ASAPQuery: 100x Drop-in Acceleration for Metrics Observability

Milind Srivastava
PhD student @ Carnegie Mellon University

github.com/ProjectASAP/ASAPQuery

May 22, 2026

Slow Grafana dashboards — looks familiar?

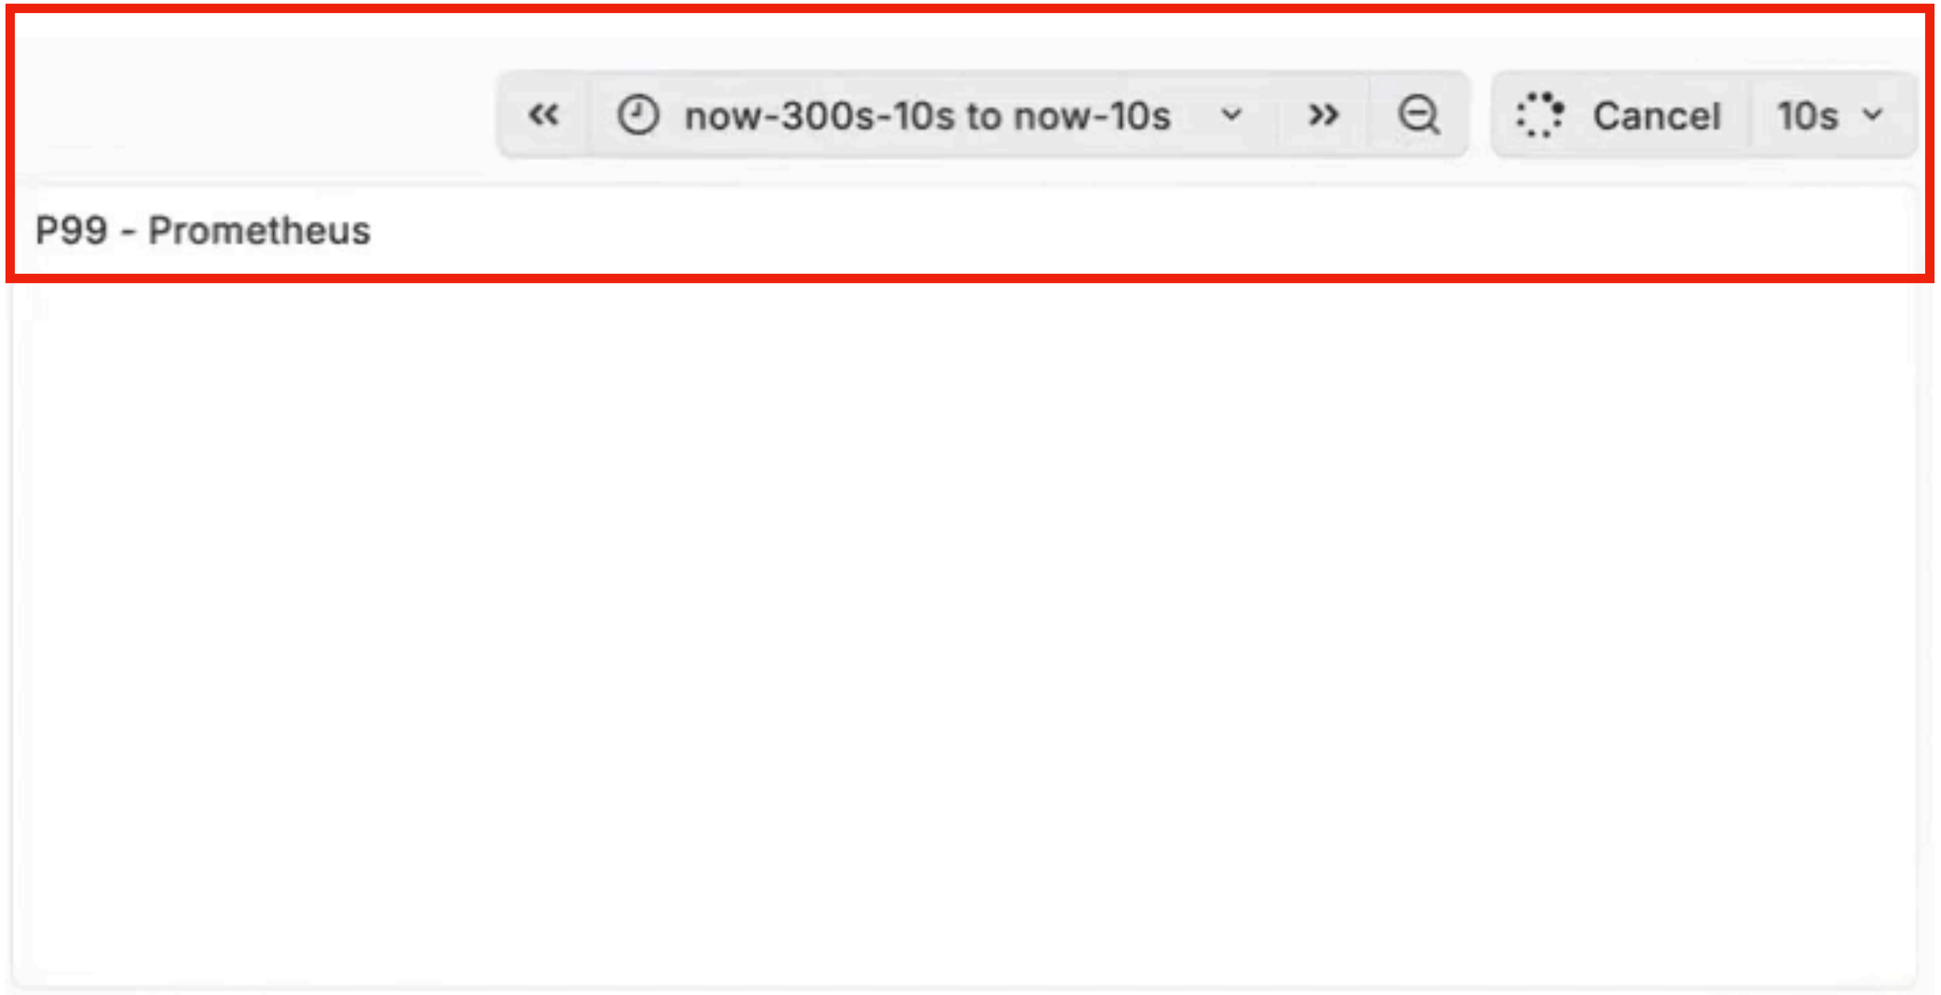

What if we could have this?

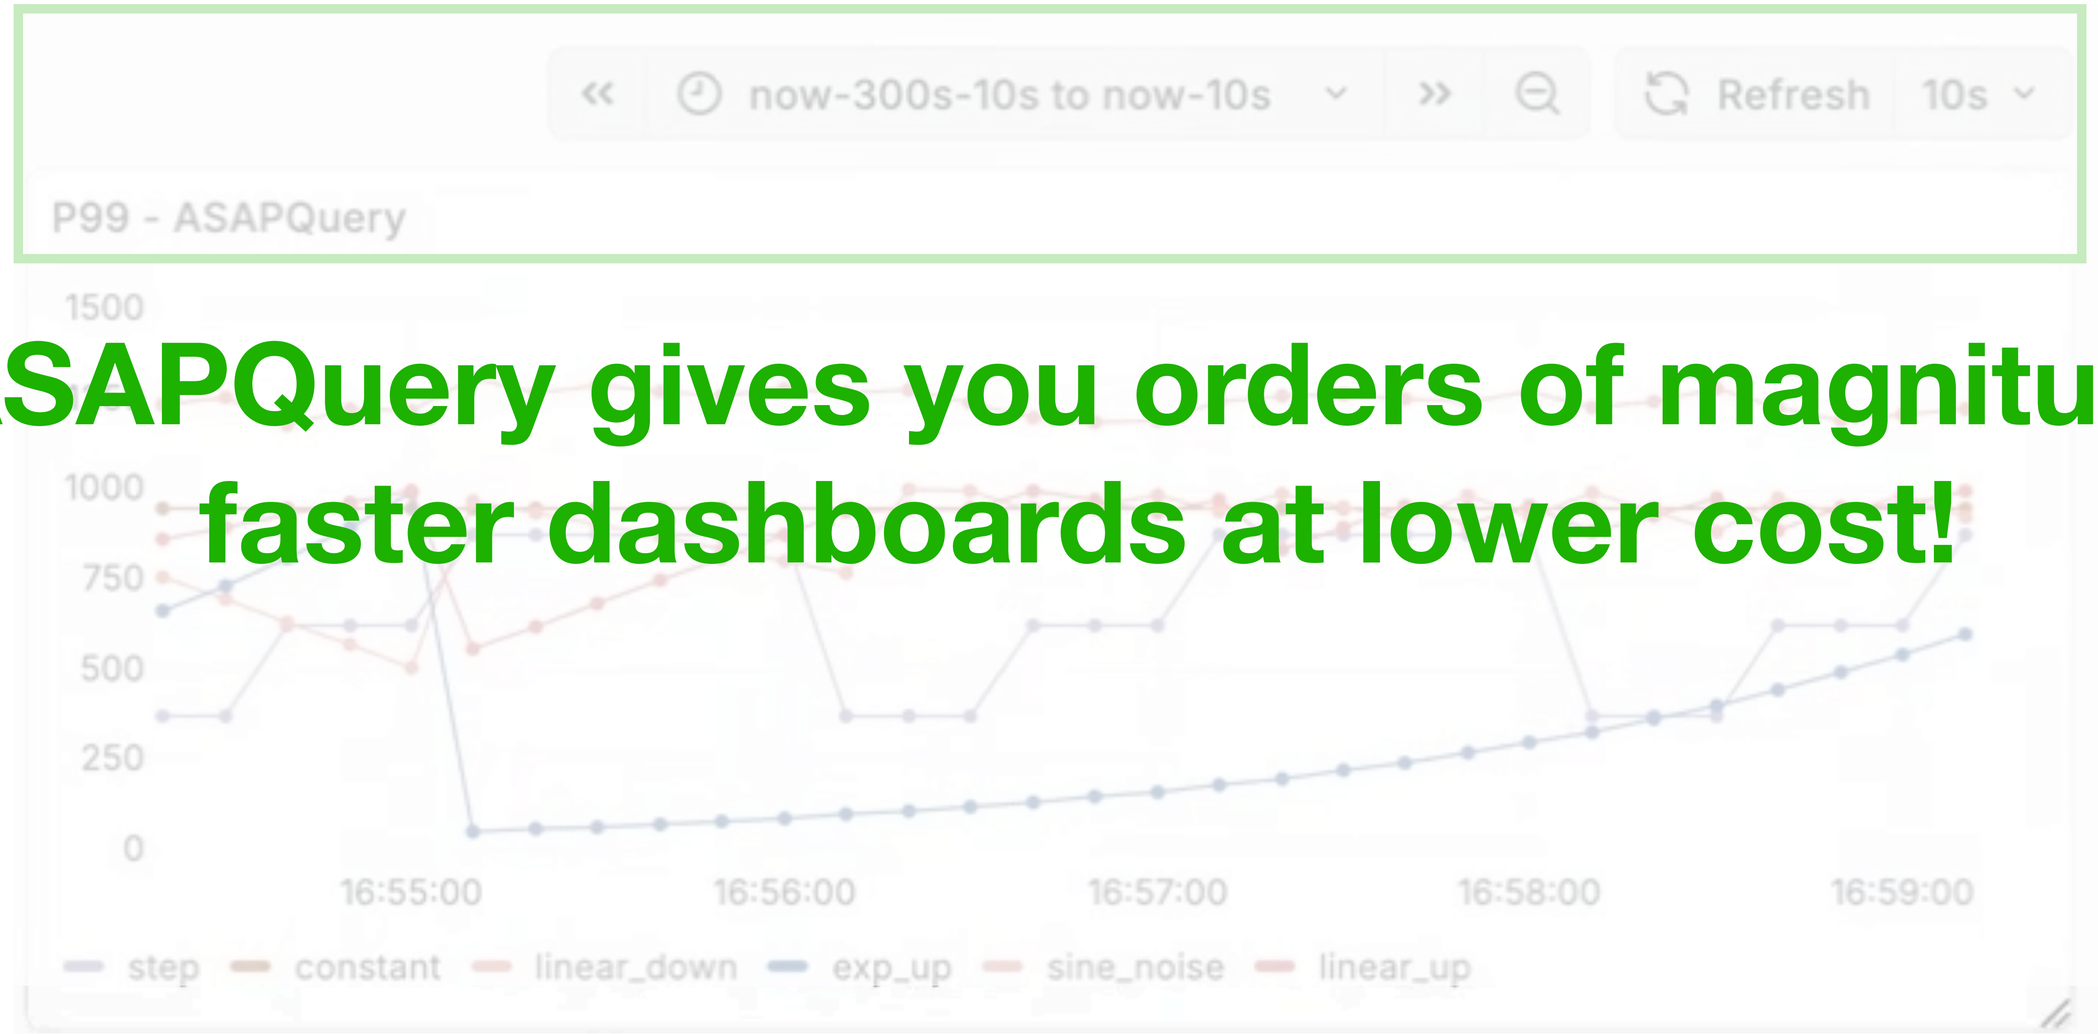

ASAPQuery gives you orders of magnitude faster dashboards at lower cost!

Who am I

Milind Srivastava

- PhD student @ Carnegie Mellon University
- Lead for ProjectASAP: projectasap.github.io
- We build systems to make observability cheaper and faster
- Our research democratizes the benefit of cool algorithms for practitioners like you

ASAPQuery: A drop-in accelerator for Prometheus-Grafana

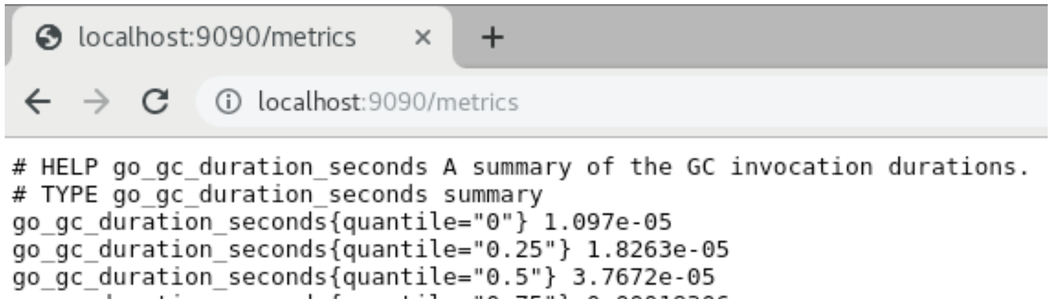

Query

exporter-1

exporter-2

exporter-3

Pull metrics

Prometheus

ASAPQuery: A drop-in accelerator for Prometheus-Grafana

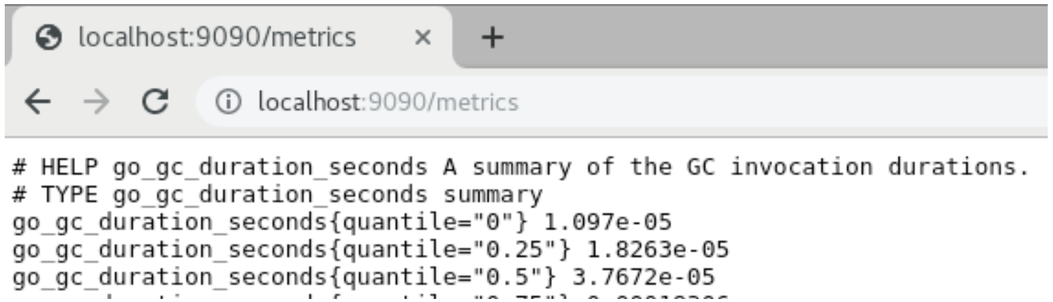

Query

remote_write
to ASAPQuery

ASAPQuery

Fallback

exporter-1

exporter-2

exporter-3

Pull metrics

Prometheus

The secret sauce: Sketches

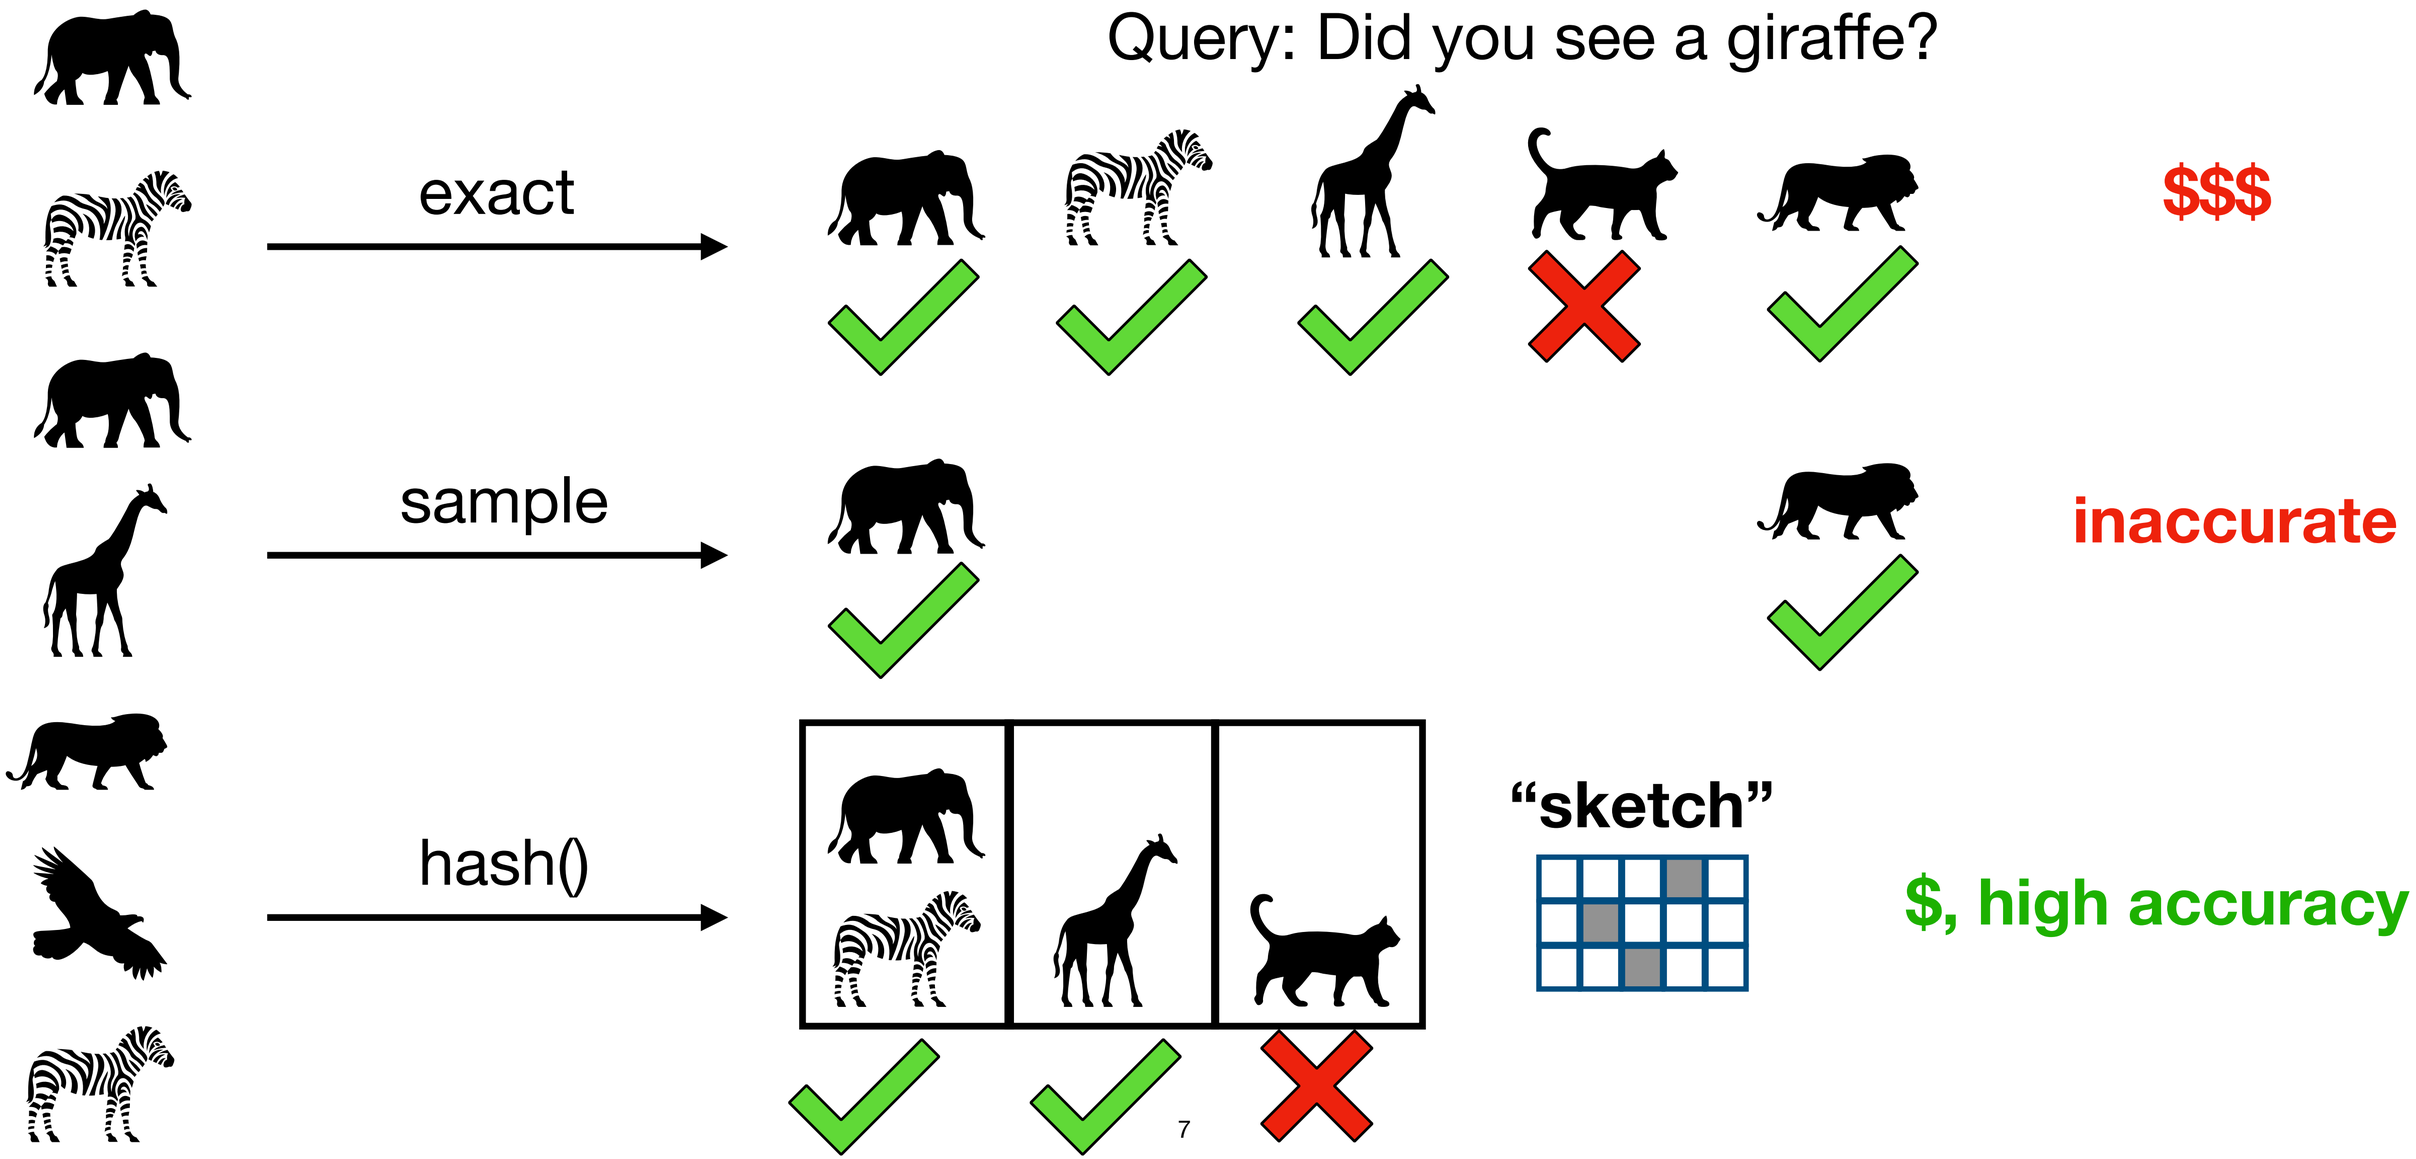

Sketches can answer many queries

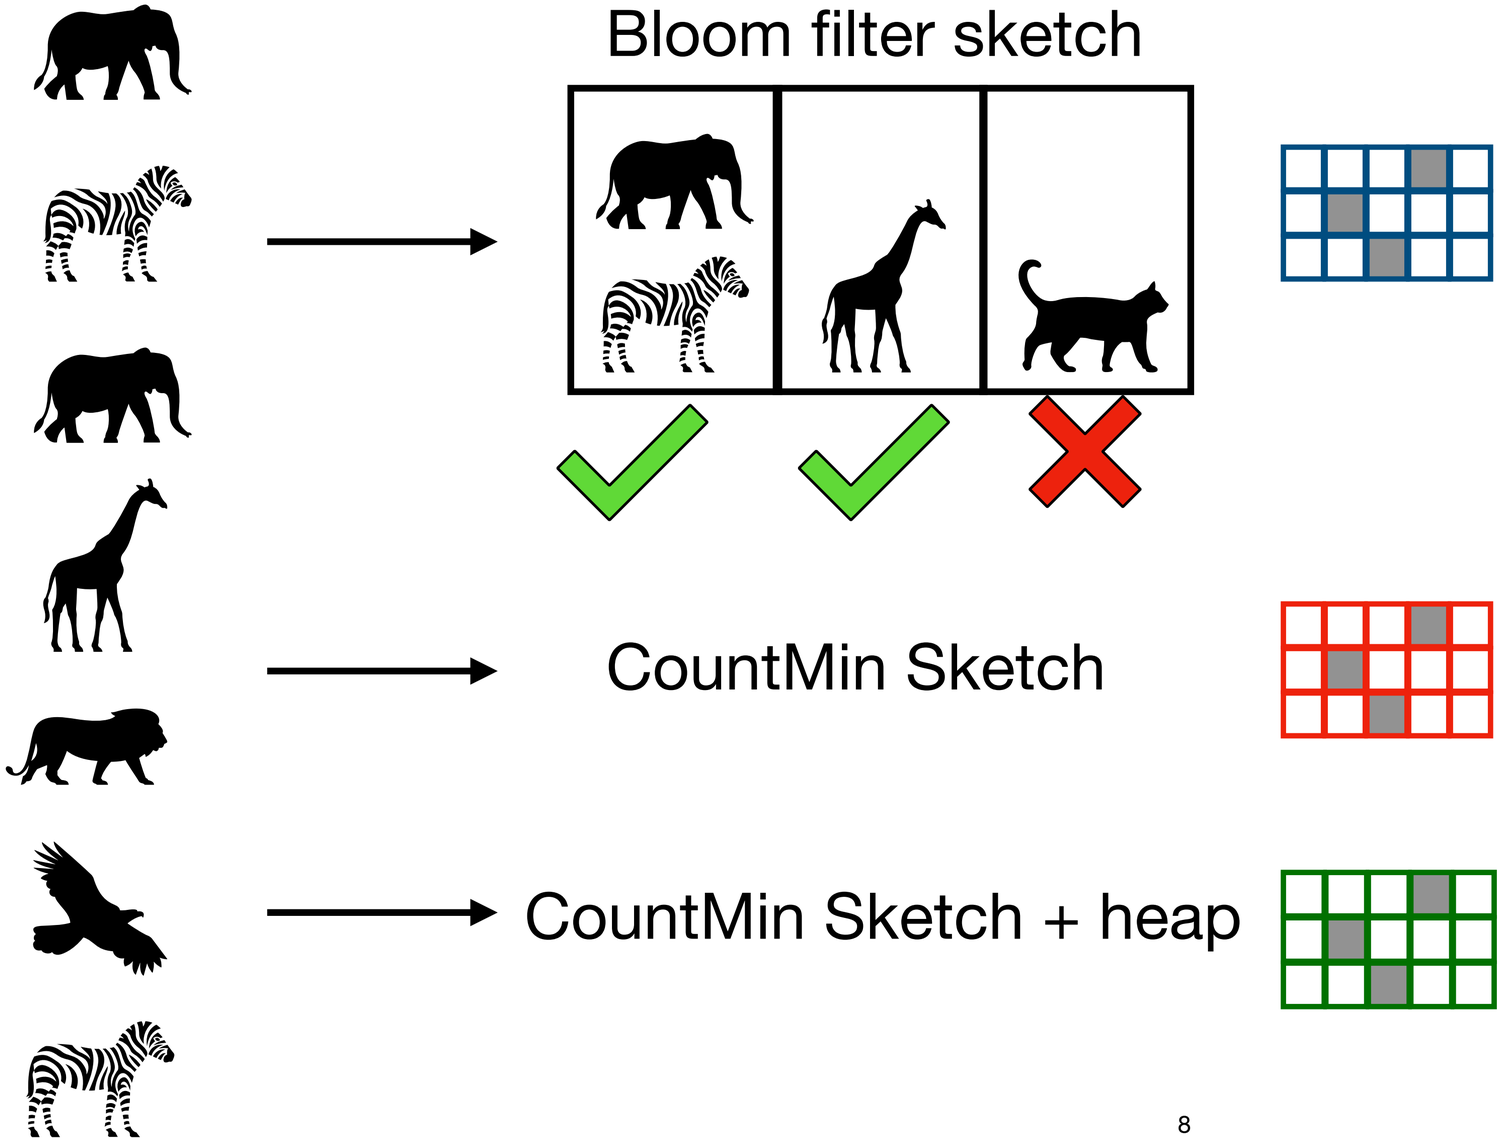

Queries on animals

Did I see a giraffe?

How many times did I see a lion?

What are the top 5 animals I saw?

From animals to metrics

Metric: CPU usage

PromQL: Queries on metrics

h1, login, 5%

h1, web, 20%

h2, proxy, 30%

h3, load-balancer, 50%

h1, web, 30%

...

...

...

→ Bloom filter Sketch

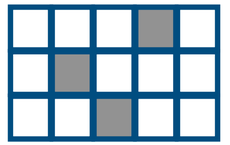

Is host X up?

→ CountMin Sketch

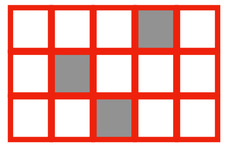

Avg CPU usage across services

→ CountMin Sketch + heap

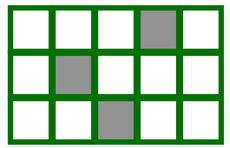

Top 5 hosts with highest CPU

→ KLL Sketch

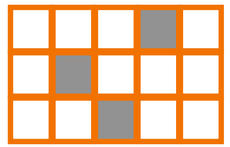

p90 CPU usage across hosts

ASAPQuery makes sketches easy to use for observability

Metrics observability

Grafana

Prometheus

Faster dashboards

Lower CPU & memory
==
Lower $$$

ASAPQuery

Sketches

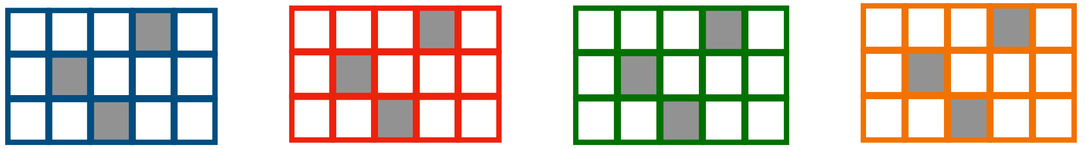

Under the hood of ASAPQuery

1. Analyze query patterns
2. Ingest metrics into sketches
3. Answer queries with sketches

Query

ASAPQuery

asap-planner

Query plan

asap-sketch-ingest

asap-query-engine

`remote_write` to ASAPQuery

Pull metrics

Prometheus

Live demo of ASAPQuery

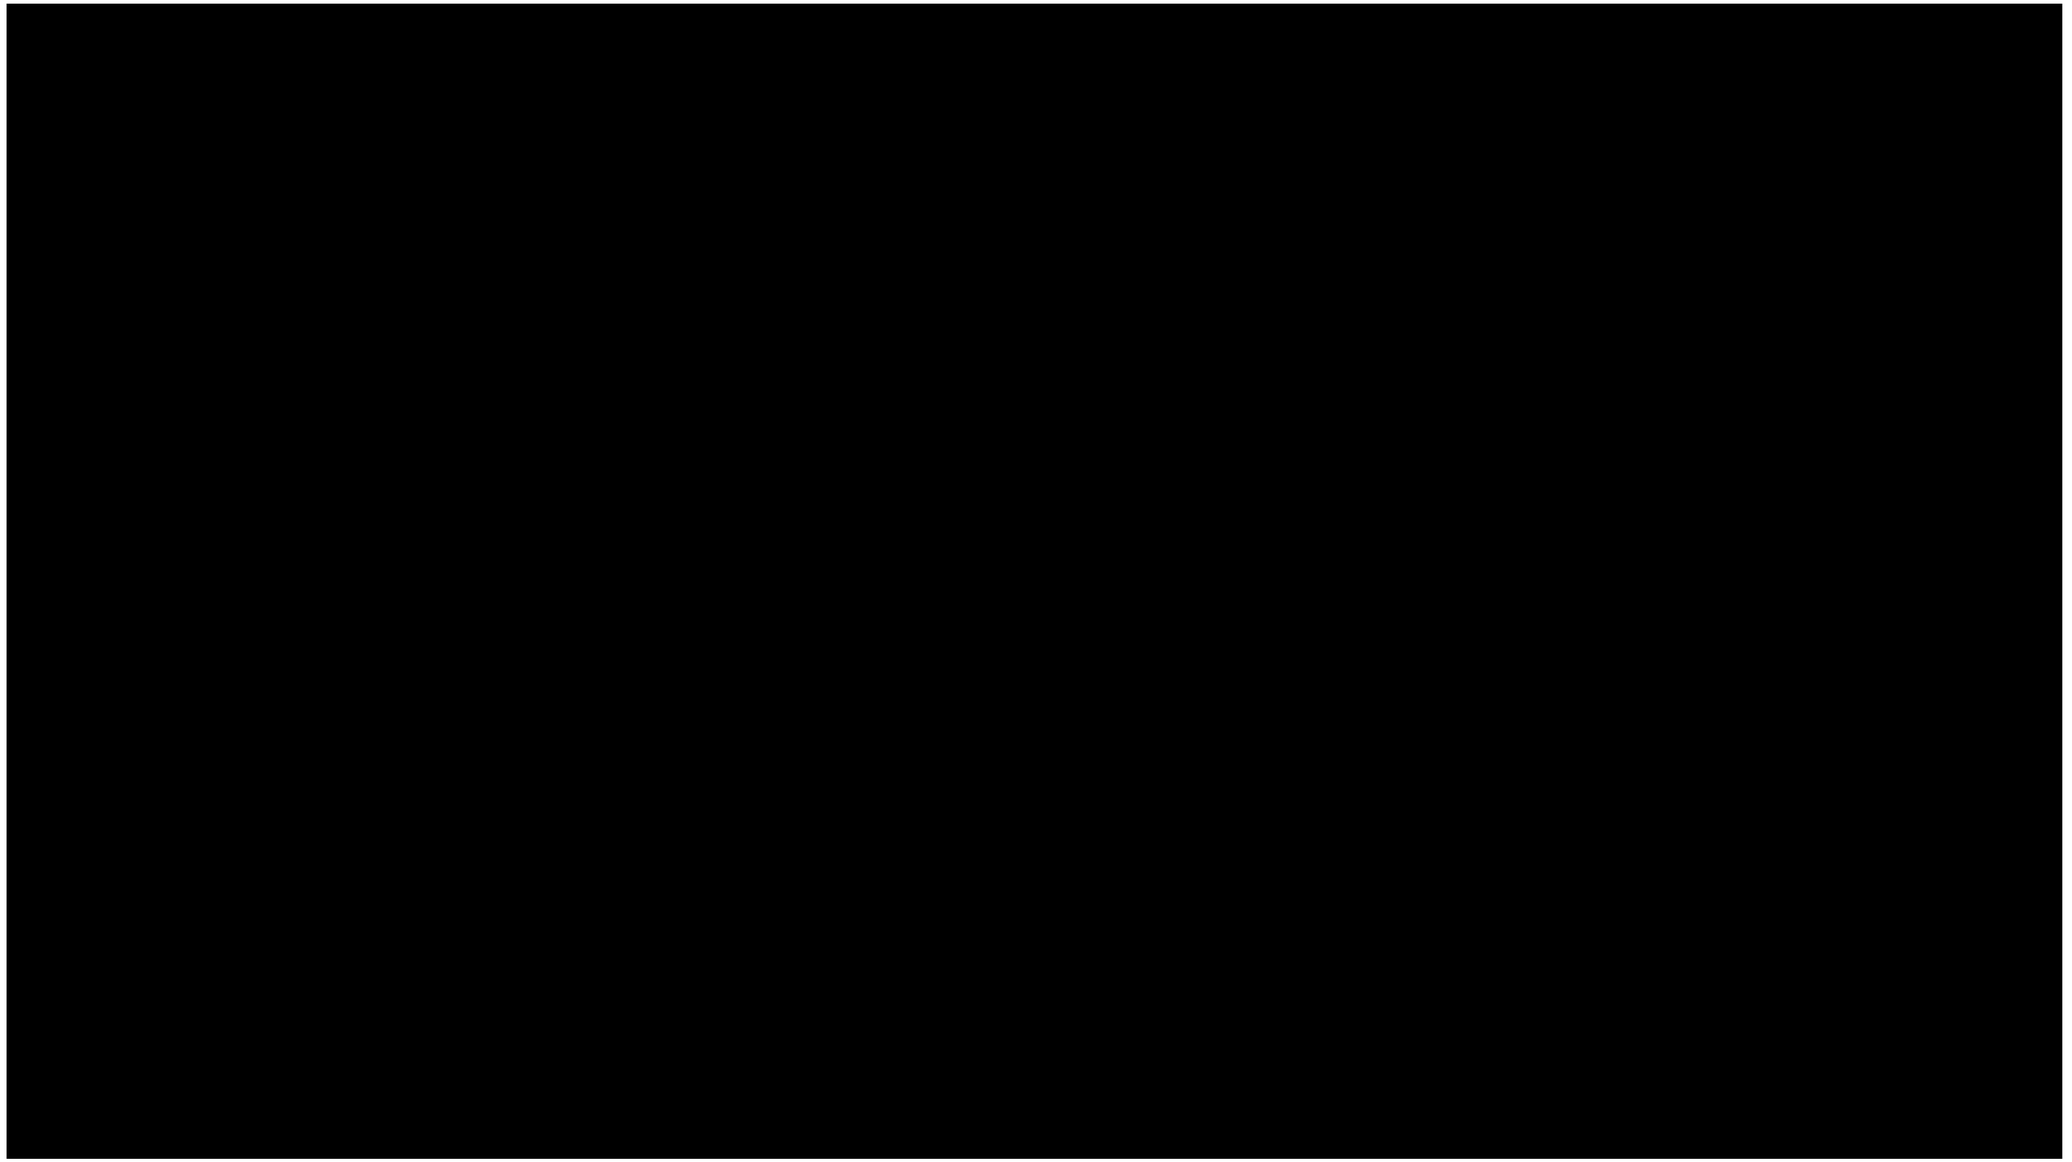

Try out ASAPQuery and share your thoughts!

If this is useful, I would appreciate a star on the repo :)

github.com/ProjectASAP/ASAPQuery

- ASAPQuery modes (instructions in Github repo)
  - quickstart: 5 min to see benefits
  - drop-in: works with your Prometheus-Grafana stack
- Upcoming blog posts:
  - Library of sketches used in ASAPQuery
  - Architecture of ASAPQuery
- Email me for help or feedback: milindsr@andrew.cmu.edu

projectasap.github.io