



History-Based Latency Prober Tuning

Jeff Borwey (jborwey@google.com)

March 15, 2022

sre.google • twitter.com/googlesre

Jeff Borwey
SREcon22 Americas
March 15, 2022



Site Reliability Engineering

Monitoring 101: Probers!

What are Probers

- Fixed workload
- Executed with a regular cadence
- Without knowledge of internals

Prober Design Decisions





Key Design Decisions

1. What code paths should be exercised?
2. How frequently should it be run?
 - a. Cost/Detection Latency trade off
3. **How do you alert on this signal?**
 - a. **What is an acceptable latency?**
 - i. Does this change over time?
 - ii. Does this change per environment?
 - b. **How do we control false-positive rate?**

How is it done in practice?

1. Identify an area of the system with low visibility
 - a. New feature
 - b. Reoccurring source of regressions/outages
2. Craft a specific workload that exercises the relevant code paths
3. Throw it into your framework of choice
4. Use “domain expertise” to pick a threshold
5. Get paged too much...
6. Increase the threshold

Now scale that to 100's of environments



Better Probing Through History!

A Better Solution: Look to History!

Alerting thresholds via **outlier detection**:

- Model historical performance
- Derive threshold for Out-Of-Distribution samples
- Adjust windowing
- Backtest!

Parametric approach


- Model latency distribution per environment
- Pick a high percentile of the distribution (99.5%, 99.95%, etc)
- Log-normal often works well for latencies
- Requires trimming outliers before fitting

$$\log(X - l) \sim N(\mu, \sigma^2)$$

$$X \sim \text{Lognormal}(\mu, \sigma^2, l)$$

μ is the mean, σ^2 is the variance, l is the location offset

Prober Latency Histogram w/ Thresholds



Non-parametric approach

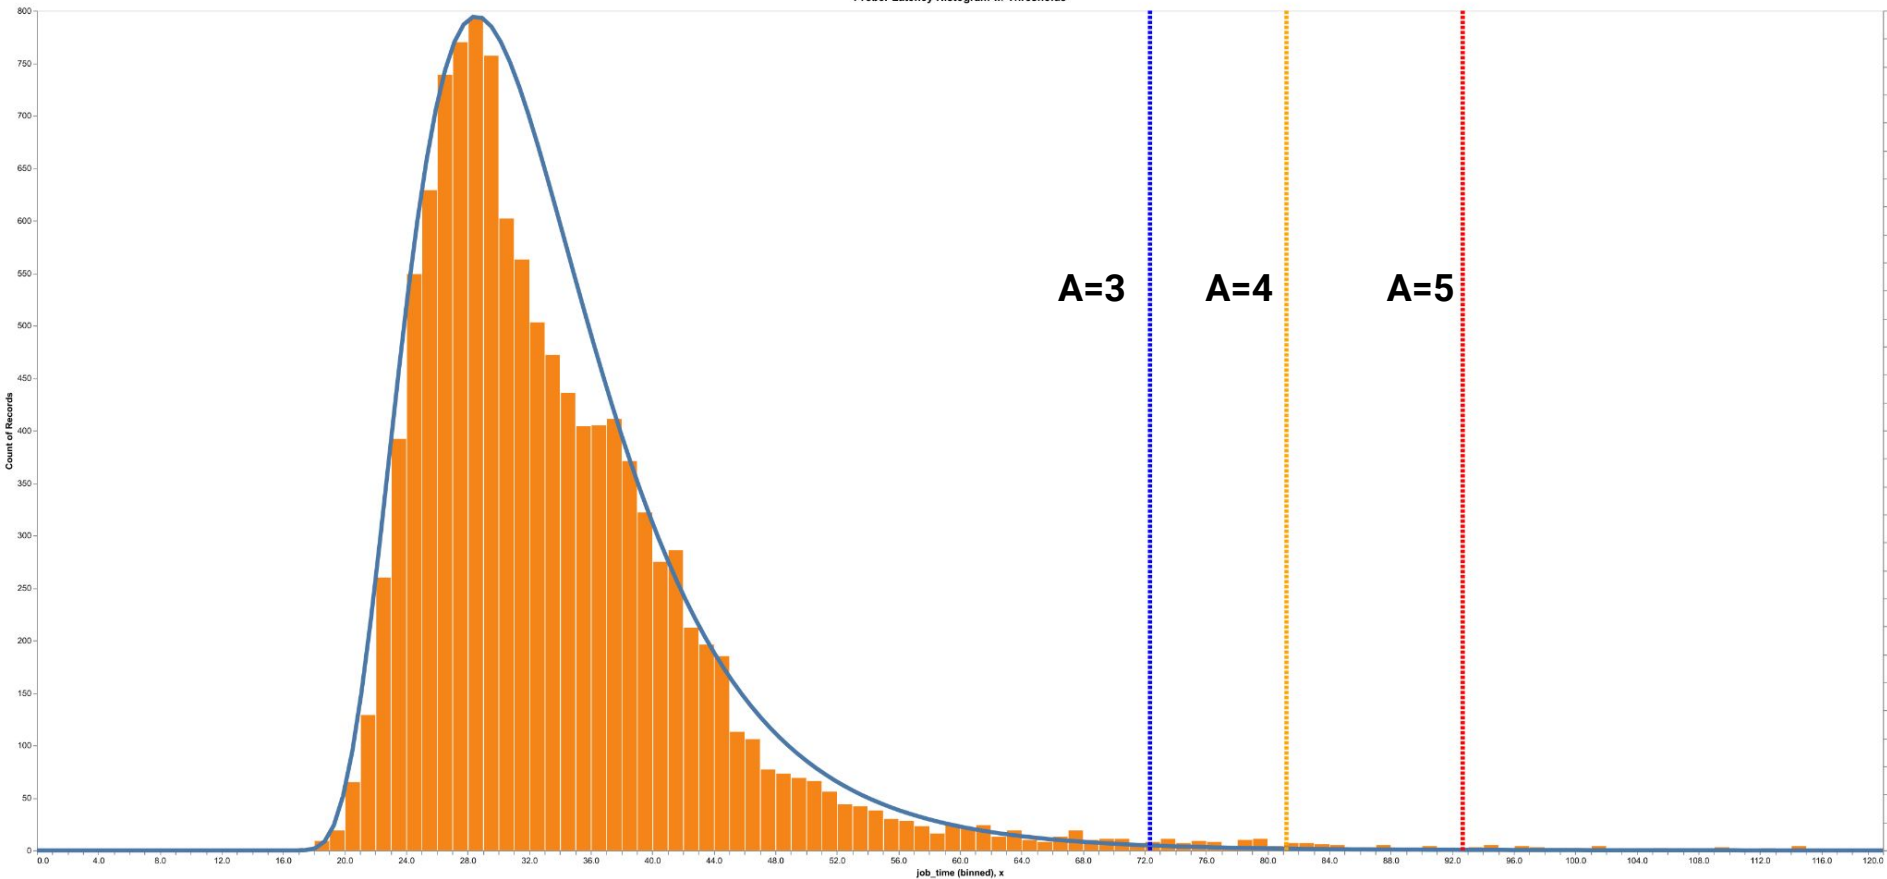
- Compute **Interquartile Range (IQR)** (75%-25%)
- Compute **skewness** using **medcouple (MC)**
- Skew-adjusted box-plot [1]

$$threshold = Q_3 + 1.5 * e^{A*MC} * IQR$$

Q₃ is the 75%, **A** is a tuneable parameter


[1] Hubert, M.; Vandervieren, E. (2008). “An adjusted boxplot for skewed distribution”. *Computational Statistics and Data Analysis*. **52** (12): 5186–5201.

Prober Latency Histogram w/ Thresholds




From Samples To Alerts


Windowing



Windowing (Average)




Windowing (Median)

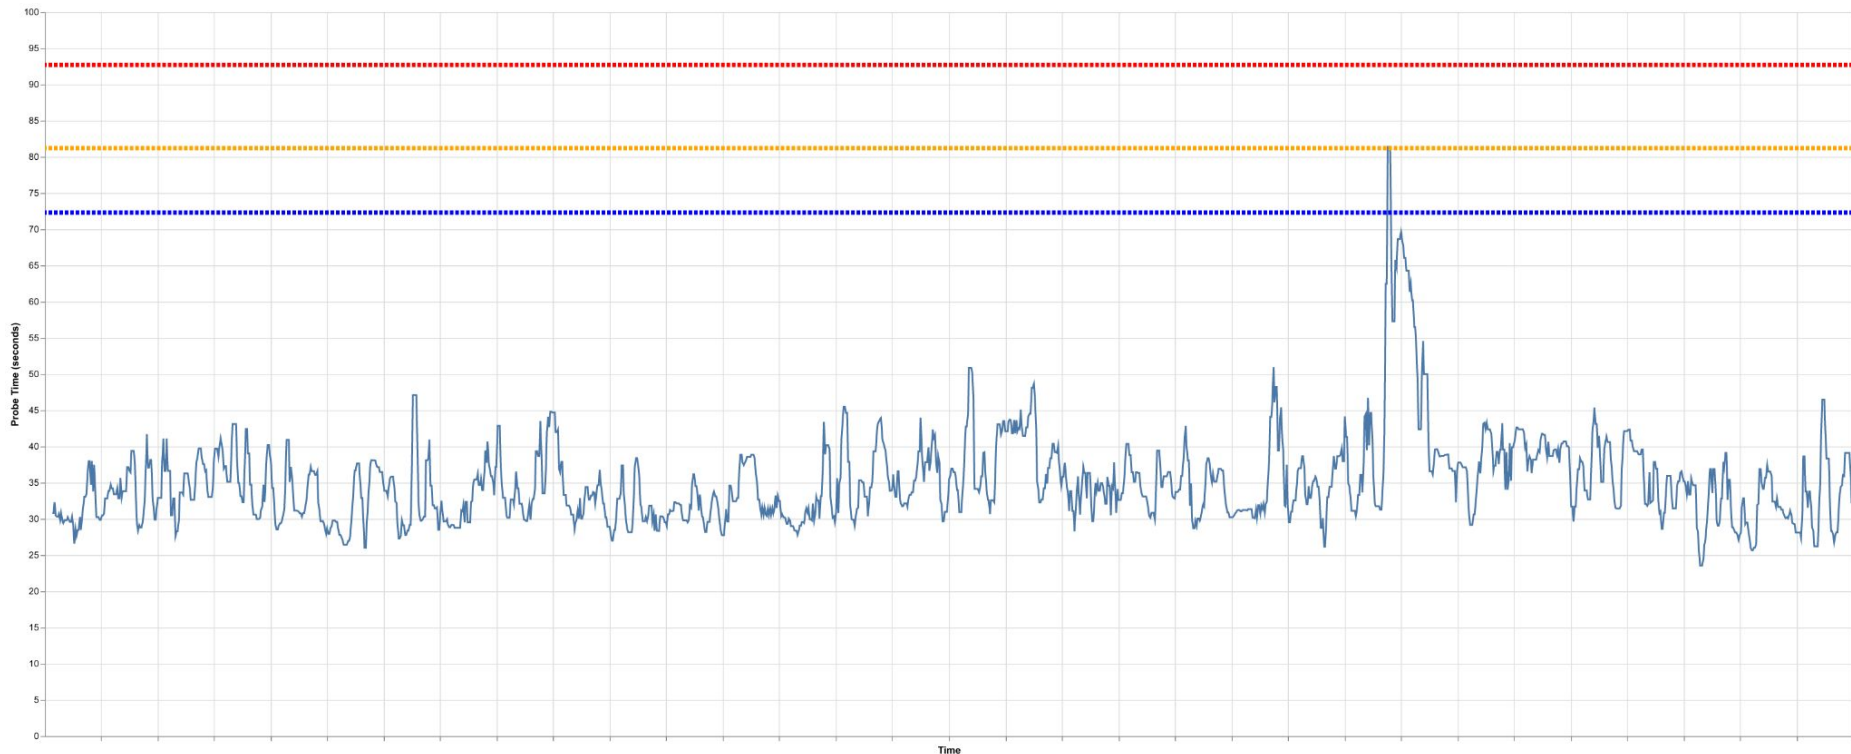


Jeff Borwey
SREcon22 Americas
March 15, 2022

Prober Latency Histogram w/ Thresholds




An Alert?




Jeff Borwey
SREcon22 Americas
March 15, 2022

Another Example




Jeff Borwey
SREcon22 Americas
March 15, 2022

Another Example




An Alert!



Jeff Borwey
SREcon22 Americas
March 15, 2022

An Alert!



Jeff Borwey
SREcon22 Americas
March 15, 2022

Conclusions

Conclusions

- Not all environments are the same!
- Probers shouldn't be pets
- Using simple statistical techniques lets us:
 - Get more out of our existing probers
 - Eliminate toil
 - Still provide knobs for tuning