Bomb Squad

Containing the Cardinality Explosion

Cody Boggs

PromCon 2018, Munich

Who?

Cody “Chowny” Boggs

- Ops Nerd of ~8 years
- Lead yak shaver @ FreshTracks for ~1 year
- Obsessed with metrics
- Pretends to write code
- Breaks things. (*All* the things.)

cody@freshtracks.io

On Deck

- What is cardinality?
- What is a “cardinality explosion”?
- Who cares?
- Charts and graphs!
- Bomb Squad live demo!

What is cardinality...

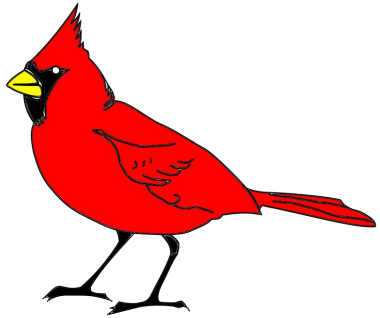

FreshTracks.io

Generally?

The number of elements in a set or group

{ b, 42, tree}

Cardinality: **3**

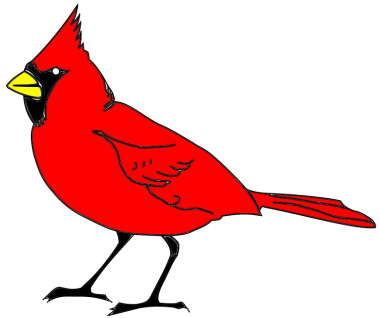

Images: https://openclipart.org/detail/133471/cardinal-remix-1

What is cardinality...

FreshTracks.io

For this talk?

The number of discrete label/value pairs (series) associated with a particular metric

cpu{host=”foo”}

cpu{host=”bar”}

cpu{host=”broken”}

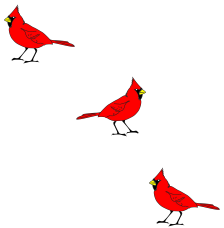

Cardinality: **3**

Images: https://openclipart.org/detail/133471/cardinal-remix-1

Words that mean things!

FreshTracks.io

Series: A discrete set of label name / value pairs containing one or more timestamped data points

Metric: A group of series sharing a “__name__” label value, eg: “api_requests”

Cardinality Explosion / High Card. Event: Sharp increase in series creation rate

Exploding Label: A label whose count of distinct values is disproportionately high compared to other labels within a metric’s series

So… Explosions?

FreshTracks.io

Rapid inflation of the number of series under one or more metrics

Examples of Causes

- Prolonged extreme pod turnover rates
- Highly elastic workloads with fine-to-medium grain labels
- Bad code deploy that sticks unique IDs, timestamps, or the like into a label value
  - This one seems to be the most common cause
  - Magnitude tends to be huge

https://openclipart.org/detail/298678/bomb-2

Why do I care?

FreshTracks.io

Areas of concern:

1. **Meaningfulness of affected data**
   a. Single “legitimate” data point per series, inability to aggregate on “exploding” labels

2. **Stability and responsiveness of Prometheus proper**
   a. Query times, memory usage, scrape durations, remote_write queue, etc.

3. **Stability of downstream receiving services**
   a. Cortex (remote write); BigTable, DynamoDB, Thanos (chunk stores); etc.

https://openclipart.org/detail/196149/fireball
https://upload.wikimedia.org/wikipedia/en/3/38/Prometheus_software_logo.svg

Impact on Prometheus

FreshTracks.io

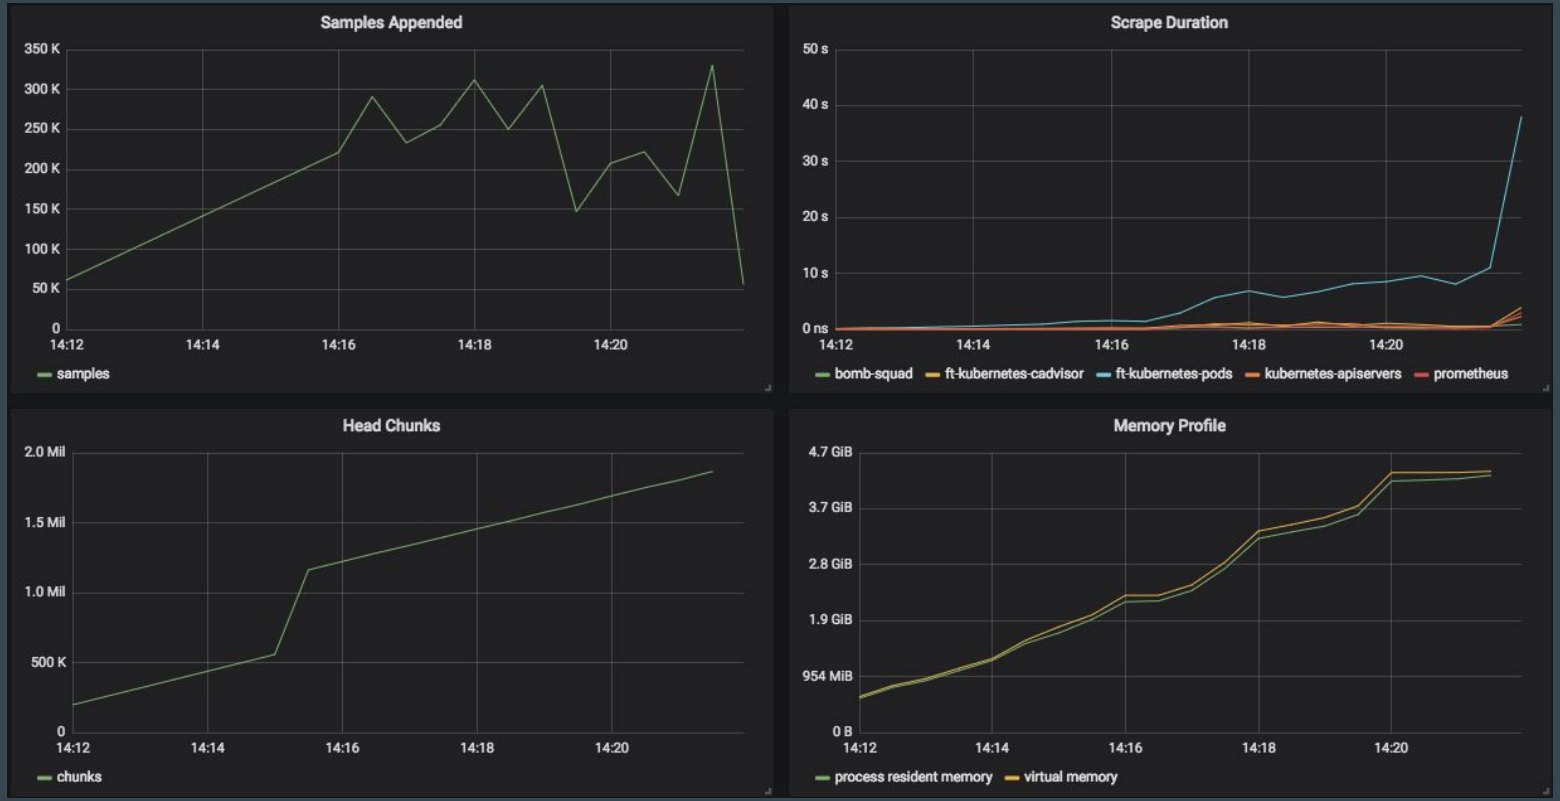

@strofcon

Impact on Cortex & BigTable

FreshTracks.io

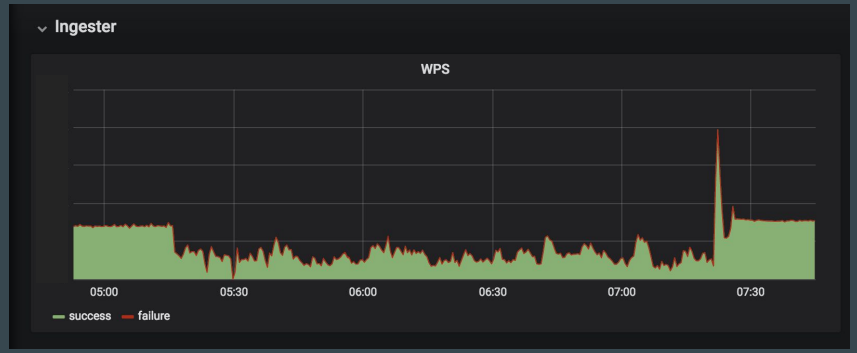

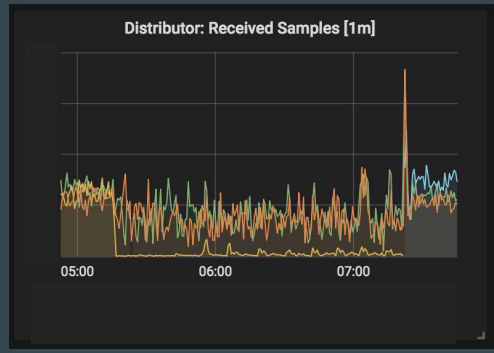

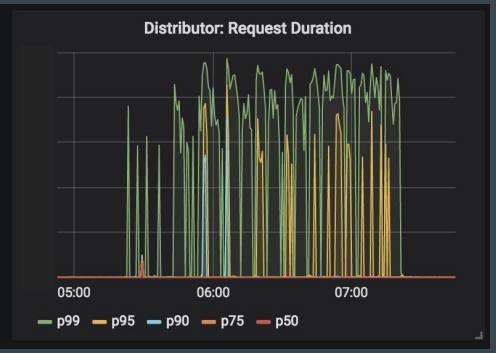

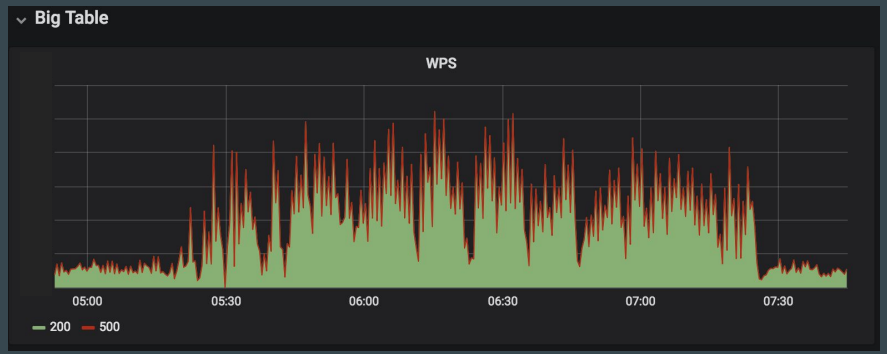

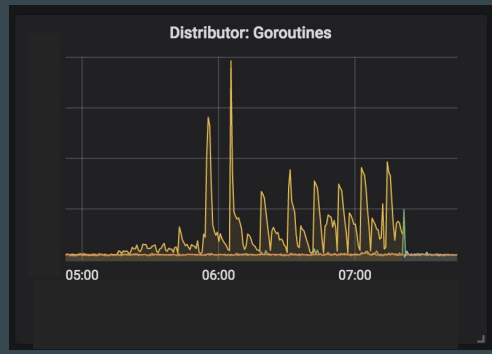

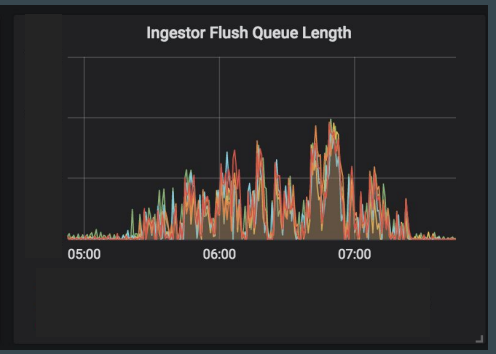

@strofcon

Who ya gonna call? Bomb Squad!

Overview:

1. Run as sidecar to Prometheus proper
2. Bootstrap recording rules into Prometheus
3. Monitor for exploding metrics
4. When found, identify exploding label
5. Insert “silencing rule” relabel config(s)

…

n. CLI commands available to list and unsilence metrics

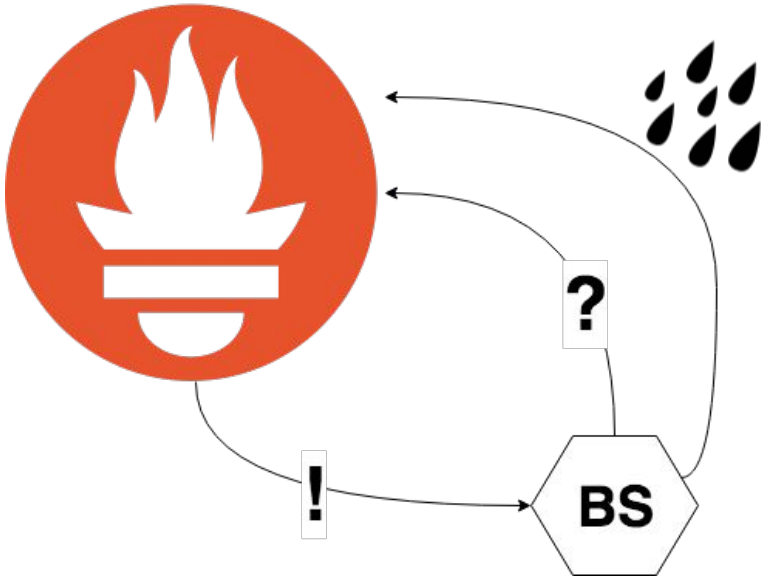

https://upload.wikimedia.org/wikipedia/en/3/38/Prometheus_software_logo.svg

In which Cody attempts a live demo...

Thanks

github.com/Fresh-Tracks/bomb-squad

cody@freshtracks.io

@strofcon