Graphite@Scale:
How to store millions of metrics per second

Booking.com

Vladimir Smirnov
System Administrator

GrafanaCon EU 2018
1 March 2017

Why you might need to store your metrics?

Most common cases:

- Capacity planning
- Troubleshooting and Postmortems
- Visualization of business data
- And more...

Graphite and its modular architecture

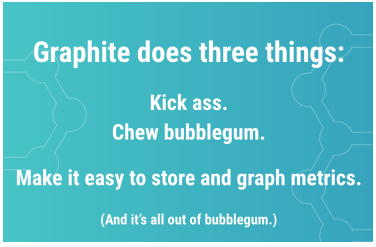

From the graphiteapp.org

- Allows to store time-series data
- Easy to use — text protocol and HTTP API

```
echo "metric.name 1.234 $(date +%s)" | nc host 2003
```

- Modular — you can replace any part of it

Graphite: Example

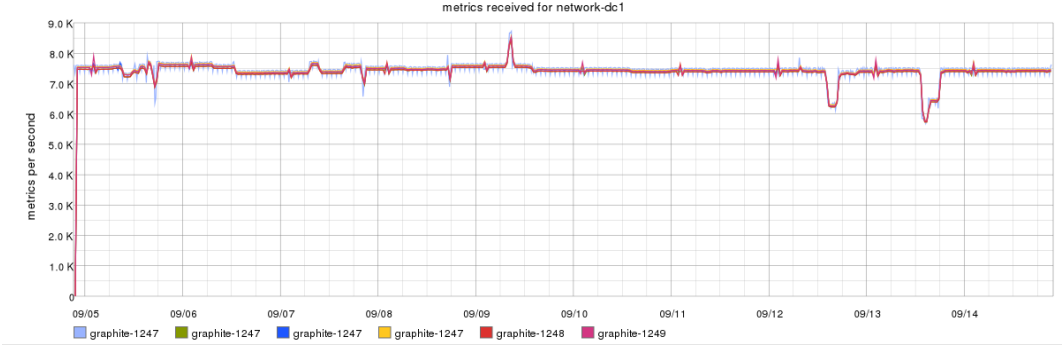

`https://host/render?target=aliasByNode(carbon.*.metricsRecevied,1)`

Our current setup

- **O(100)** Storage servers in multiple DCs
- **O(10)** of Frontend Servers
- **O(100)** TB of data in total
- **O(100 M)** unique metrics
- **O(10 M)** unique points per second
- **O(10 k)** RPS on Frontend
- **O(10 k)** of Individual Metric Requests per second
- **O(10 M)** points fetched from storage every second.

Our current setup

- **O(100)** Storage servers in multiple DCs
- **O(10)** of Frontend Servers
- **O(100)** TB of data in total
- **O(100 M)** unique metrics
- **O(10 M)** unique points per second
- **O(10 k)** RPS on Frontend
- **O(10 k)** of Individual Metric Requests per second
- **O(10 M)** points fetched from storage every second.

Our current setup

- **O(100)** Storage servers in multiple DCs
- **O(10)** of Frontend Servers
- **O(100)** TB of data in total
- **O(100 M)** unique metrics
- **O(10 M)** unique points per second
- **O(10 k)** RPS on Frontend
- **O(10 k)** of Individual Metric Requests per second
- **O(10 M)** points fetched from storage every second.

Our current setup

- **O(100)** Storage servers in multiple DCs
- **O(10)** of Frontend Servers
- **O(100)** TB of data in total
- **O(100 M)** unique metrics
- **O(10 M)** unique points per second
- **O(10 k)** RPS on Frontend
- **O(10 k)** of Individual Metric Requests per second
- **O(10 M)** points fetched from storage every second.

Our current setup

- **O(100)** Storage servers in multiple DCs
- **O(10)** of Frontend Servers
- **O(100)** TB of data in total
- **O(100 M)** unique metrics
- **O(10 M)** unique points per second
- **O(10 k)** RPS on Frontend
- **O(10 k)** of Individual Metric Requests per second
- **O(10 M)** points fetched from storage every second.

Our current setup

- **O(100)** Storage servers in multiple DCs
- **O(10)** of Frontend Servers
- **O(100)** TB of data in total
- **O(100 M)** unique metrics
- **O(10 M)** unique points per second
- **O(10 k)** RPS on Frontend
- **O(10 k)** of Individual Metric Requests per second
- **O(10 M)** points fetched from storage every second.

Our current setup

- **O(100)** Storage servers in multiple DCs
- **O(10)** of Frontend Servers
- **O(100)** TB of data in total
- **O(100 M)** unique metrics
- **O(10 M)** unique points per second
- **O(10 k)** RPS on Frontend
- **O(10 k)** of Individual Metric Requests per second
- **O(10 M)** points fetched from storage every second.

Our current setup

- **O(100)** Storage servers in multiple DCs
- **O(10)** of Frontend Servers
- **O(100)** TB of data in total
- **O(100 M)** unique metrics
- **O(10 M)** unique points per second
- **O(10 k)** RPS on Frontend
- **O(10 k)** of Individual Metric Requests per second
- **O(10 M)** points fetched from storage every second.

Original stack

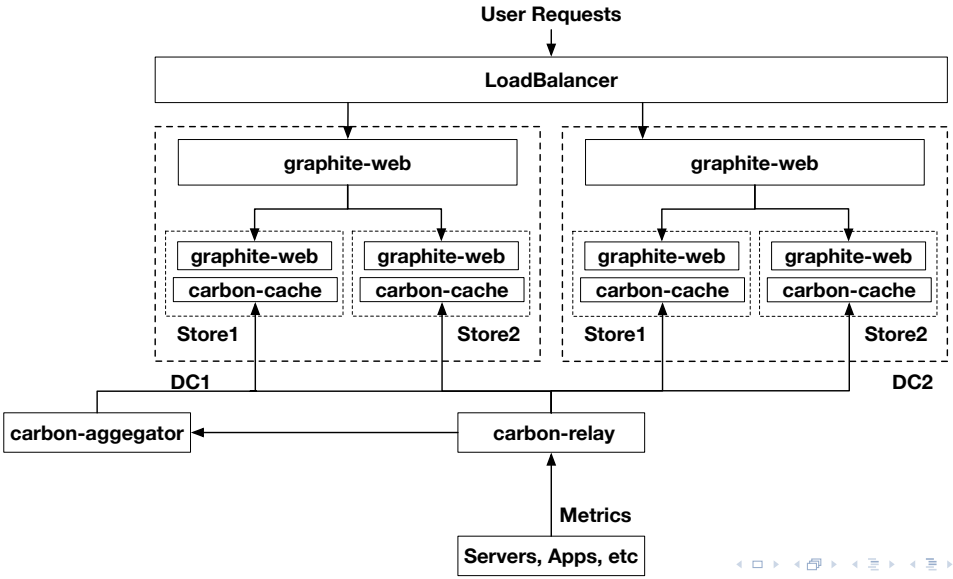

Problems: Relay

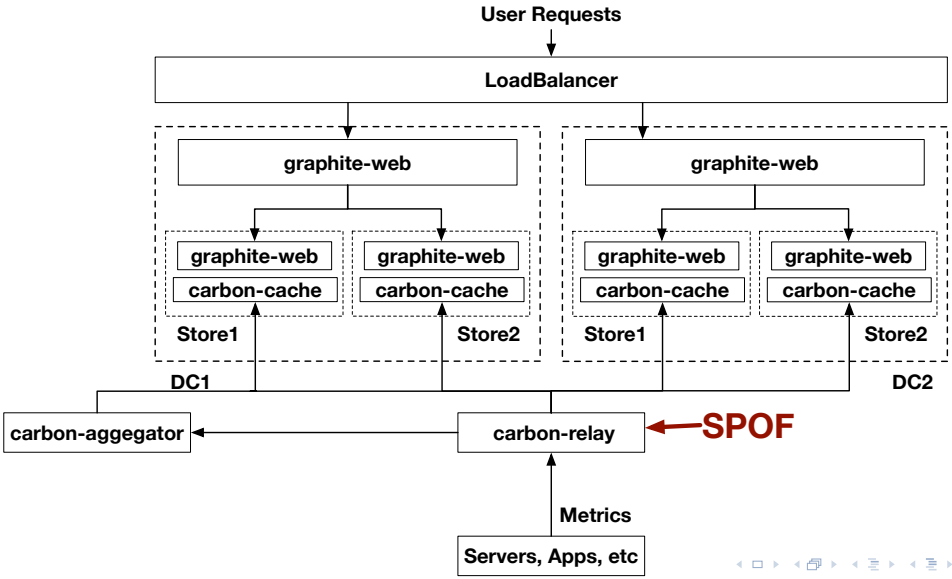

Problems: Scalability

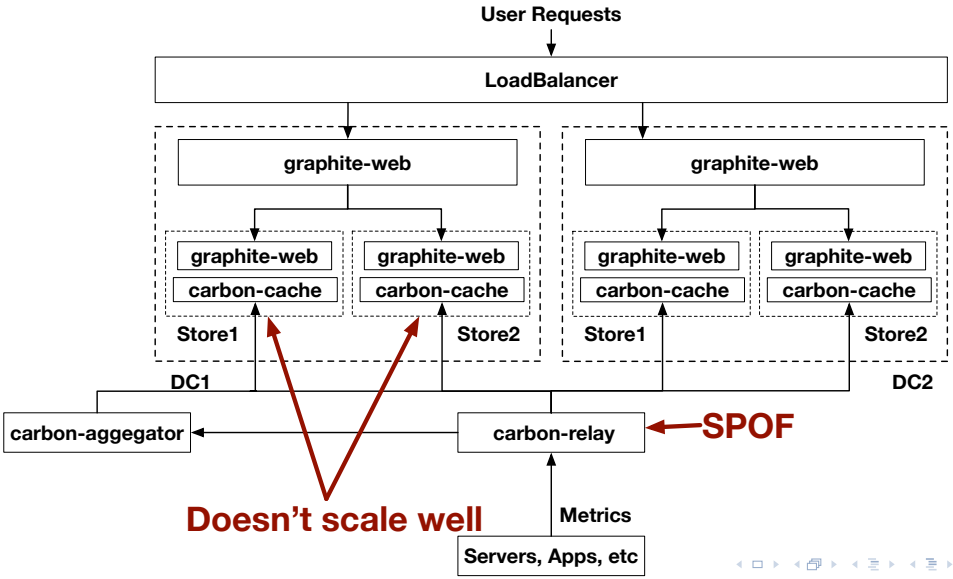

Problems: Consistency

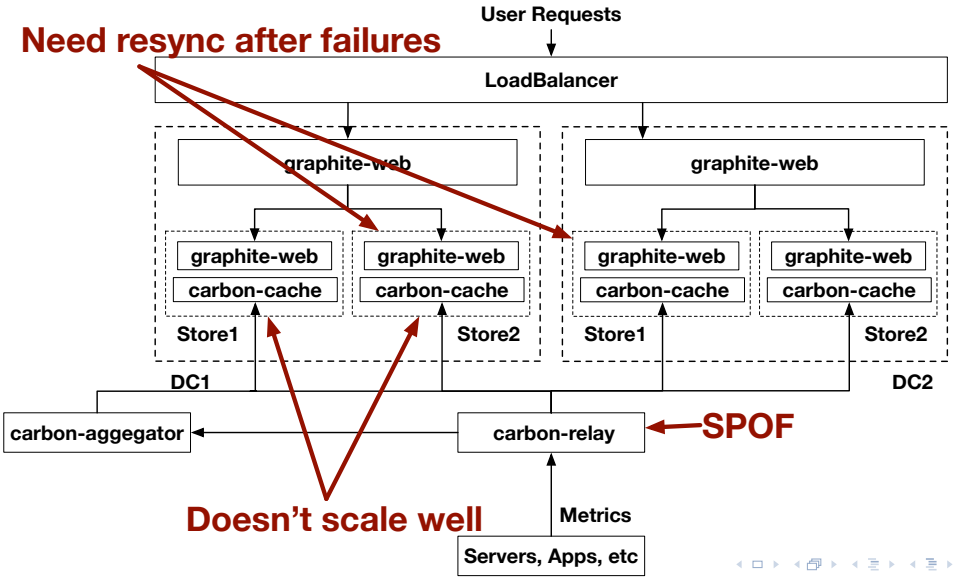

Problems: Render time

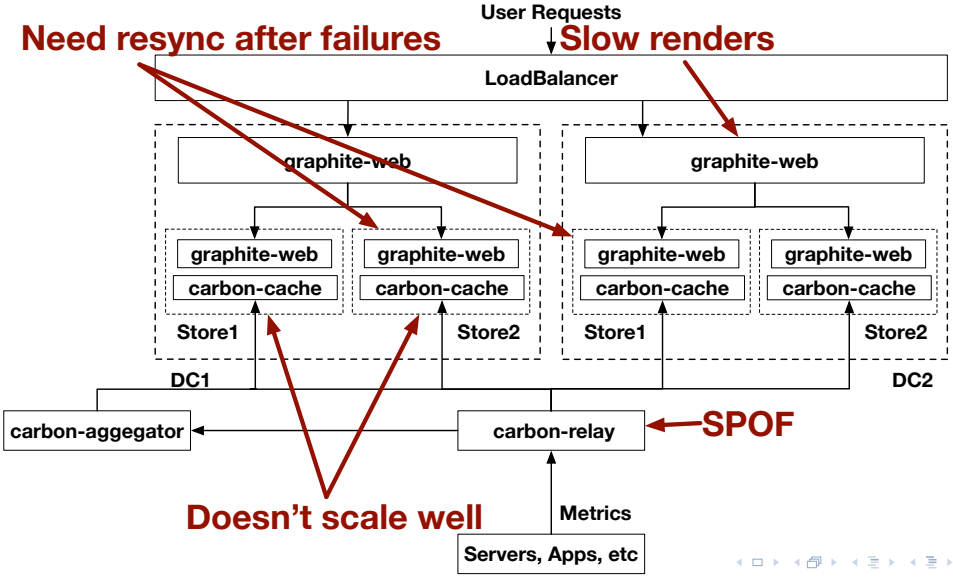

Replacing carbon-relay

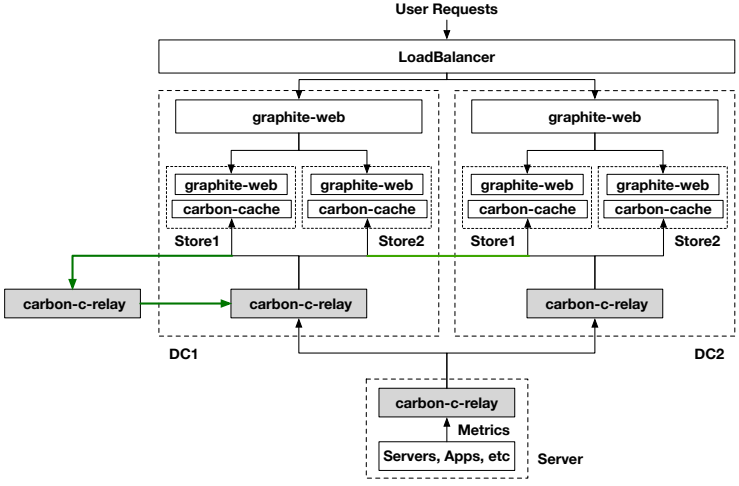

Replacing carbon-relay

carbon-c-relay:

- Written in **C**
- Routes **1M** data points per second using only **2** cores
- L7 LB for graphite line protocol (RR with sticking)
- Can do aggregations
- Buffers the data if upstream is unavailable

Zipper stack: Solution

Query: target=sys.server.cpu.user

Result:

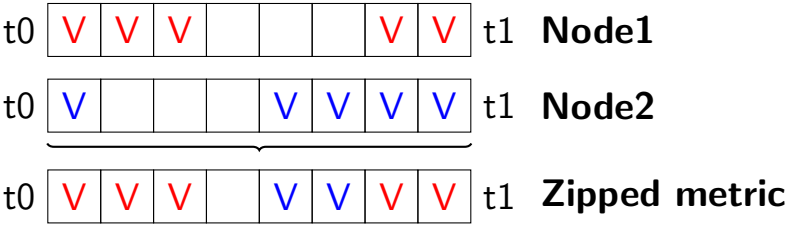

Zipper stack: architecture

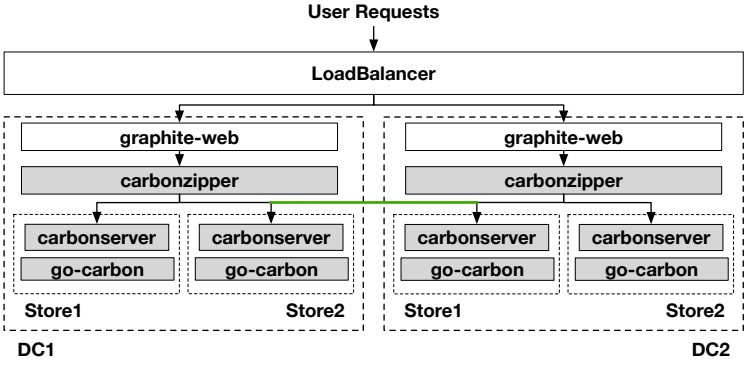

Zipper stack: results

- Written in **Go**
- Can query store servers in **parallel**
- Can "Zip" the data
- carbonzipper $\Leftrightarrow$ carbonserver — **2700** RPS
  graphite-web $\Leftrightarrow$ carbon-cache — **80** RPS.
- carbonserver is now part of go-carbon (since December 2016)

Metric distribution: how it works

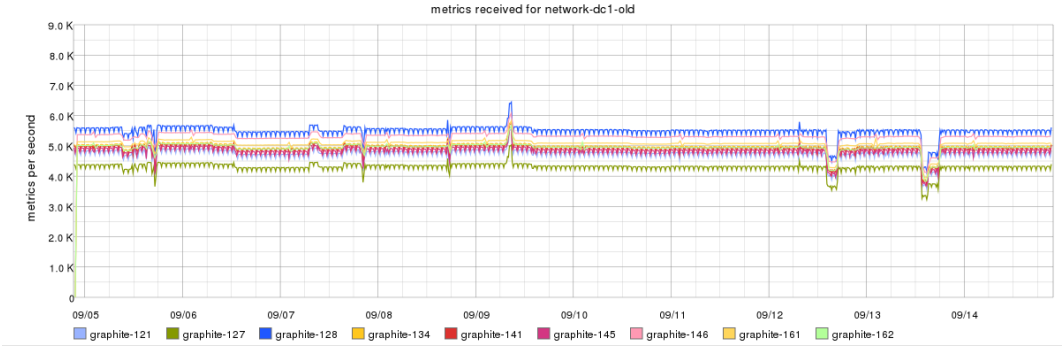

Up to **20%** difference in worst case

Metric distribution: jump hash

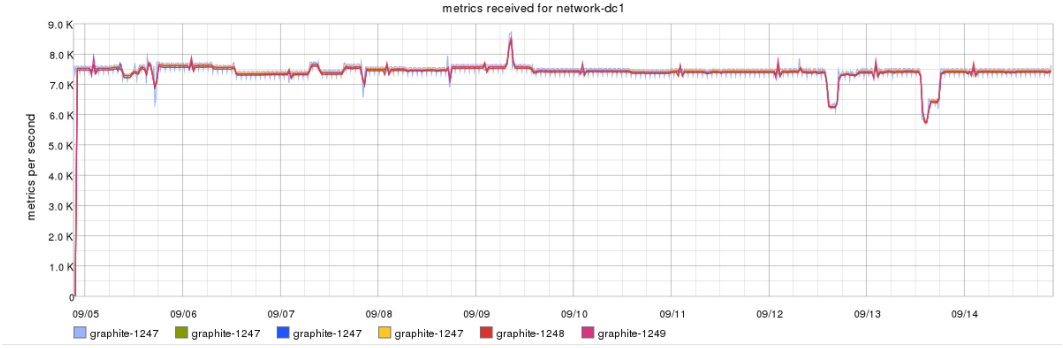

arxiv.org/pdf/1406.2294v1.pdf

Rewriting Frontend in Go: carbonapi

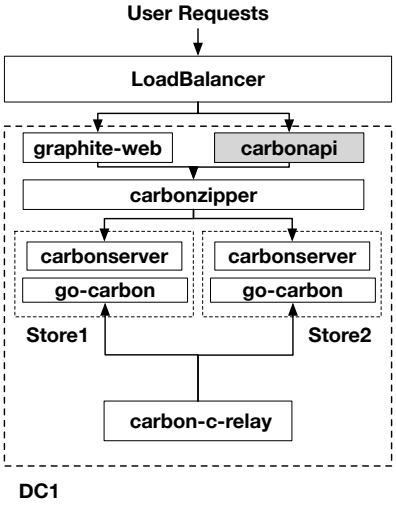

Rewriting Frontend in Go: result

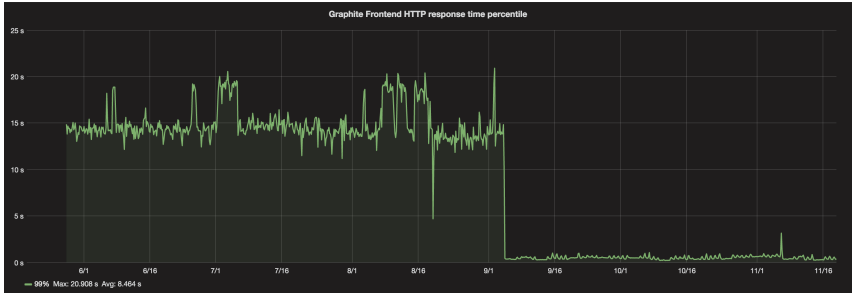

- Significantly reduced response time for users (**15s $\Rightarrow$ 0.8s**)
- Allows more complex queries because it's faster
- Easier to implement new heavy math functions
- Parsing and functions are available as separate libraries.

Replication techniques and their pros and cons

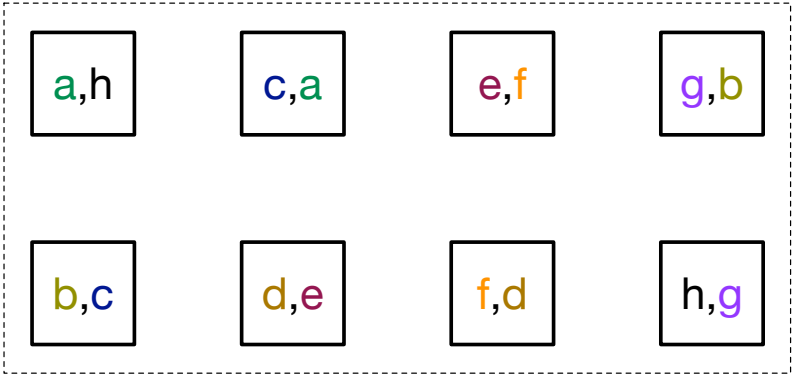

Replication Factor 2

Replication techniques and their pros and cons

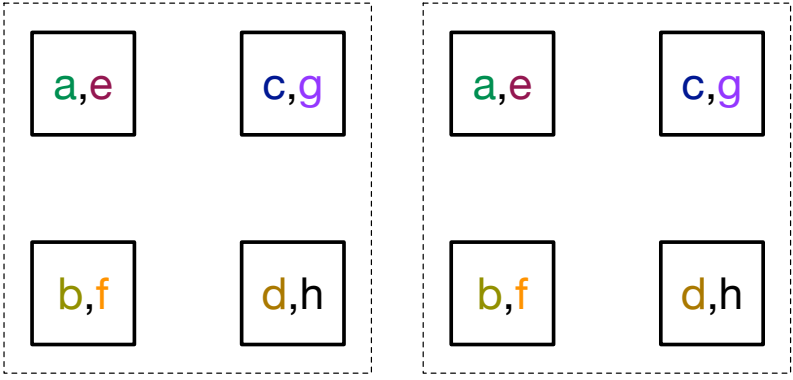

Replication Factor 1

Replication techniques and their pros and cons

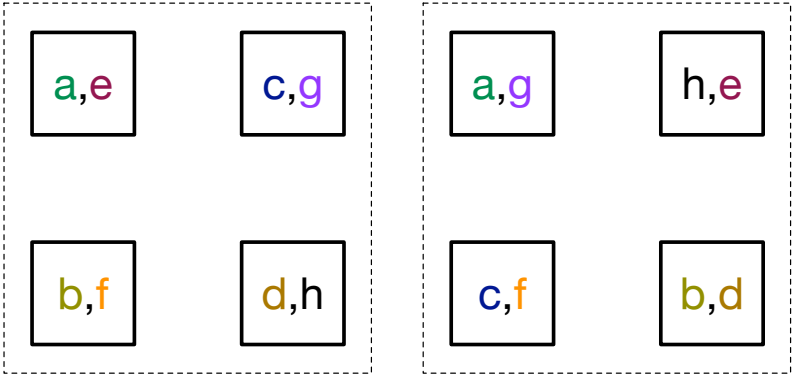

Replication Factor 1, randomized

Replication techniques and their pros and cons

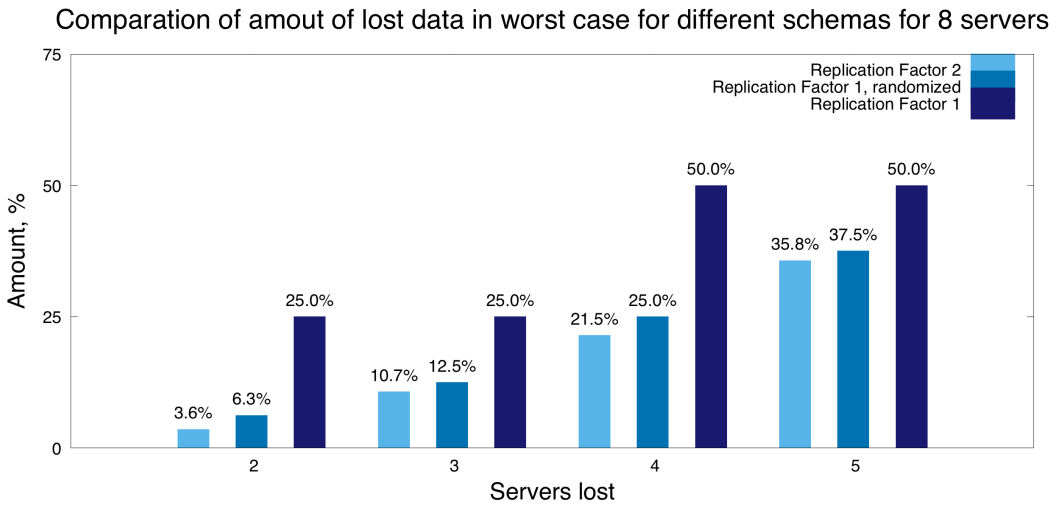

Replication techniques and their pros and cons

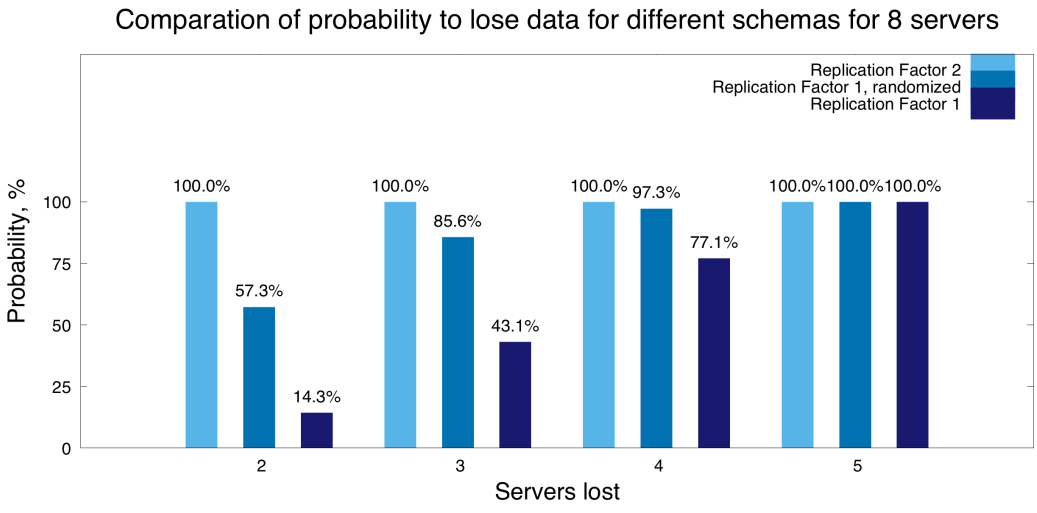

Adding simple tags

Example:
target=sum(virt.v1.*.dc:datacenter1.status:live.role:graphiteStore.text-match:metricsReceived)

- Separated tags stream and storage
- No history
- No negative match support (yet)
- Only "and" syntax

What's next?

- Find a replacement for Whisper (in progress)
- Replace graphite line protocol between components (in progress)
- Migrate to streaming protocol between backends (in progress).
- Implement differential flamegraphs
- Continue to work on collecting traces

It's all Open Source!

- carbon-c-relay — github.com/grobian/carbon-c-relay
- carbonzipper — github.com/go-graphite/carbonzipper
- go-carbon — github.com/lomik/go-carbon
- carbonapi — github.com/go-graphite/carbonapi
- carbonsearch — github.com/kanatohodets/carbonsearch
- gorelka — github.com/go-graphite/gorelka
- flamegraphs — github.com/Civil/ch-flamegraphs
- replication factor test — github.com/Civil/graphite-rf-test

Several major users: Booking.com, eBay Classifieds Group and Slack

Questions?

vladimir.smirnov@booking.com

civil.over@gmail.com

Twitter: @Civilus

Facebook: civilus

Telegram: Civiloid

LinkedIn: vladsmirnov

Thanks!

We are hiring SREs in Amsterdam!
https://workingatbooking.com

Bonus: Instrumenting: FlameGraphs: Before

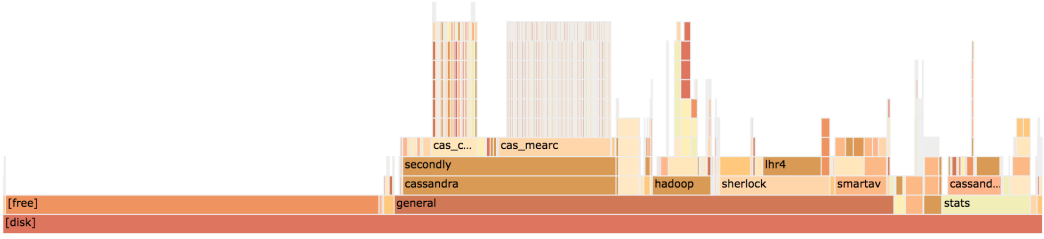

Bonus: Instrumenting: FlameGraphs: After

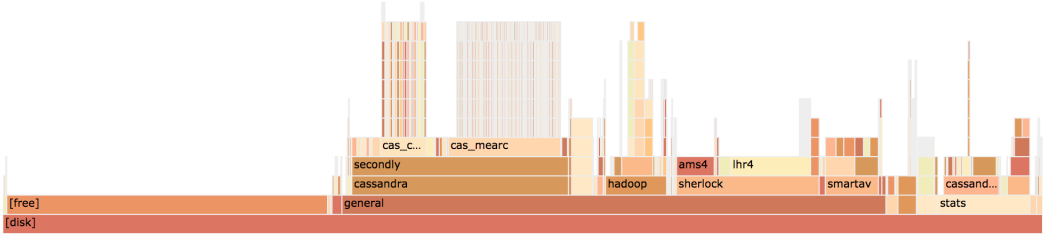

Bonus: Instrumenting: Conclusion

- Collect and Store information about every metric
- Database: Clickhouse
- Stores raw data about each metric: name, size, mtime, access time, etc.

Bonus: Instrumenting: Profiling stack