



# Alert on all the things

Integrating Quicksilver with Prometheus

Lorenz Bauer ([imb@cloudflare.com](mailto:imb@cloudflare.com)), 2017-08-18, PromCon 2017

# Cloudflare

- Web performance / security as a Service
- Global CDN
  - 4.8m HTTP req/s avg
  - 1.2m DNS queries/s avg
- [blog.cloudflare.com](https://blog.cloudflare.com)

# Quicksilver

- Transfers customer configuration from central location to the “Edge”
- “Replicated” database
- Runs on a few thousand machines
- Team: Geoffrey Plouviez, Sami Kerola, me

# Quicksilver

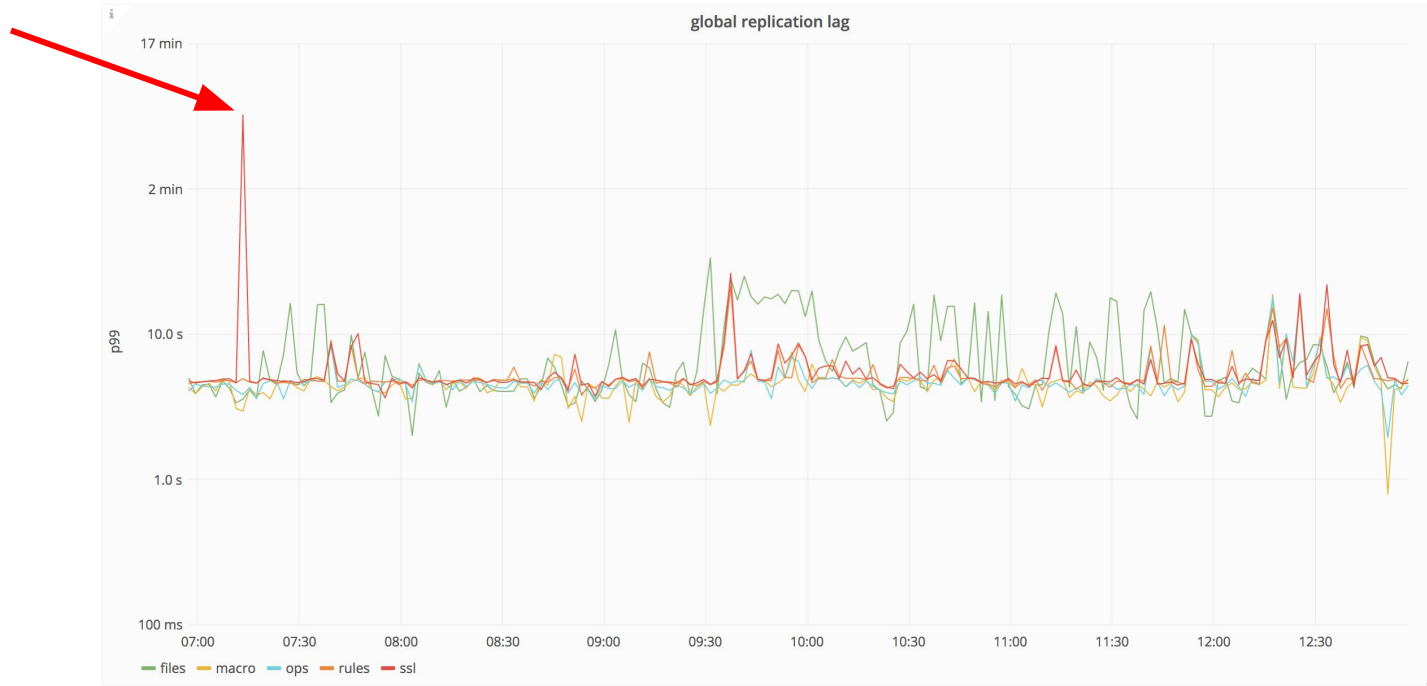


Leaflet | © OpenStreetMap © CartoDB

# Replication lag



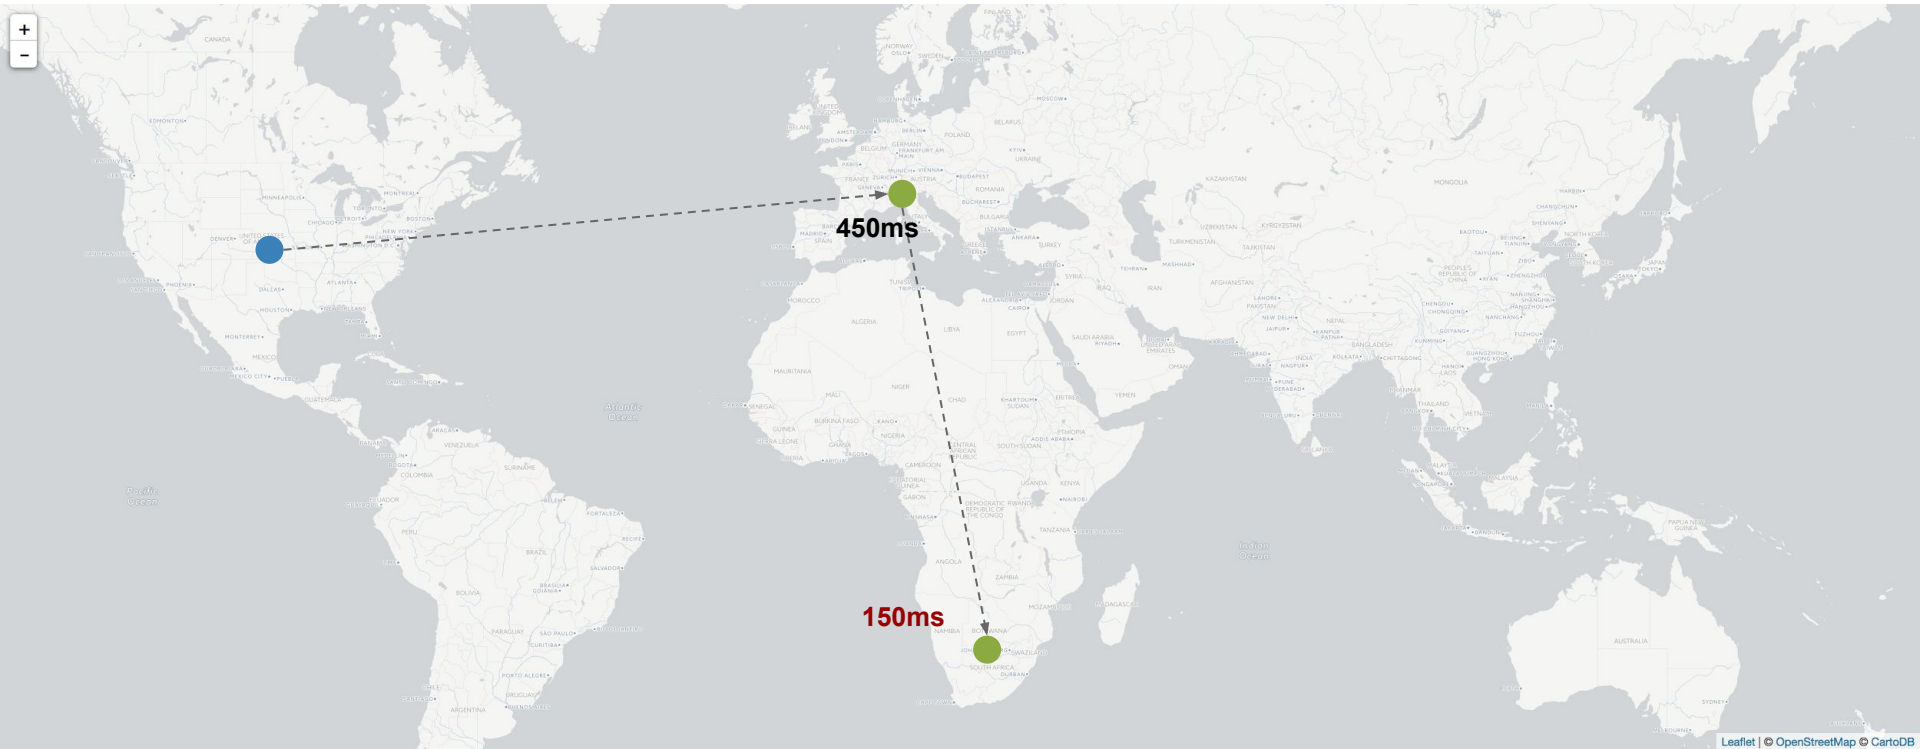
# Replication lag



# In reality, we lack information

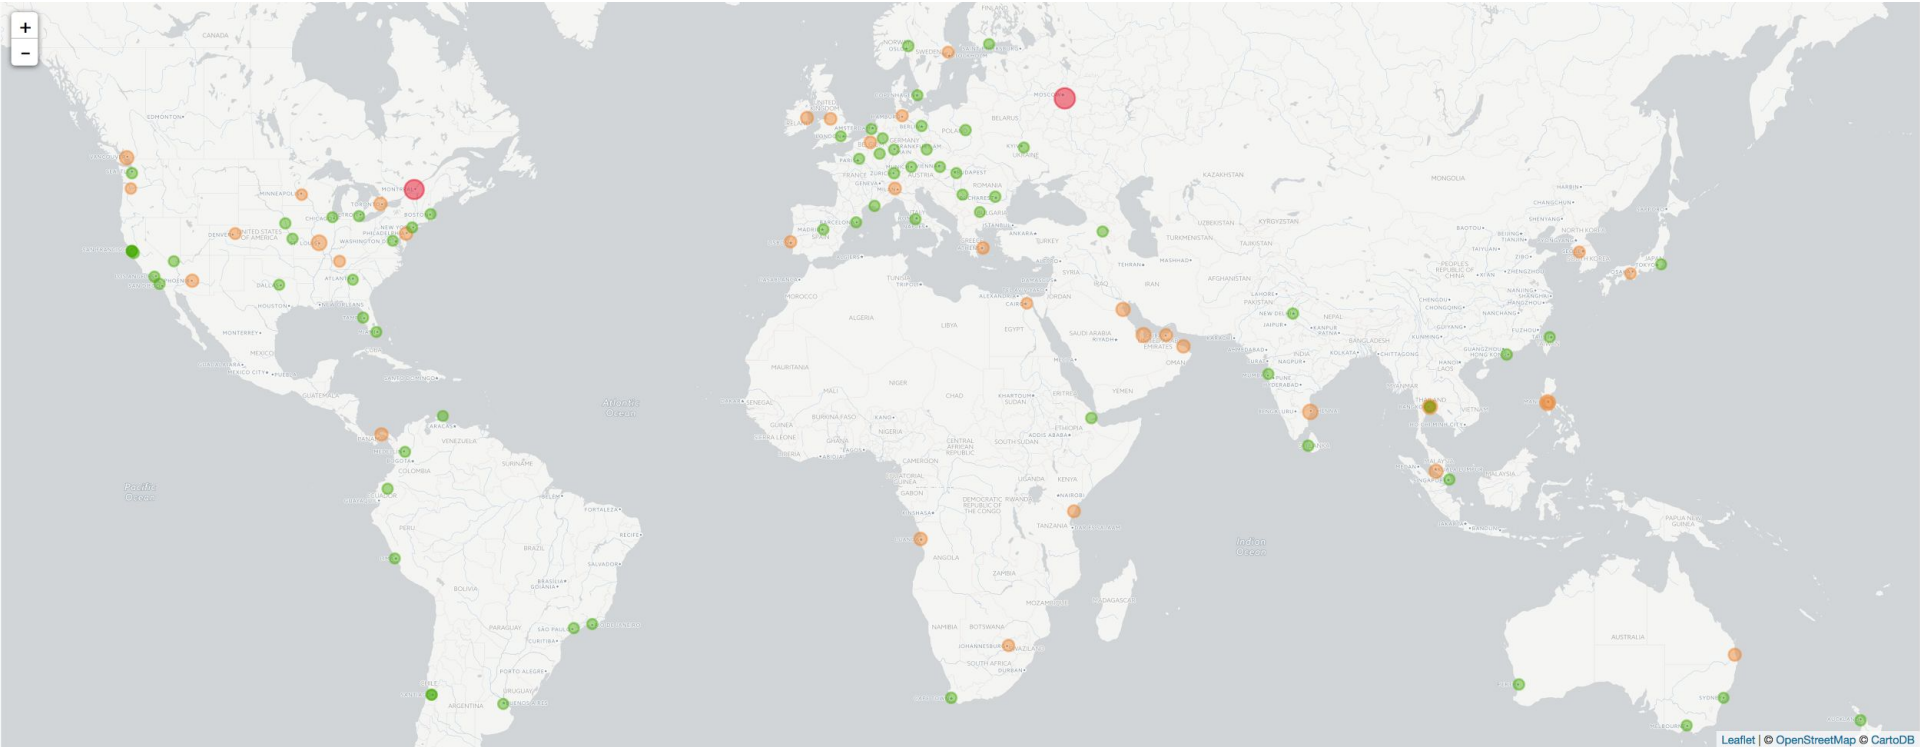


# Network lag





# Network lag



Leaflet | © OpenStreetMap © CartoDB

# #1 Alert cascade

- An alert close to the root triggers alerts in children  
→ Inhibition

## #2 Inhibition oh my

Alertmanager inhibition rule:

- target\_match:
  - alername: No\_Heartbeat\_Received
- source\_match\_re:
  - alername: .+

Alert inhibits itself

## #2 Inhibition oh my

### Alerts can inhibit themselves #666

Edit

New issue



Open

lmb opened this issue on Mar 23 · 9 comments



lmb commented on Mar 23 • edited



I have a bunch of alerts in the form of Prefix\_Some\_Name. I've set up the following inhibition rule:

Assignees

lmb


Labels

## #2 Inhibition oh my


- Sharp edges included
- Consequence of implementation, not design (?)

# #3 Aggregated histograms are buggy

Per machine  
Max sample  $\leq 5s$



Aggregated  
Max sample  $\leq ?!$



## #3 Aggregated histograms are buggy

- Much improved by [Jack Neely's PR](#)
- Fixed in 2.0 (?)
- More problems to come?

## #4 Federation + HA: not a dream team

- End up with multiple scrapes per job at any point in time
- Deduplication has to happen in queries
- How? Depends on your metric



## #4 Federation + HA: deduplicating histograms

```
histogram_quantile(0.99,  
  sum by(label, le) (  
    max without(common1,common2) (irate(bucket[3m]))  
  )  
)
```

- Aggregate per label, le
- Choose a single scrape (min works too)

## #5 SD + up: to be or to be absent

- Scenario

- Service running in Marathon
- Service discovery used to find targets

```
sum(up{job="something"}) by (job) < 1  
or absent(up{job="something"})
```

## In summary

- Prometheus is great!1!!!
- Inhibition needs love
- Querying federated HA requires extra care