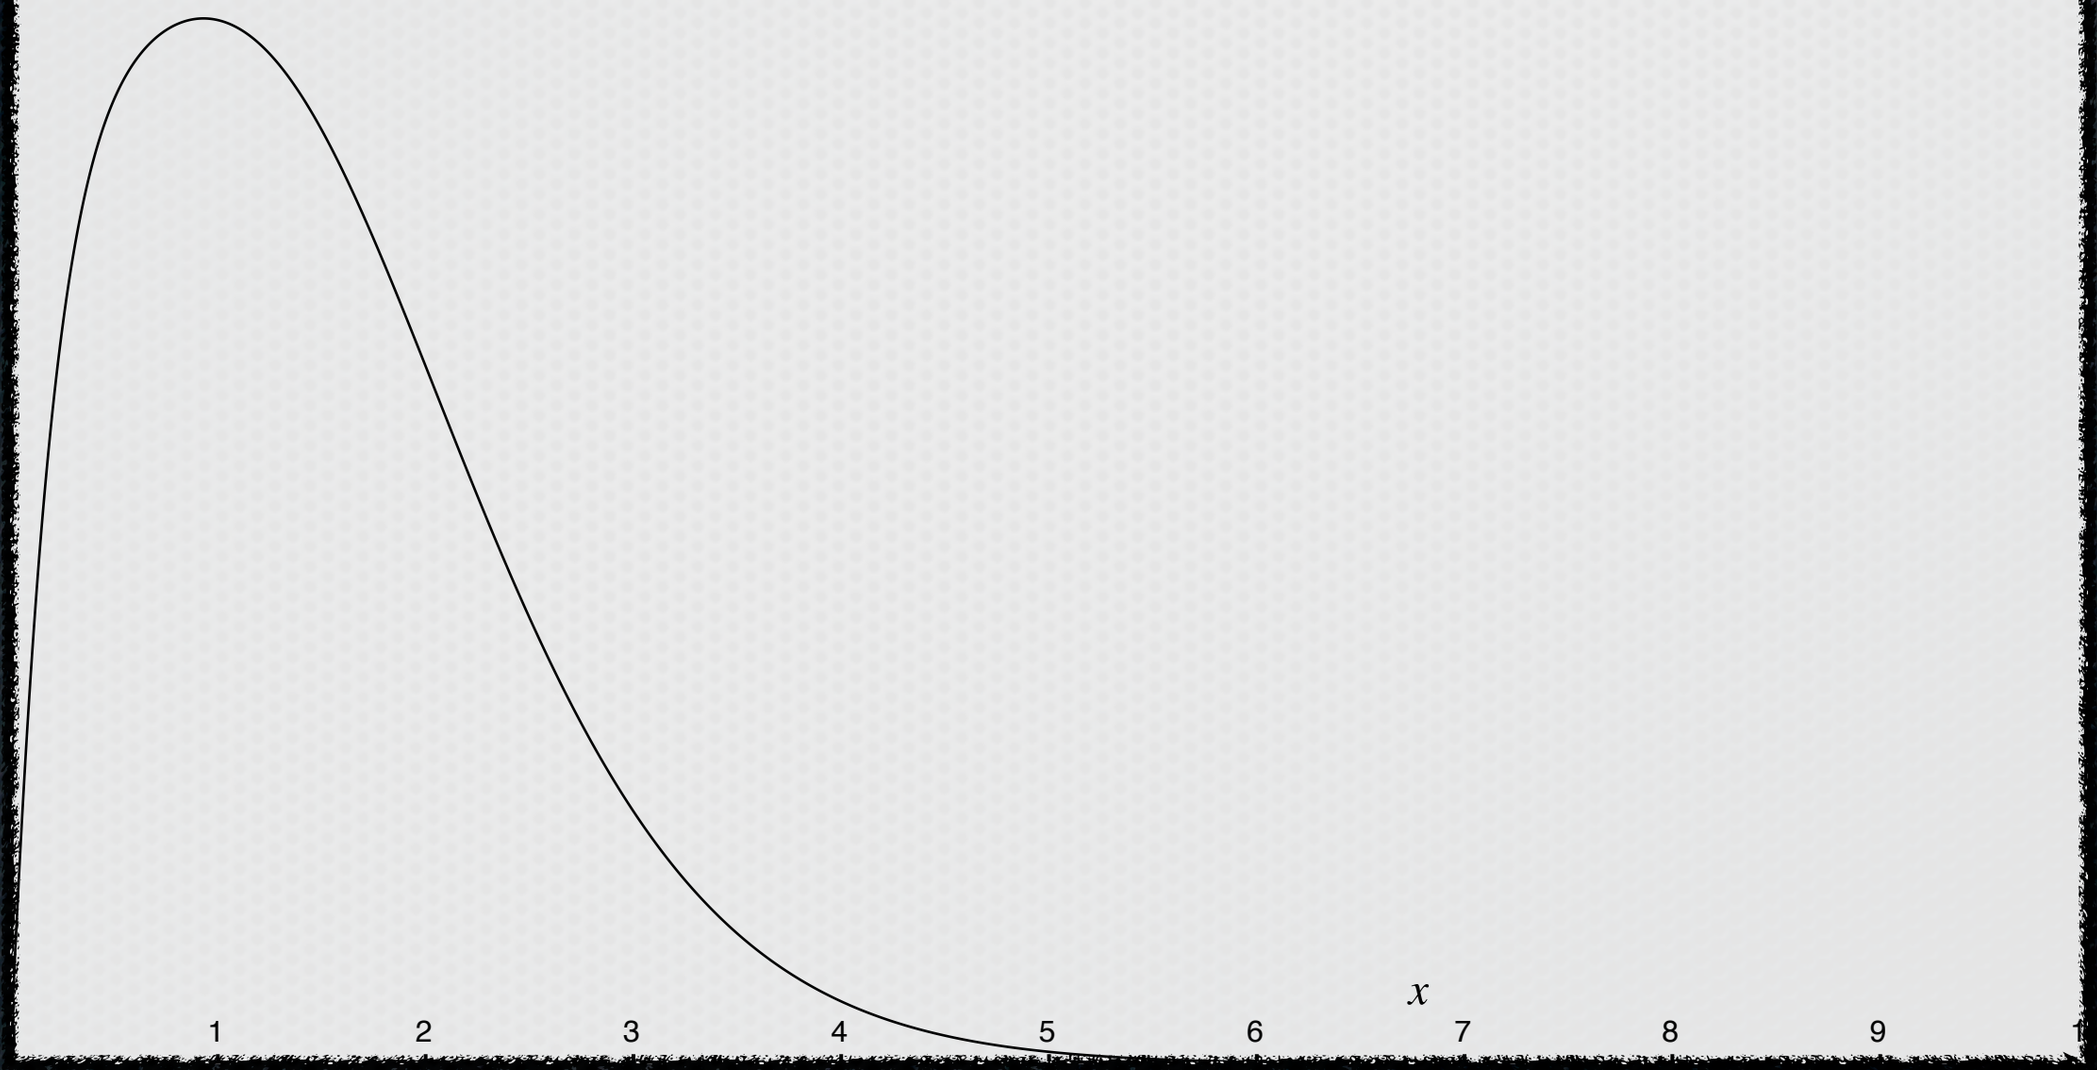

Gamma Gamma Gamma

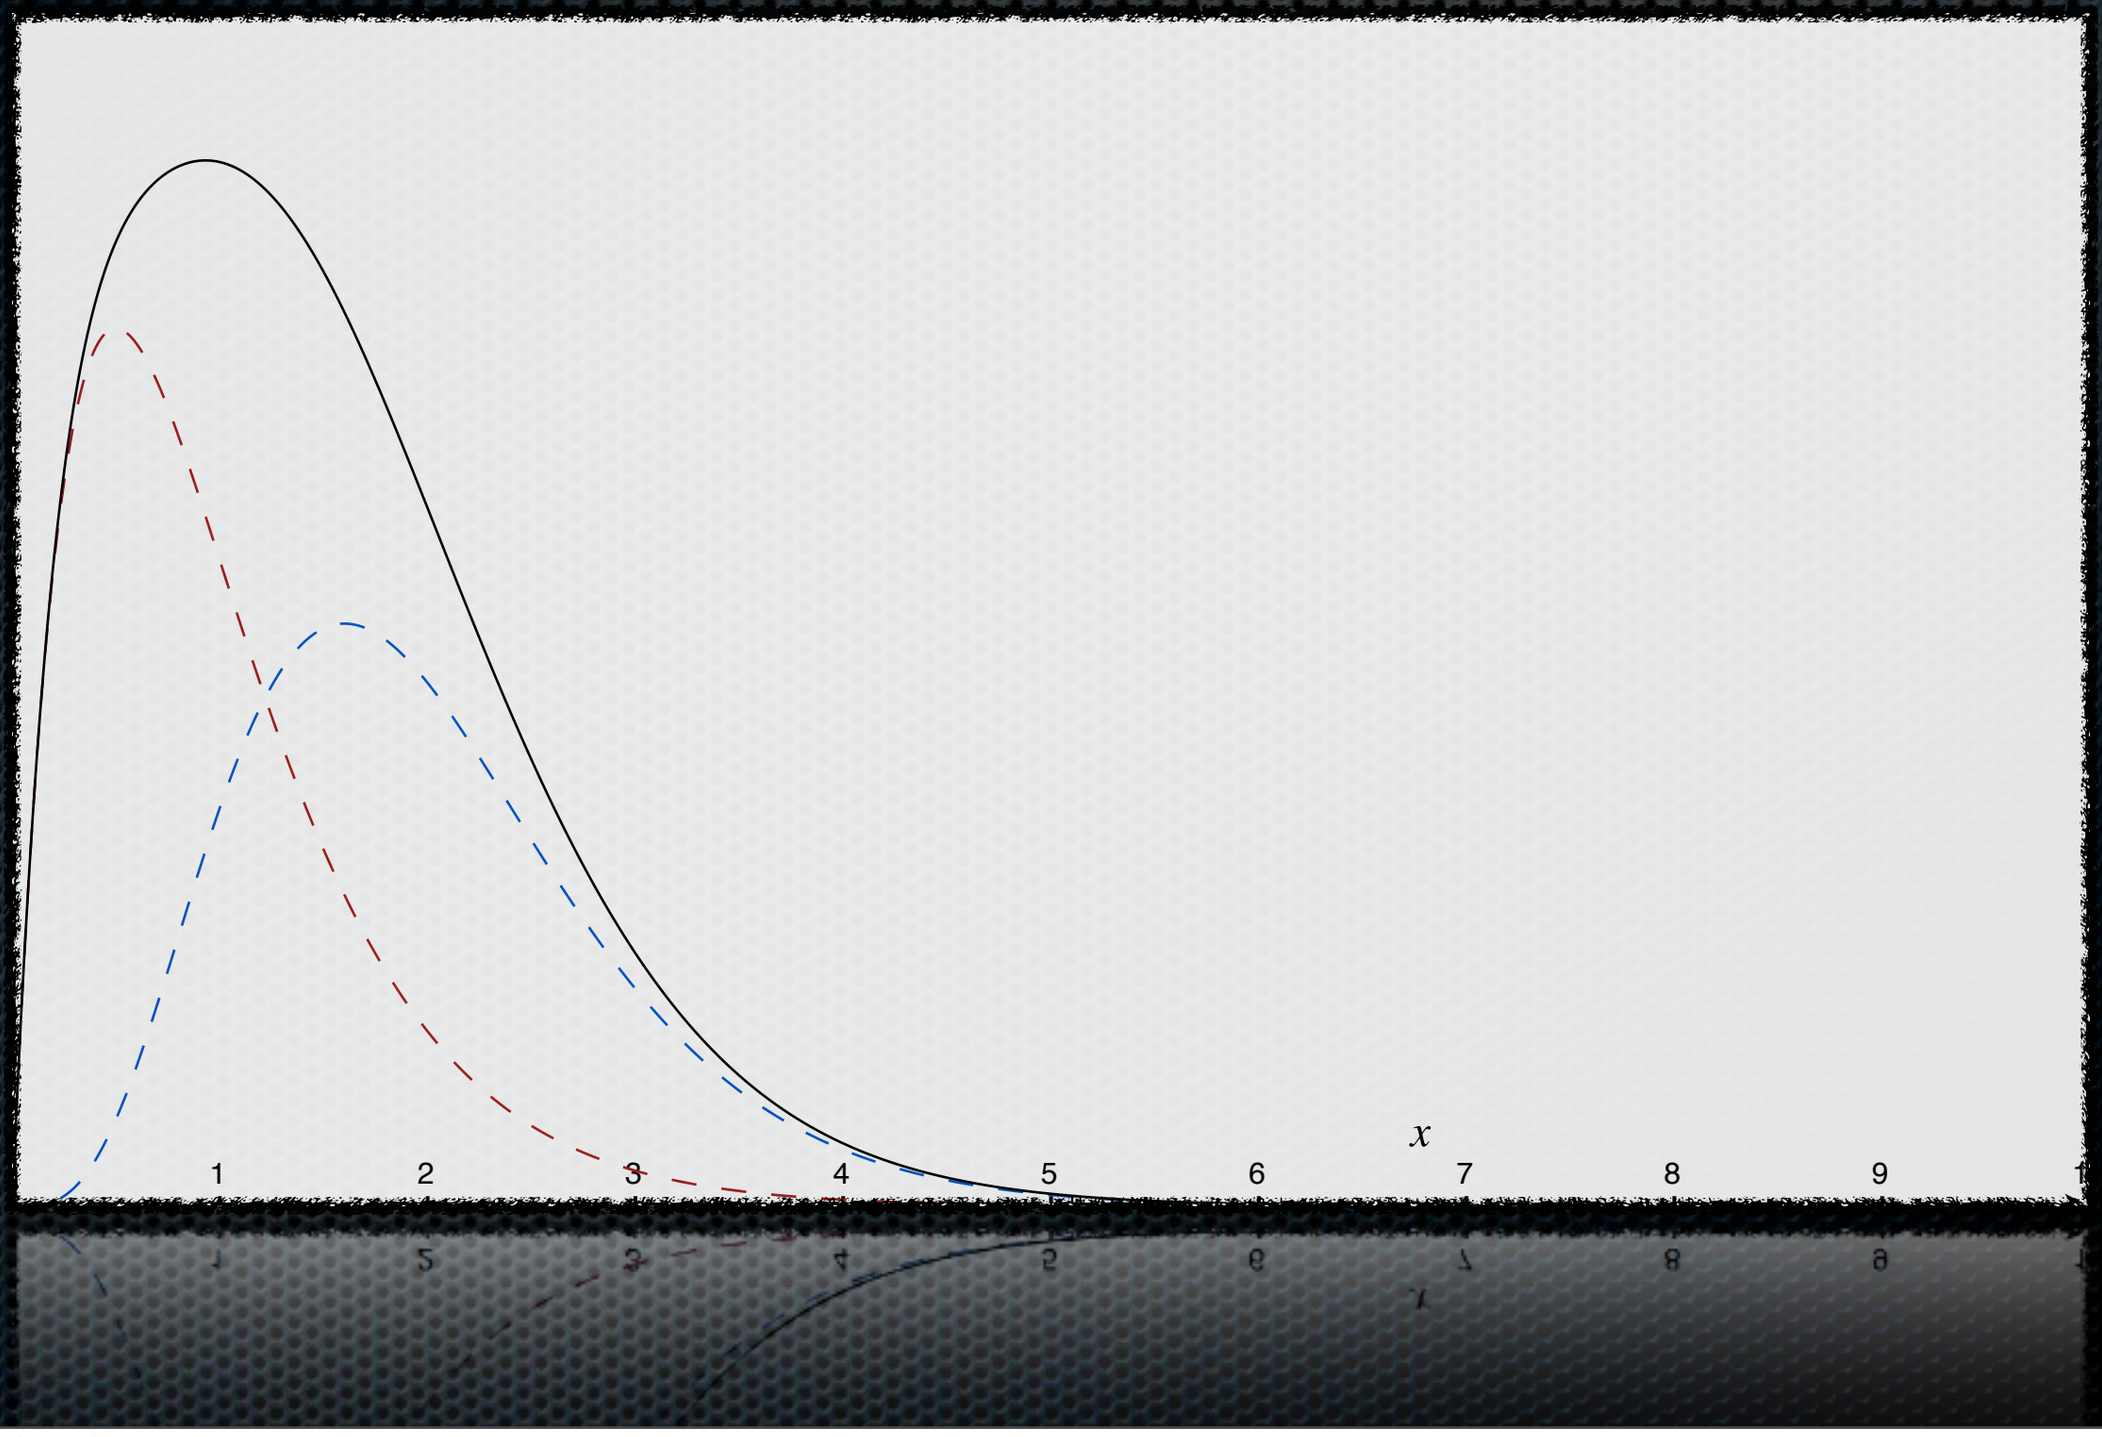

Gamma Gamma Gamma

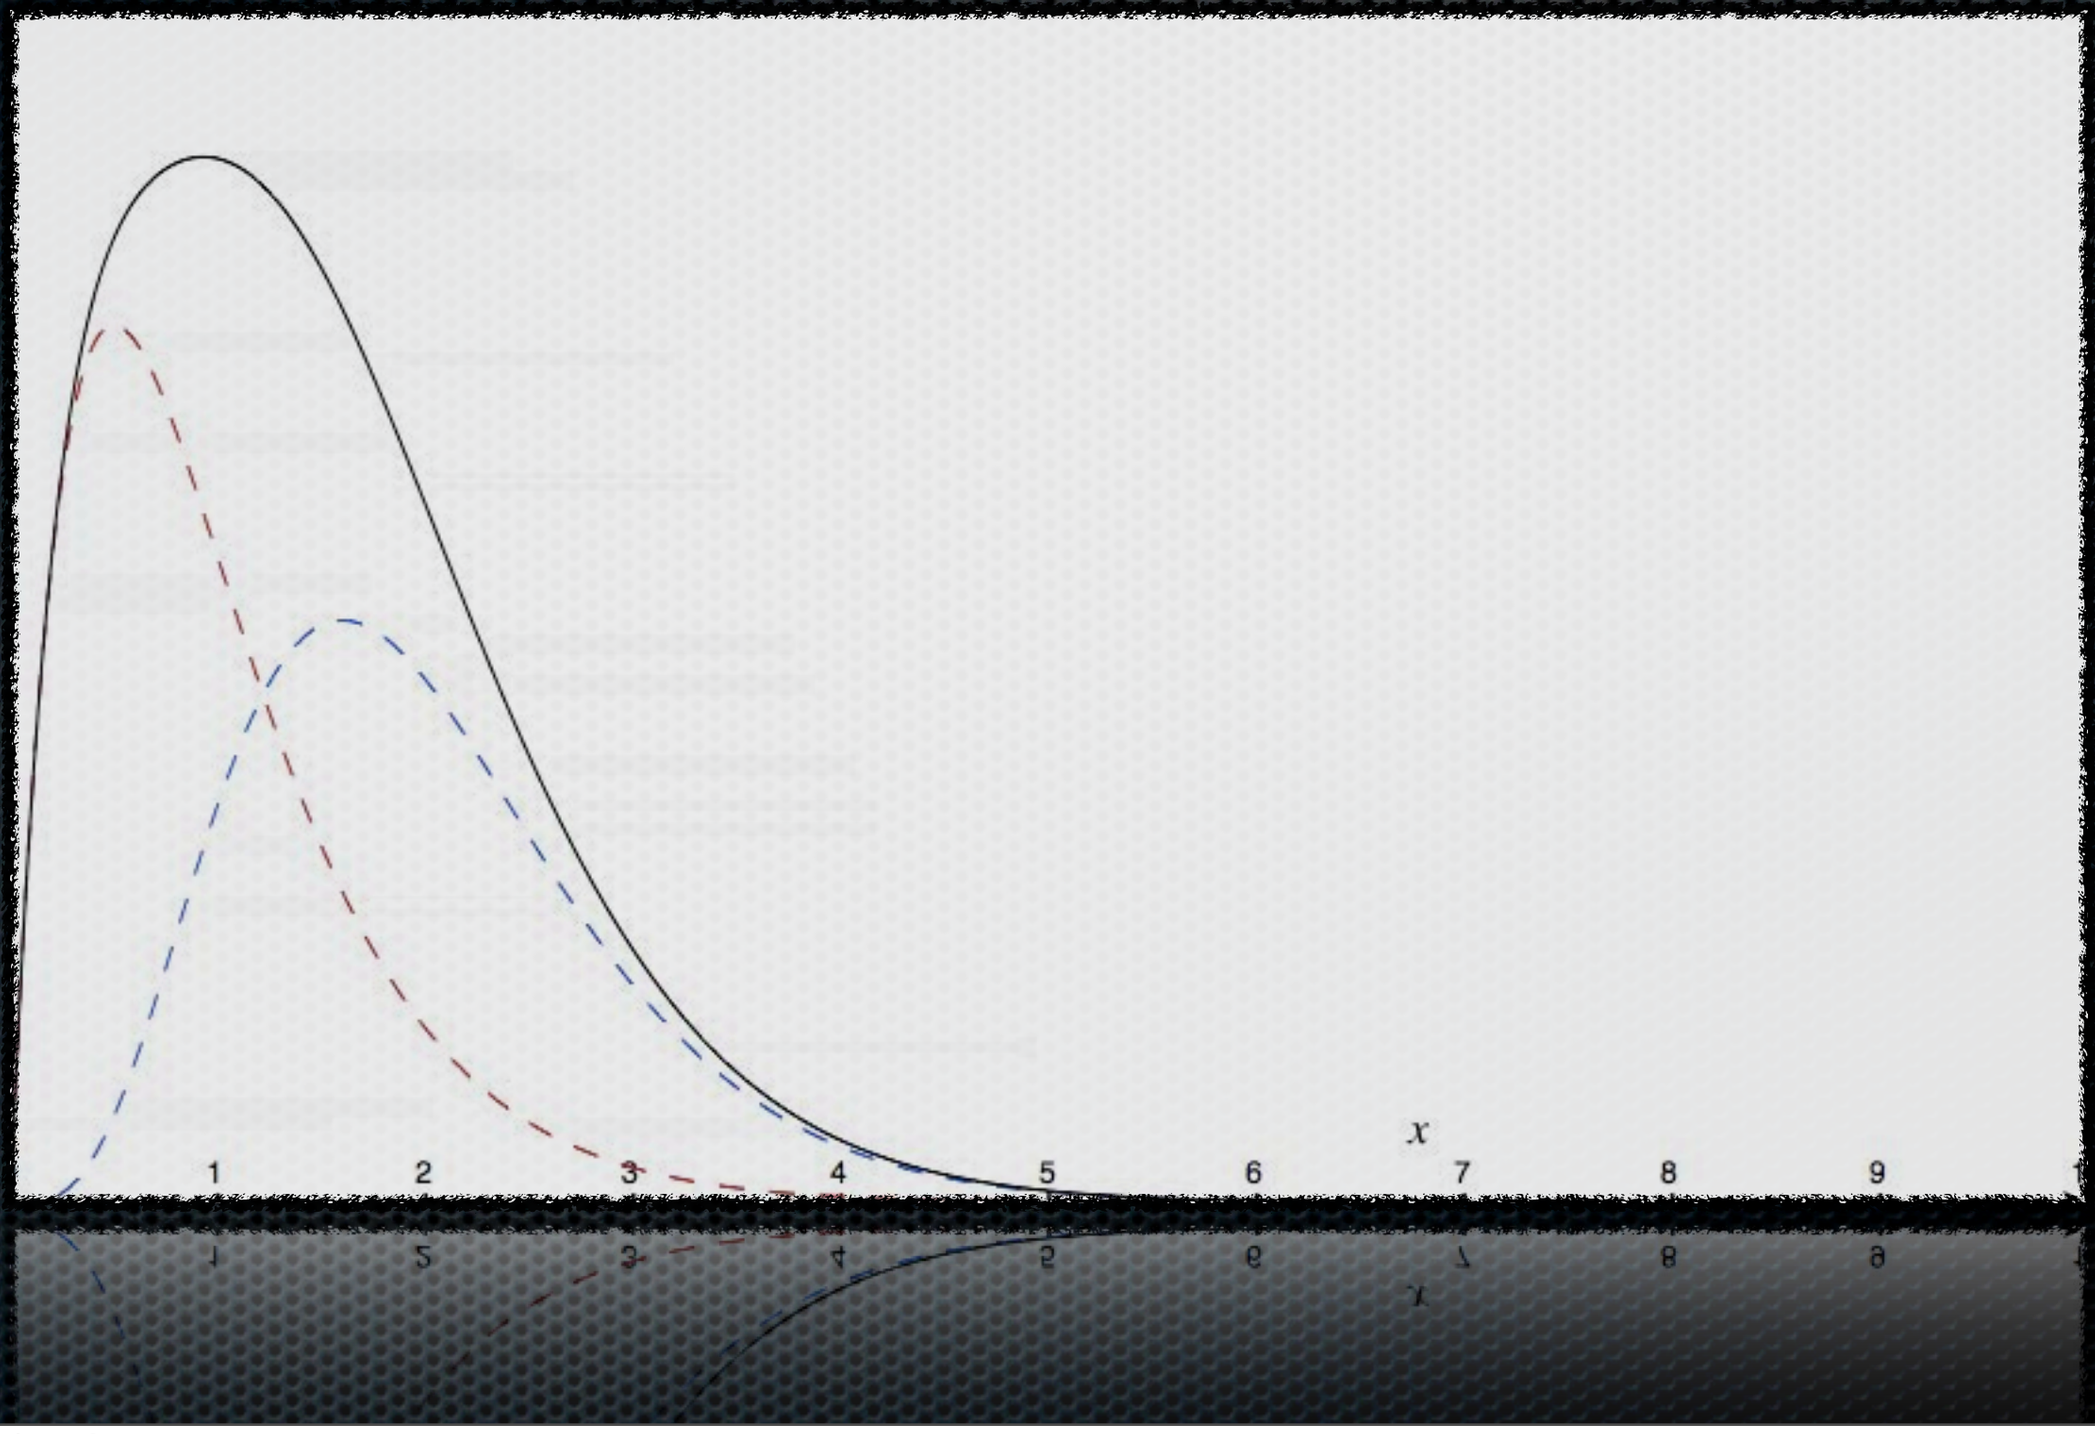

Gamma Gamma Gamma

average: 4.67

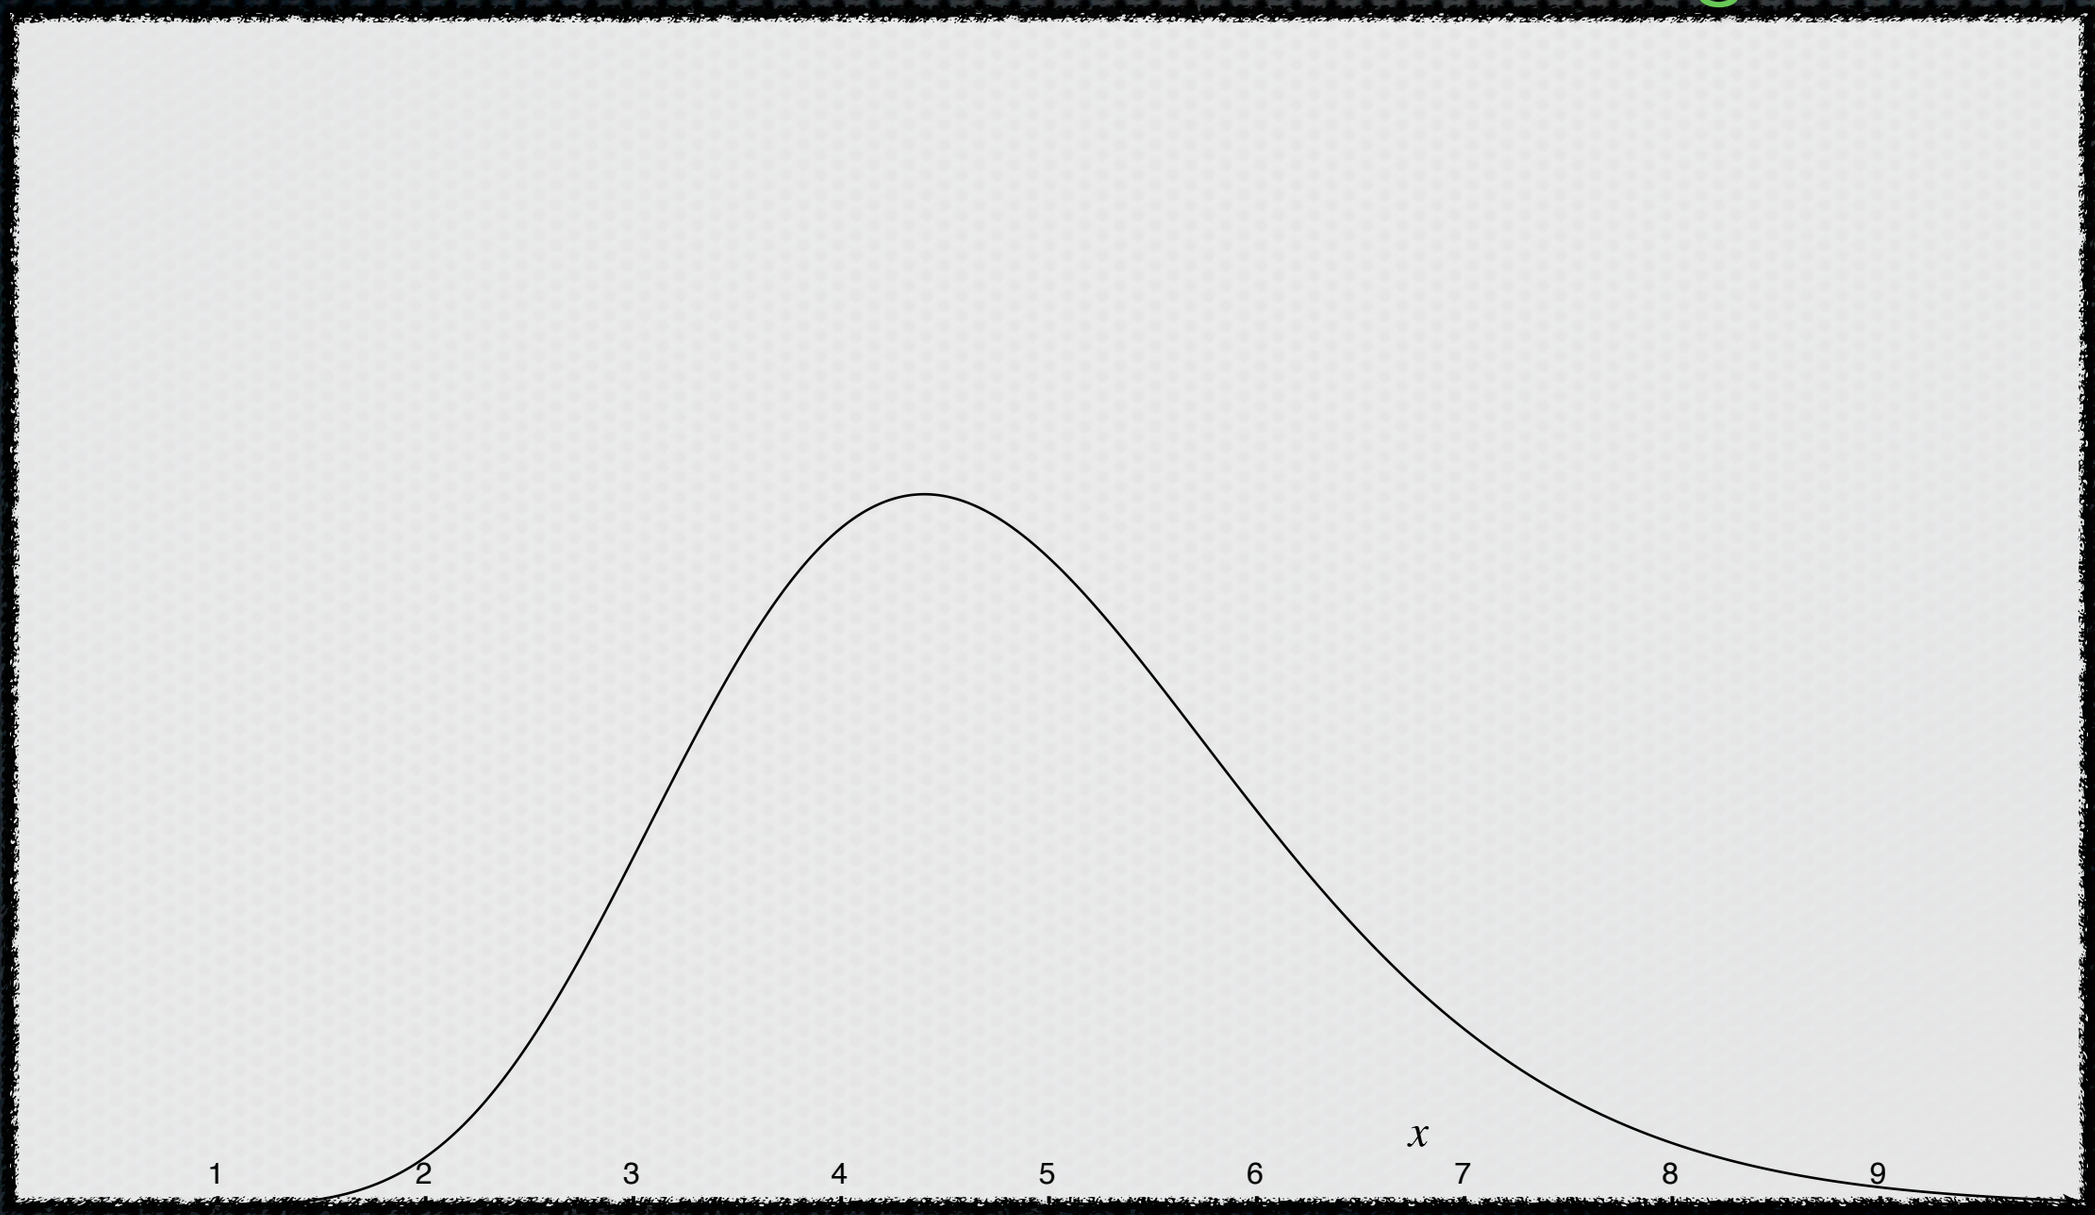

Do we need more?

- Wait. Stop.
- In some systems we know enough.
- We know their distributions,

  ... the leather seat, the inertia of the vessel.
- A sample here and there, an average, etc.
- Good enough: temperature, load, payments, bugs.
- Not enough: disk IOPS, memory use, CPU time.

We do need more.

- we have a lot of data sources from which
- we typically compose an average of many values, and
- understanding the distribution is extremely valuable.

Histograms: the need.

- Given an average of 4, and a cardinality of 1000 maybe you had 1000 samples of exactly 4? Dunno.

- If you have a stddev of 1, about 68.2% of your data points will be between 3 and 5.

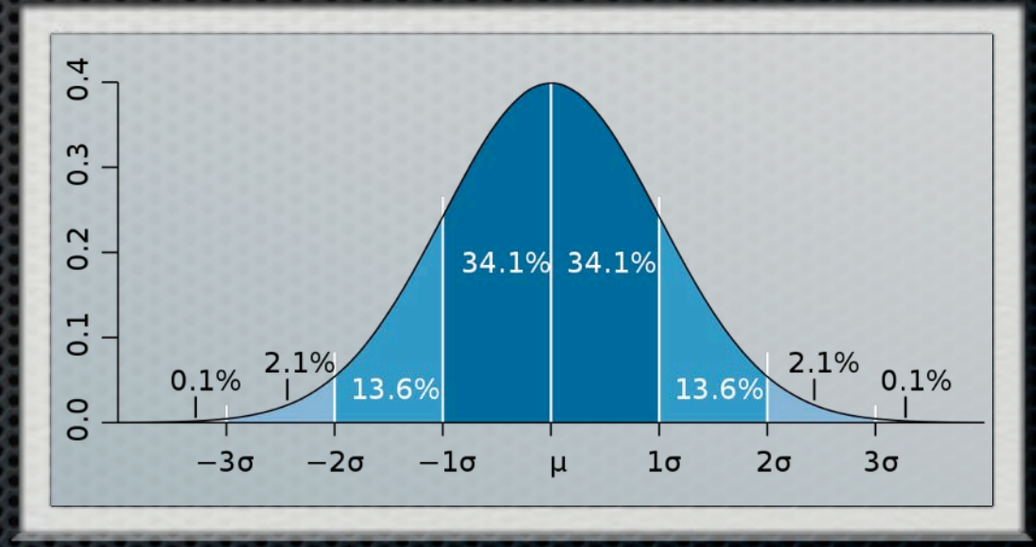

- This will never be as detailed as a histogram of the full dataset.

Histograms: the want

- Ideally, we'd want histograms bucketed as we see fit with all data represented.
- Too much data: this isn't tractable.
- So, how do we bucket?

Bucketing Basics

- If I had numbers bounded by a small space:
  - between 0ms and 10ms
- I could bucket [0-1),[1-2),...[9,10).
  - eleven buckets, fairly manageable.
- What do I do when I get 20? or 200? or 20000?
- What do I do when I get lots of data between 0.001 and 1?

How we do more.

- We make an assumption (yes: ass: u & mption)
- We assume the same thing floating point assumes:
  - The difference between “very large” numbers and the difference between “very small” numbers can have accuracy inversely proportional to their magnitude.
- We take this and build histograms.

Solution: log buckets

- [0.001,0.01),
  [0.01,0.1),
  [0.1,1),
  [1,10),
  [10,100),
  [100,1000),
  [1000,10000),
  [10000,100000)

- $[2^{-10},2^{-9}),[2^{-9},2^{-8}),[2^{-8},2^{-7}),$
  $[2^{-7},2^{-6}),[2^{-6},2^{-5}),[2^{-5},2^{-4}),$
  $[2^{-4},2^{-3}),[2^{-3},2^{-2}),[2^{-2},2^{-1}),$
  $[2^{-1},1),[1,2),[2,4),[4,8),$
  $[8,16),[16,32),[32,64),$
  $[64,128),[2^{7},2^{8}),[2^{9},2^{10}),$
  $[2^{10},2^{11}),[2^{11},2^{12}),[2^{12},2^{13}),$
  $[2^{13},2^{14}),[2^{14},2^{15})$

Engineers might like these,
but humans do not.

Histograms: the compromise

- Log-linear:
  - 0.1 , 0.15 , 0.2 , 0.25 , 0.3 , 0.35 , 0.4 , 0.45
    0.5 , 0.6 , 0.7 , 0.8 , 0.9
    1.0 , 1.5 , 2.0 , 2.5 , 3.0 , 3.5 , 4.0 , 4.5
    5.0 , 6.0 , 7.0 , 8.0 , 9.0
    10...
- I first witnessed this in DTrace courtesy of Bryan Cantrill

Where are we now?

- For some data, we store simple statistical aggregations
- Other bits of data we use log-linear histograms
  - quite insightful, while
  - still small enough to store
- A single “five minute average” can look more like this:

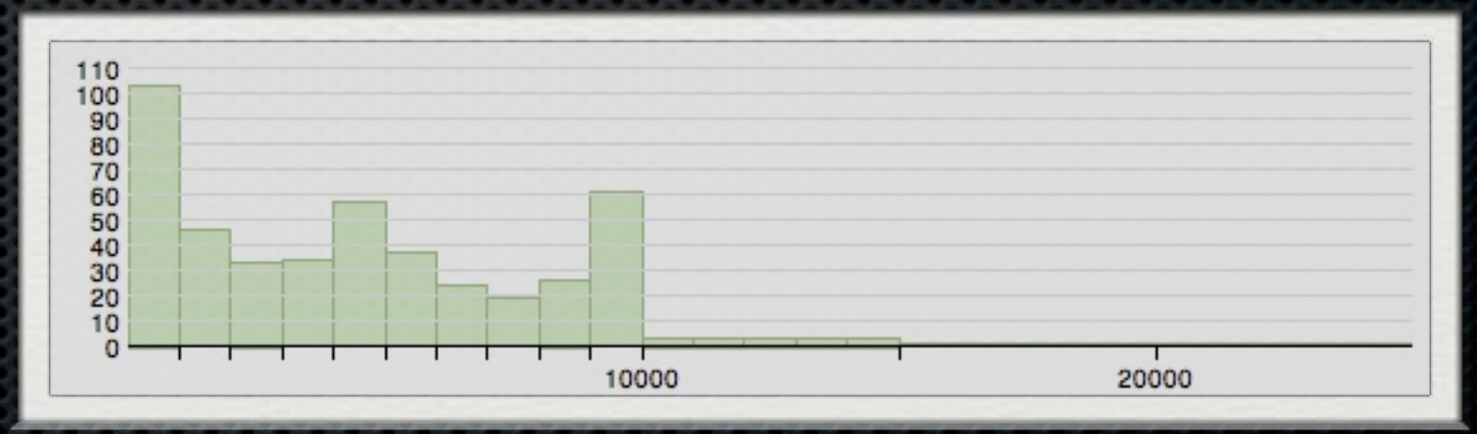

You've Been Deceived.

1. Denial and Isolation
2. Anger
3. Bargaining
4. Depression
5. Acceptance

let's just skip this one, okay?

If you're not angry, get angry:
it's the next step in healing

This is where it gets interesting.

We are still here...
at least in the open source world.

Reconnoiter

- Reconnoiter:
  - https://labs.omniti.com/labs/reconnoiter
  - https://github.com/omniti-labs/reconnoiter
- Open source (BSD)

Reconnoiter: backend

- C (as God intended),
  with lua bindings to ease plugin development
- Java (external) bridge for making things like JMX easier
- Esper (Java) for streaming queries (fault detection)
- AMQP and REST for component intercommunication
- PostgreSQL backing store
- Very actively developed

Reconnoiter: frontend

- HTML + CSS
- Javascript: JQuery, flot
- PHP for AJAX APIs
- Needs engineers.

Features (1)

- Data Collection:
  - most standard protocols (SNMP, REST, JDBC, etc.)
  - some nice bridges to legacy systems: (nagios, munin, etc.)
- Data Storage:
  - supports sharding
  - never delete old data

Features (2)

- We don't do histograms now
- Because all the old data is stored, we could
  - If we have enough data to make this worthwhile, we will likley need an on-the-fly approach
- We store raw data and rollups including:
  - average, cardinality, stddev, derivatives (and counters)

Features (3)

- separates data from visualization
- drag-n-drop graphing
- real-time data:
  - you can run checks “often” (every 100ms?)
  - you can see that data in the browser
  - you can compose a graph from disparate data

Active vs. Passive

- Reconnoiter supports passive metric submission
  - what version of code is deployed
- Reconnoiter's design is focused on active collection
  - how may users have signed up
  - how many visitors to your site
  - how full is your disk
  - how many open bugs
  - how many hours spent

Bias

- Reconnoiter doesn't solve all your problems (we're busy trying to do that with Circonus)

- I believe Reconnoiter is the best tool available for heterogeneous, multi-datacenter, full-architecture monitoring.

- I am seriously biased.

You've Been Deceived.

1. Denial and Isolation

2. Anger

3. Bargaining

4. Depression

5. Acceptance

let's just skip this one, okay?

If you're not angry, get angry:
it's the next step in healing

This is where it gets interesting.

Here we are.

This is on you. Good luck.

Fonts Served Globally

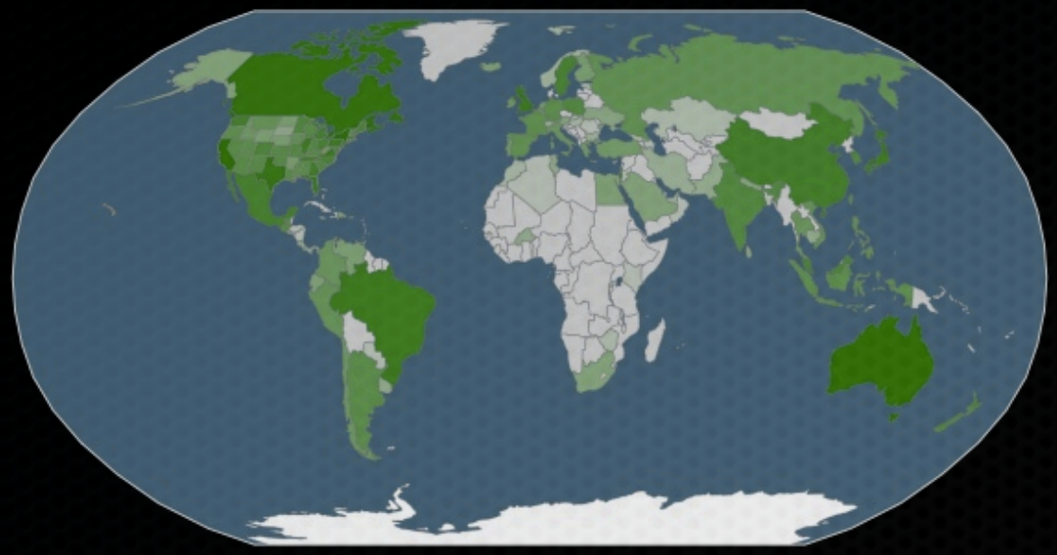

FS Bandwidth

Refusals

Font Usage Leaderboard

FOUNDRY	TYPEFACE	FONT	VIEWS
URW++	Classico	Medium	1723
Neufville Digital	Futura	Bold	1531
URW++	Classico	Regular	1494
URW++	Classico	Italic	1465
URW++	Classico	Medium Italic	1372
Parachute	Square Sans Pro	Bold	1300
Mark Simonson Studio	Proxima Nova	Semibold	1206
Neufville Digital	Futura	Medium	1009
URW++	Clarendon URW	Bold	940
Fontsmith	FS Joey Web	Regular	842

from **UK**

from **Central Europe**

from **US East**

from **US West**

Thank You!

Questions?