



Monitoring 101

THE BASICS



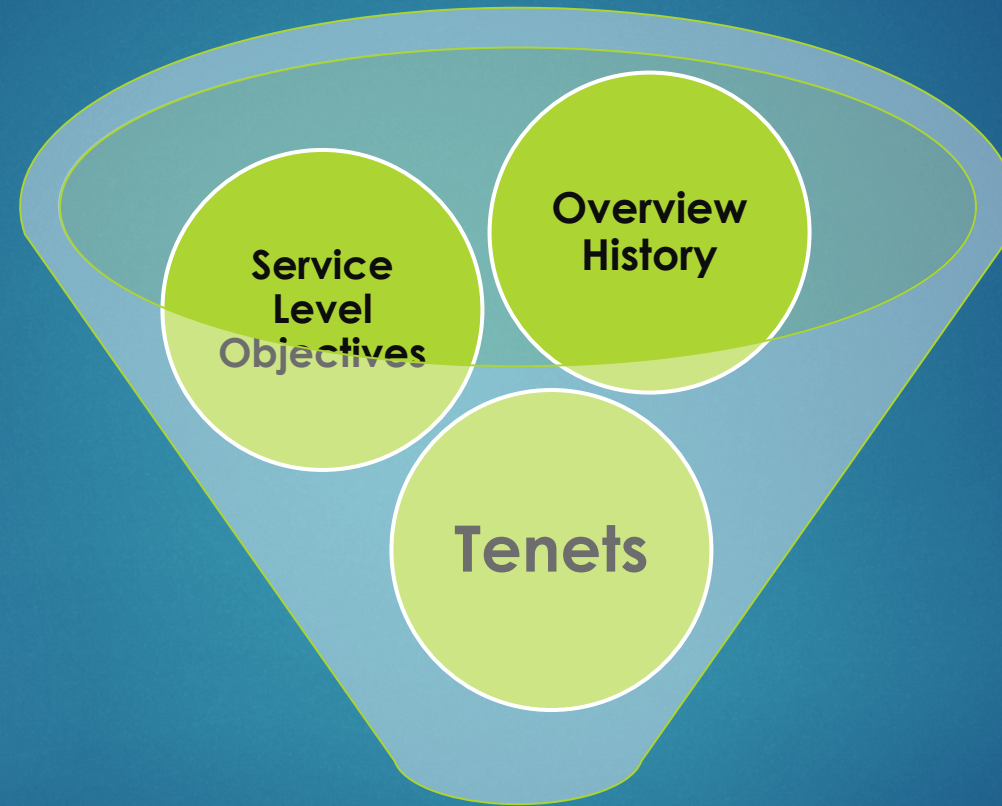
Theo Schlossnagle

CEO @Circonus
Twitter: @postwait



CIRCONUS

Agenda



Navigation Skills



“ Monitoring is the action of observing and checking static and dynamic properties of a system.

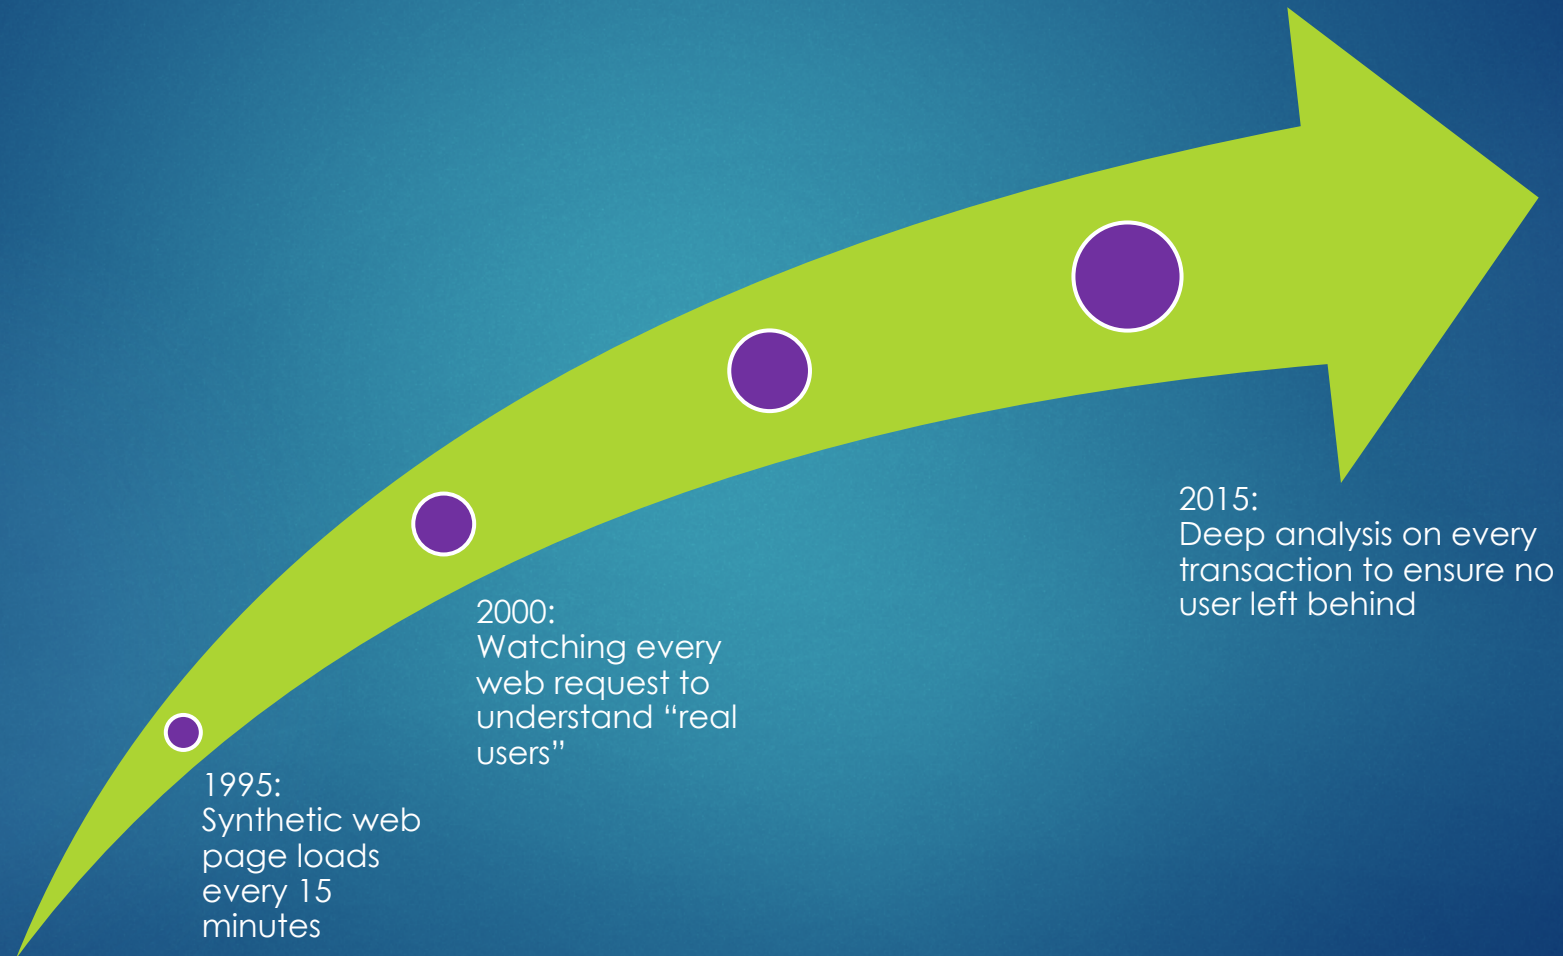
”

- HEINRICH HARTMANN (<http://l42.org/GwE>)

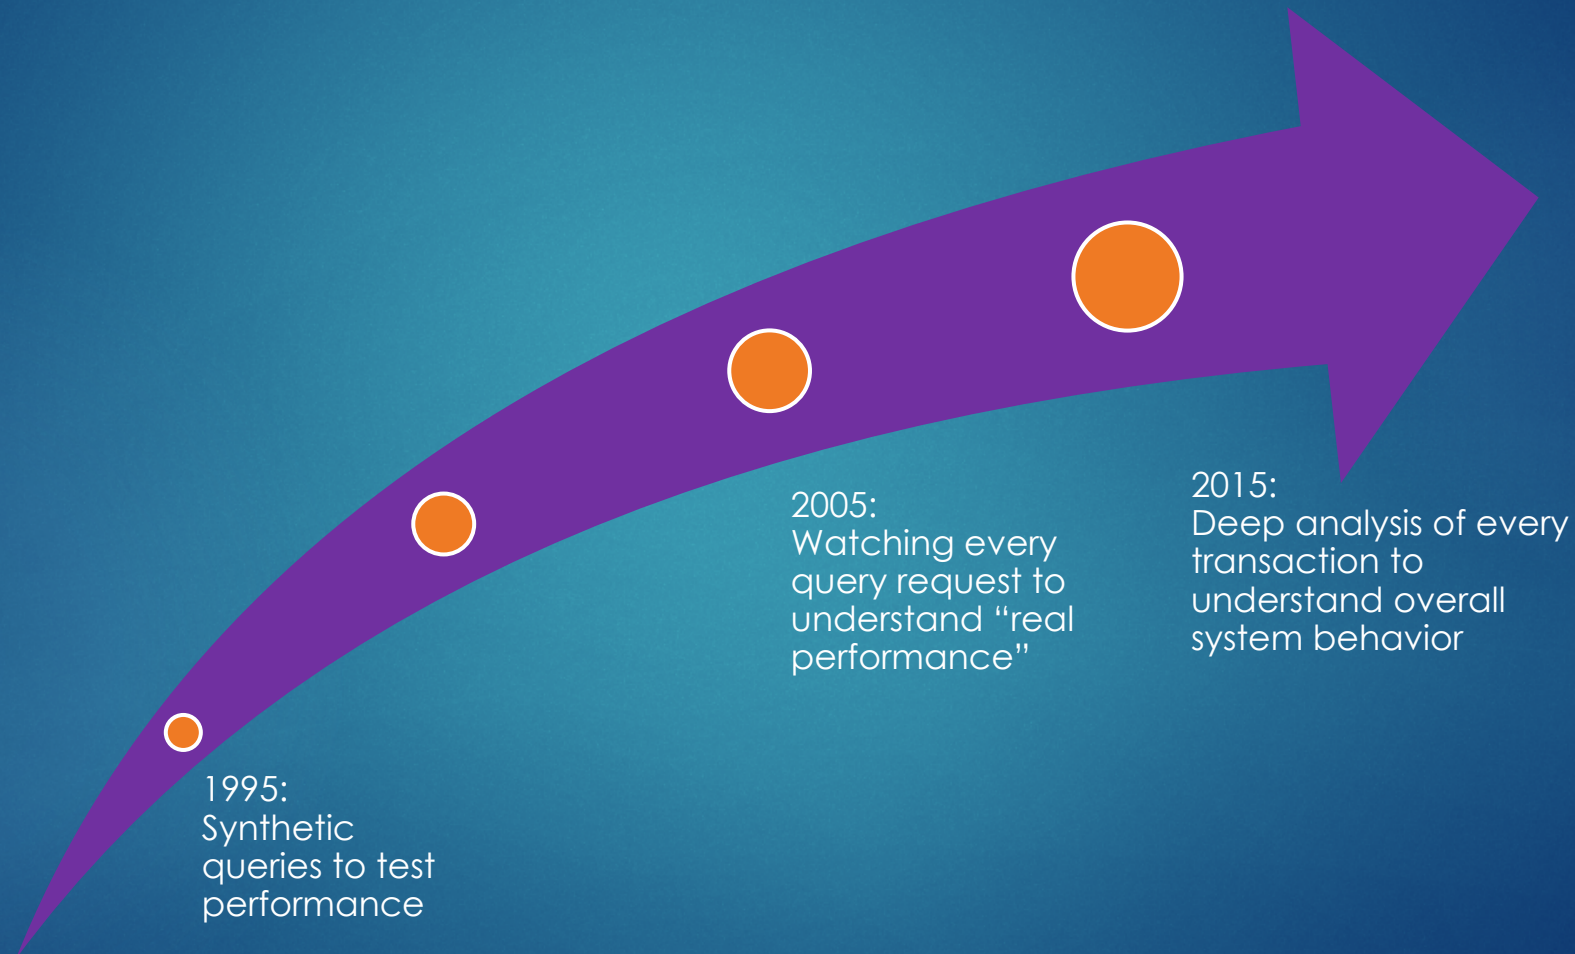
Your System Is Larger Than Your “Systems”



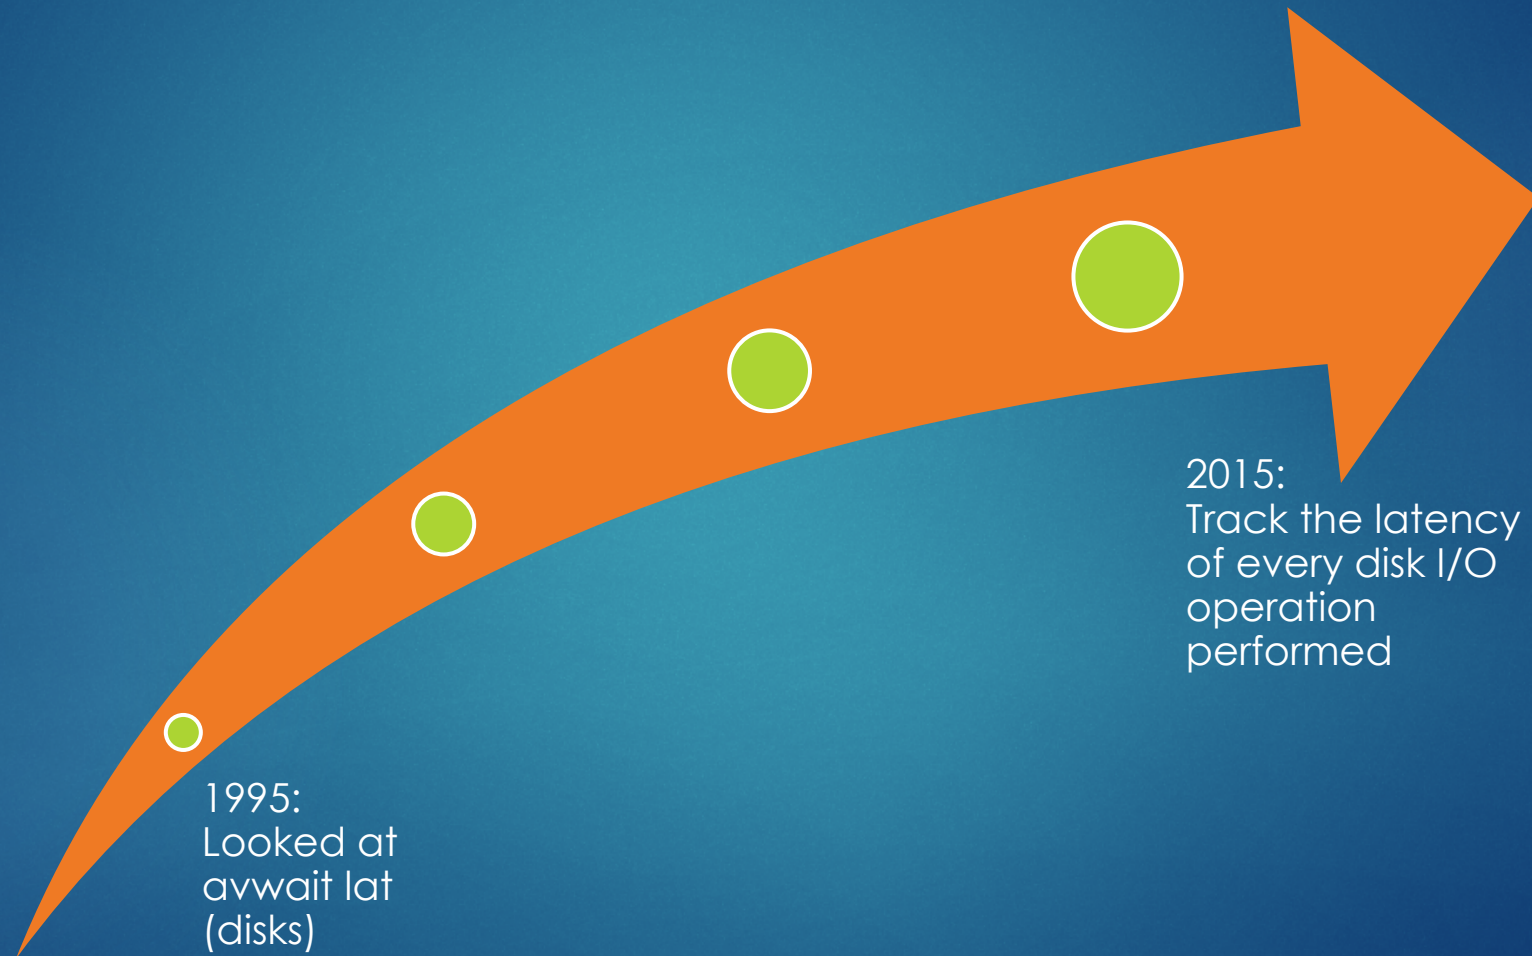
Evolution of Web Monitoring



Evolution of Database Monitoring



Evolution of Systems Monitoring



Monitoring Is Sophisticated

Increased Telemetry Volume

Advances in Time Series Databases to store trillions of samples in a billion streams.

Advances in Stream Analytics to handle velocity at scale for real-time analysis and alerting.

More Valuable Operational Questions

Data Science is the future.

Increased volume mandates computer assistance where “ops dashboards” once worked.

Most sophisticated modeling: stats, machine learning, AI, etc.

Increased Organizational Velocity

Systems are decoupled, distributed and changing faster.

Understanding overall systems behavior is like looking at sand dunes.



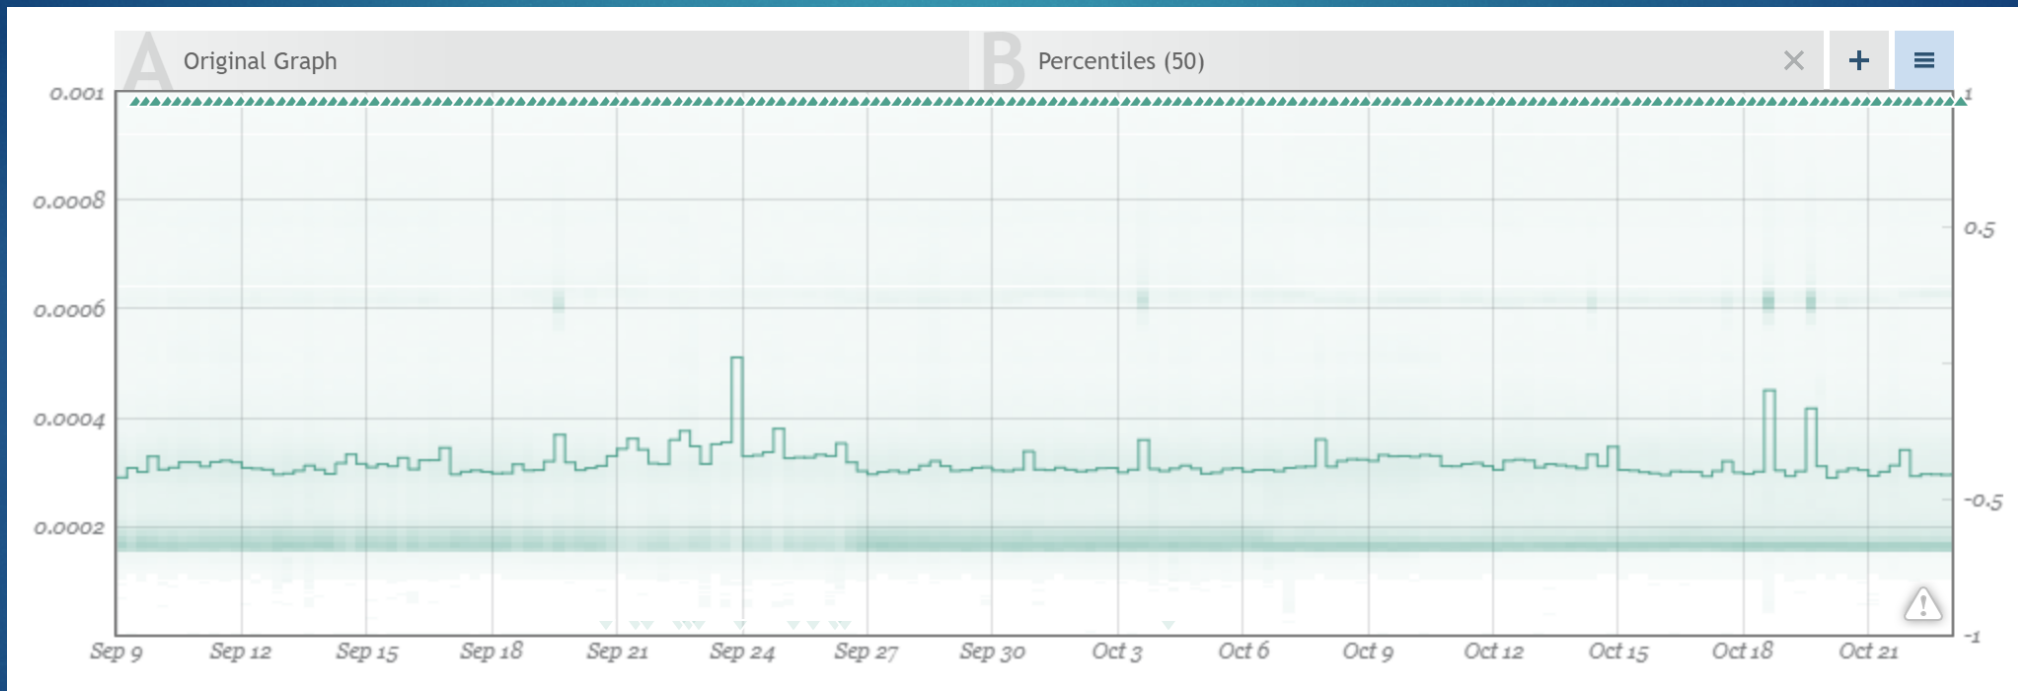
Service Level Objectives

SLOS ARE WHAT DRIVES SRES

SLO: usually based on percentiles


- ▶ E.g. 95th percentile less than 10ms
 - ▶ “simply” 95% of all samples should 10ms or less, 1% can be arbitrarily bad
- ▶ Not “simple”
 - ▶ Calculated over what period of time (or worse, number of samples)?
 - ▶ Why 95% and not 99% or 99.9% or 99.34860943%?
 - ▶ Why 10ms?
- ▶ The tragedy of the not-a-histogram histogram:
 - ▶ There are no right answers, and rarely good ones.

Median Latency Over 5m Stepping Window



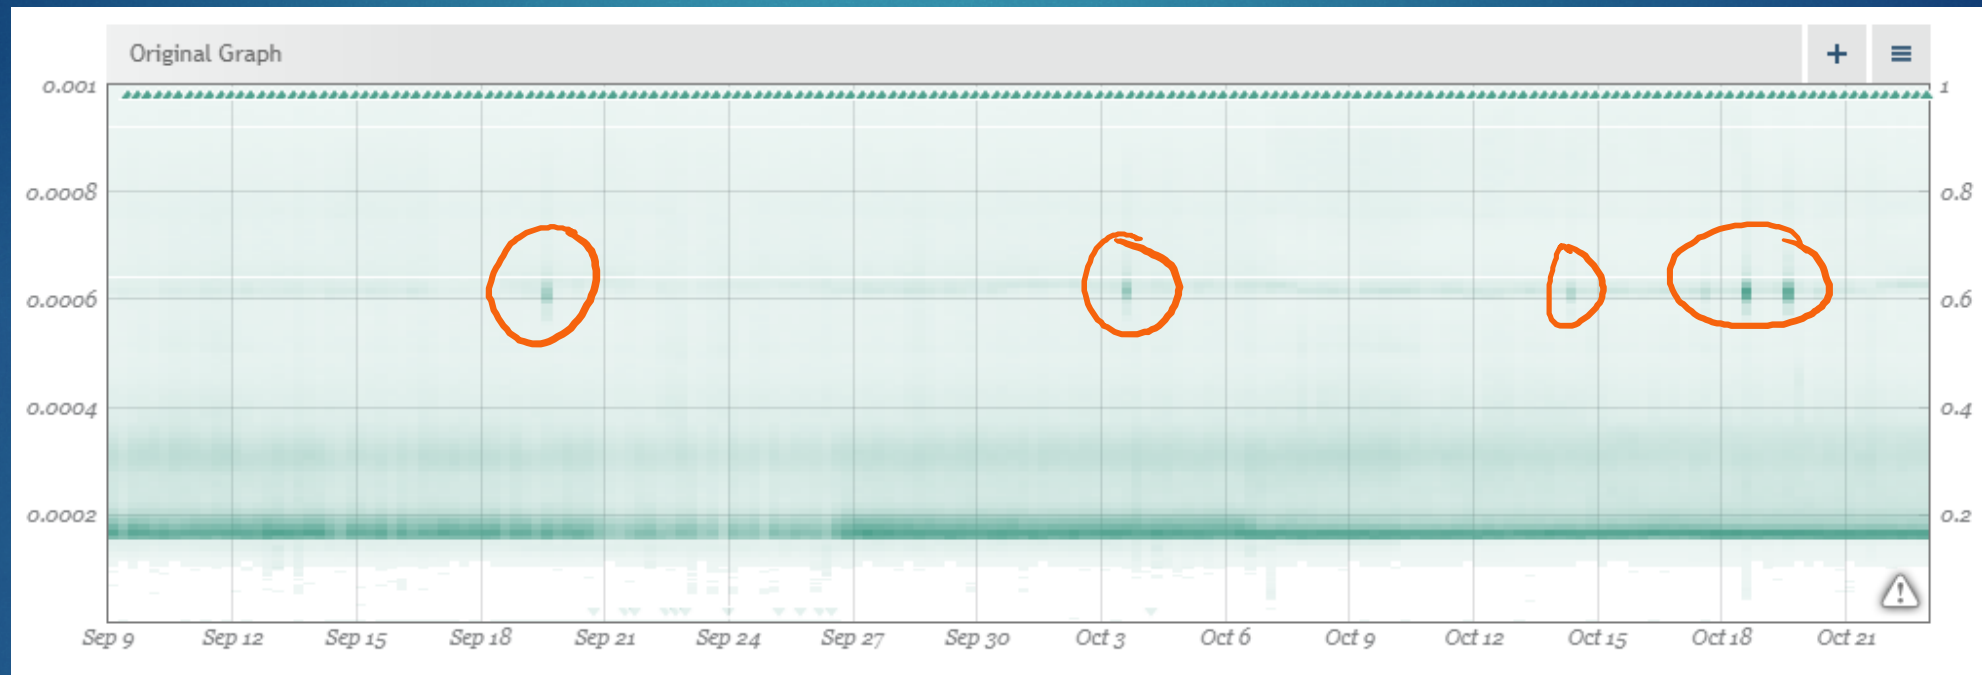
Summary Histogram

30days and 36mm samples




Time-series Histogram

30days and 36mm samples



Time-series Histogram


30days and 36mm samples




Time-series Histogram

30days and 36mm samples

Let's zoom in




Average Latency Over 5m Stepping Window




Summary Histogram

2days and 1.6mm samples




Latency Over 5m Stepping Window




Latency Over 5m Stepping Window


5m time slice
has 80k samples




p(95) Latency Over 5m Stepping Window



$p^{-1}(10\text{ms})$ Latency Over 5m Stepping Window



p^{-1} (50ms) Latency Over 5m Stepping Window





Time Matters

The time quantum you use to assess is your minimum window of failure.



Uncertainty Matters

You will certainly want to revise your goals,
likely in all parametric space.



Histograms Matter

You cannot manage percentile-based SLOs at scale without histograms.

#1

Do not measure rates.

You can derive the rate of change over time at query time.

#2



Monitor outside the tech stack.

Your tech stack would not exist without happy customers and a sales pipeline.

Monitor that which is important to the health of your organization.

#3

Do not silo data.

The behavior of the parts must be put in context.

Correlating disparate systems and even business outcomes is critical.

#4

Value observation of real work
over the measurement of synthesized work.

#5

Synthesize work to ensure function

for business critical, low-volume events.

#6

Percentiles **are not** histograms.

For robust SLO management
you need to store histograms for post-processing.

#7

History is critical;

not weeks or months, but years of detailed history.

Capacity planning, retrospectives, comparative analysis, and modelling all rely on accurate, high-fidelity history.

#8

Alerts require documentation.

No ruleset should trigger an alert without:

- human-readable explanation
- business impact description
- remediation procedure
- escalation documentation

#9



Be outside the blast radius.

The purpose of monitoring is to detect changes in behavior and assist in answering operational questions.

#10

Something is better than nothing.

Don't let perfect be the enemy of good.

You have to start somewhere.



Thank You!