



wide event analytics

@igonwhilefalse

hello!



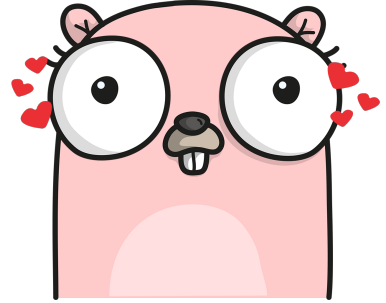
@igonwhilefalse

gentle constructive rant

debugging large scale  
systems using events

understanding  
system behaviour

app

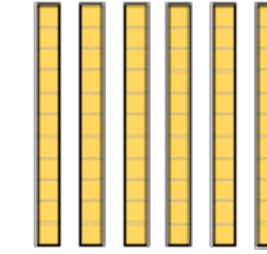


events

{ k: v }



column store



analytical  
queries

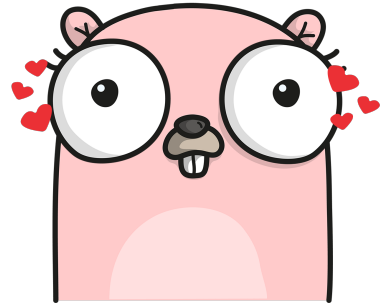
SELECT ...  
GROUP BY



users



app

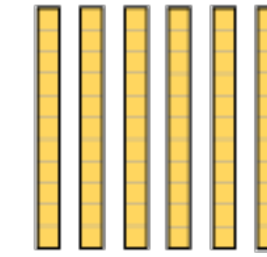


events

{ k: v }



column store



analytical  
queries

SELECT ...  
GROUP BY

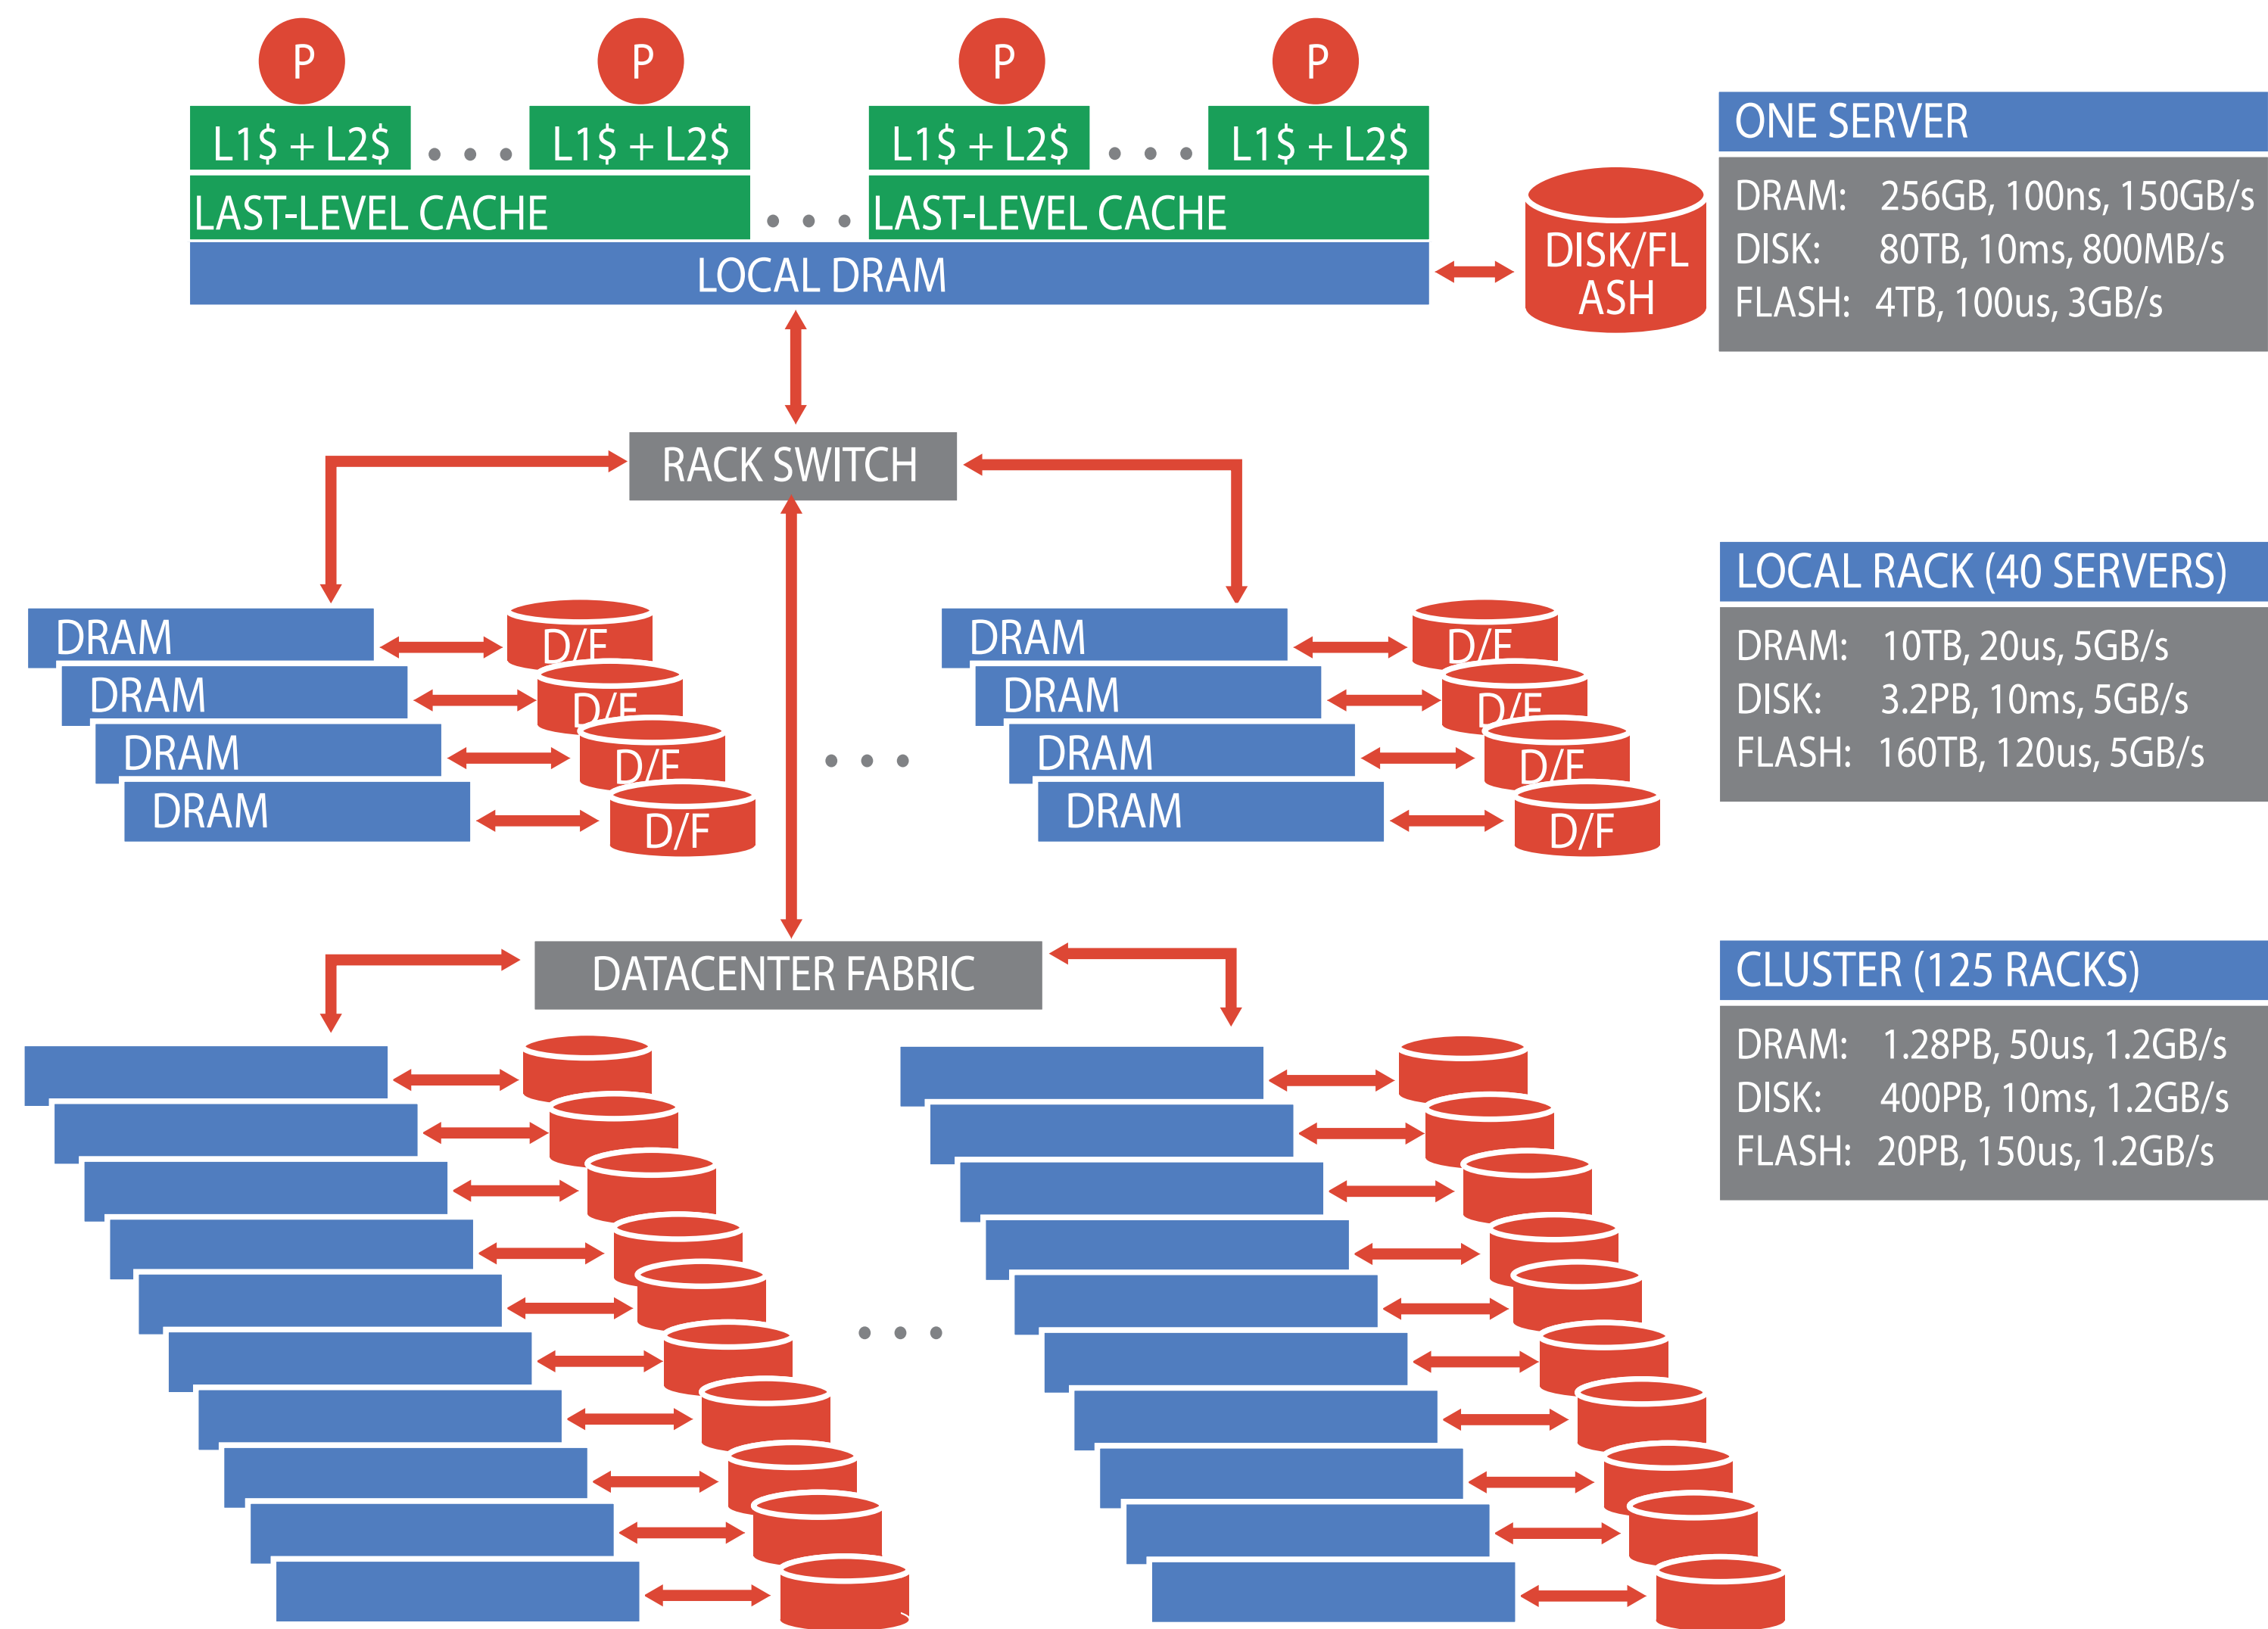


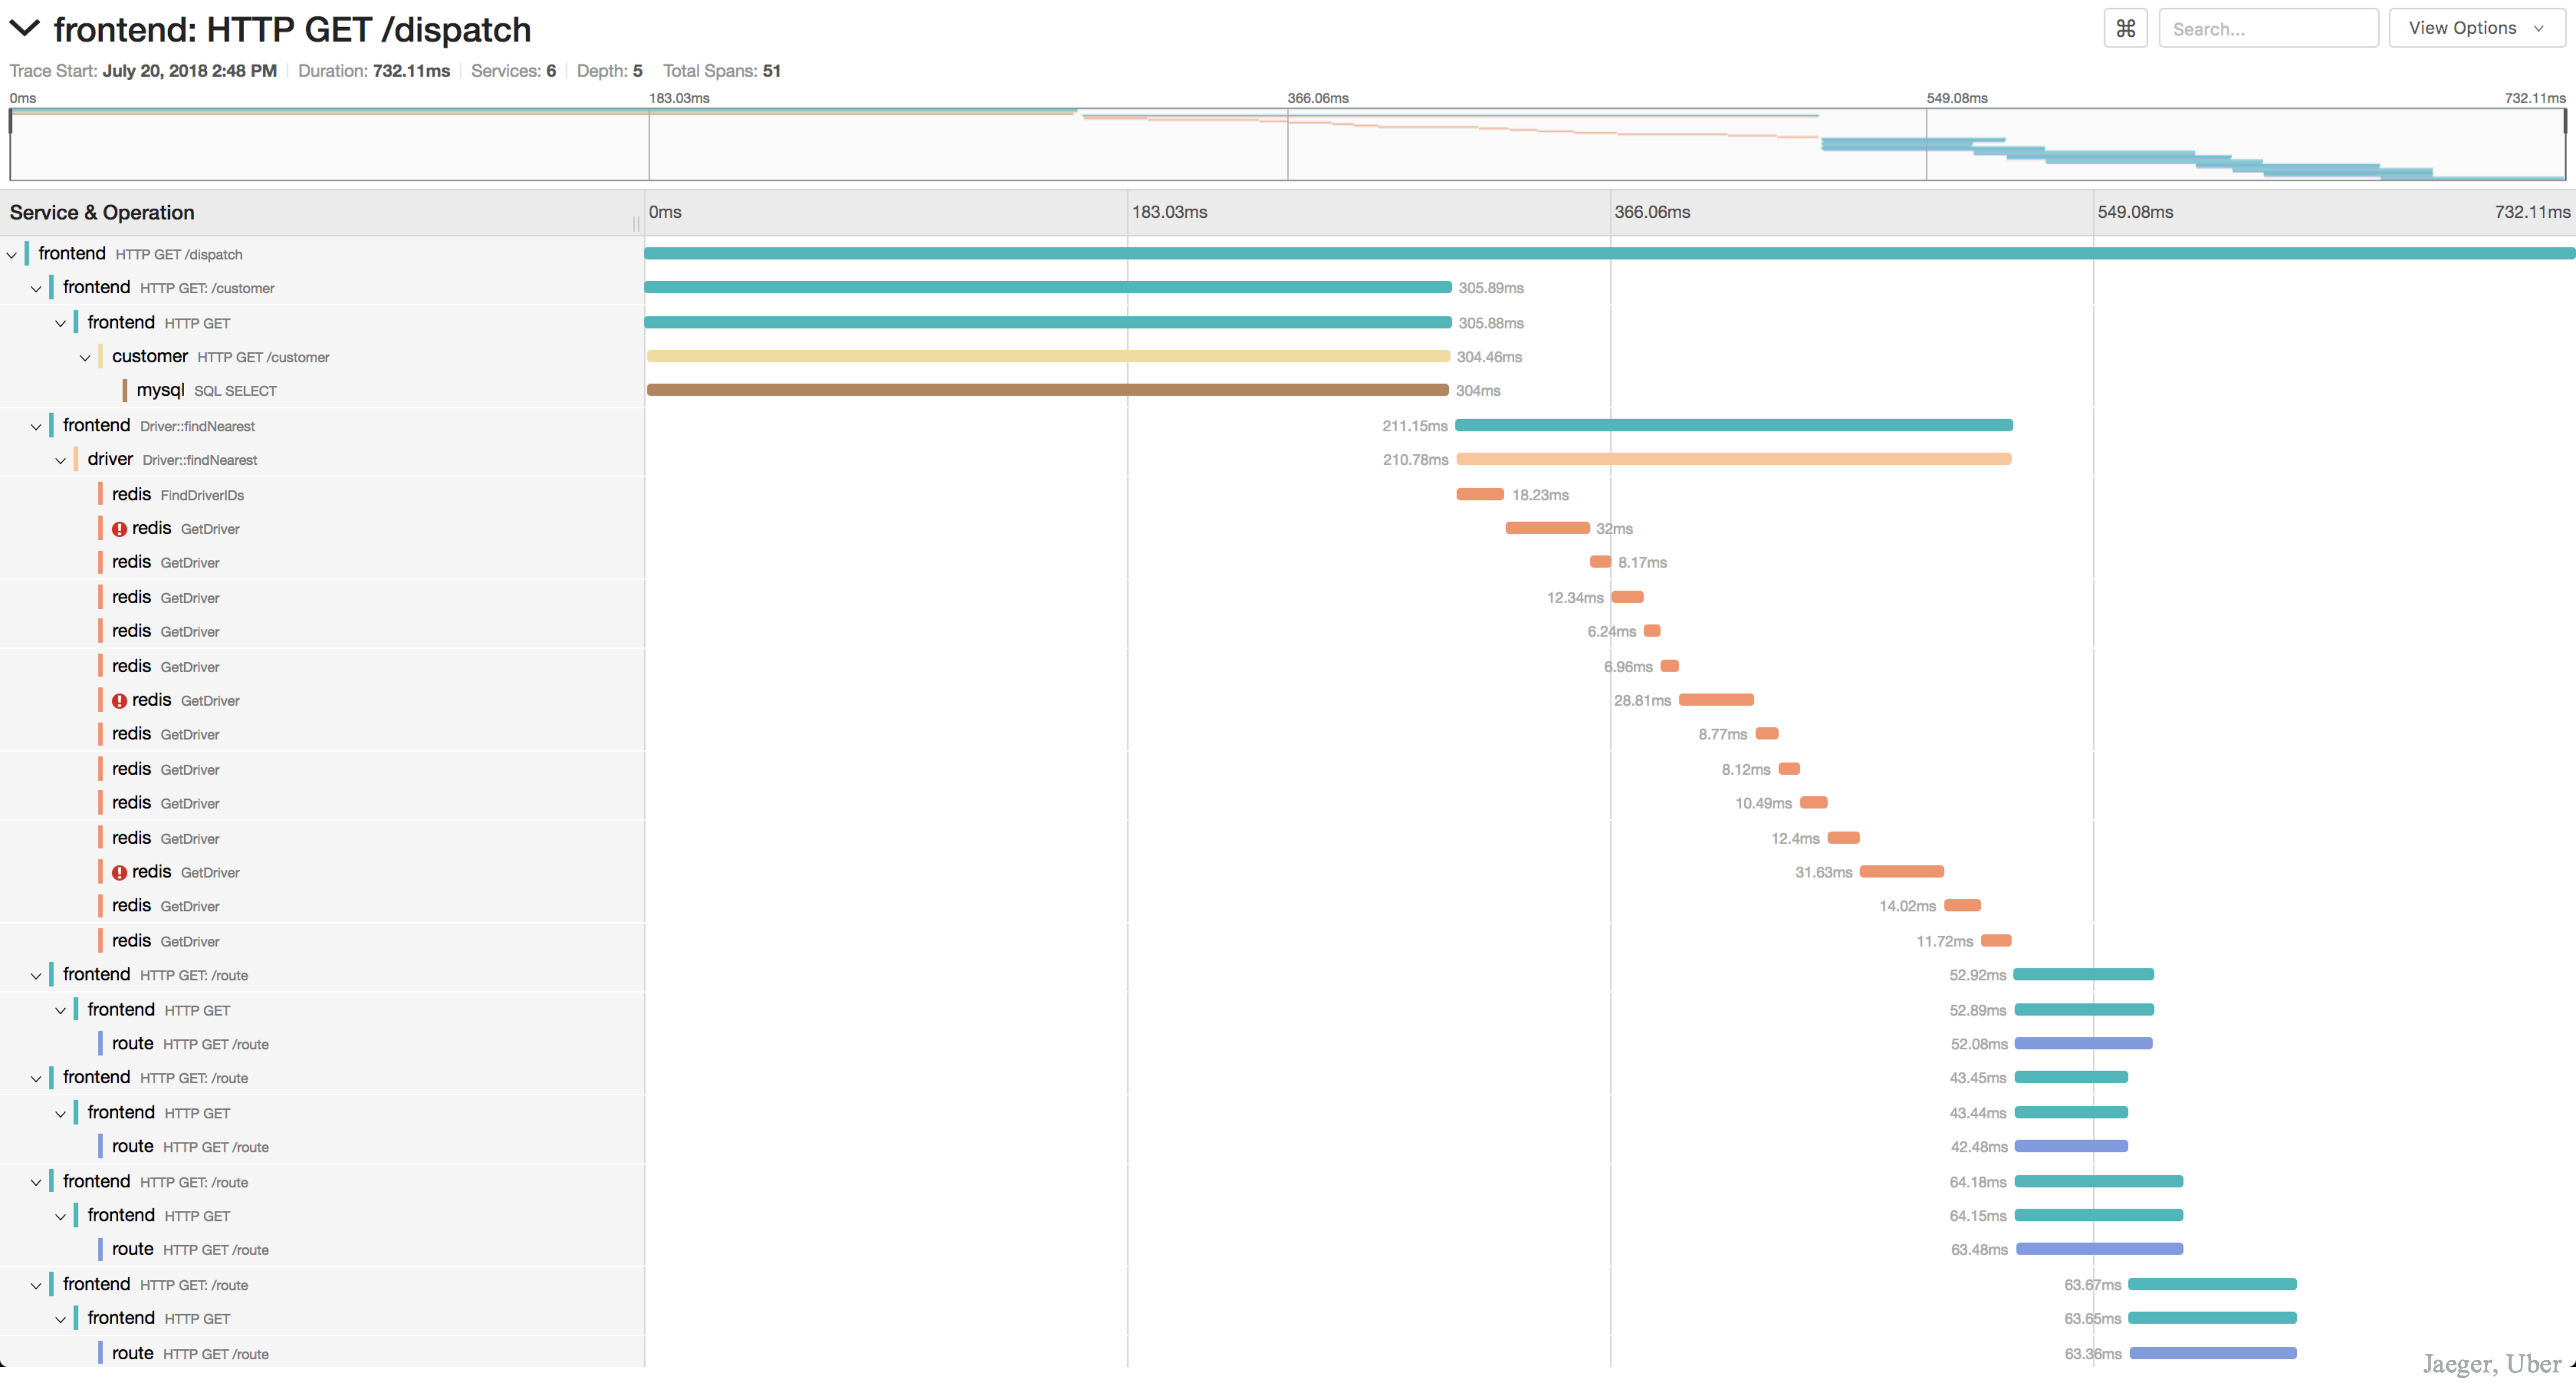
users



you are here

software is becoming  
increasingly complex







PREPARE TO STOP

NO DANGEROUS GOODS IN TUNNEL

LOW CLEARANCE 4.4m

CAHILL TUNNEL

NO DANGEROUS GOODS IN TUNNEL

LOW TUNNEL CLEARANCE 4.4m

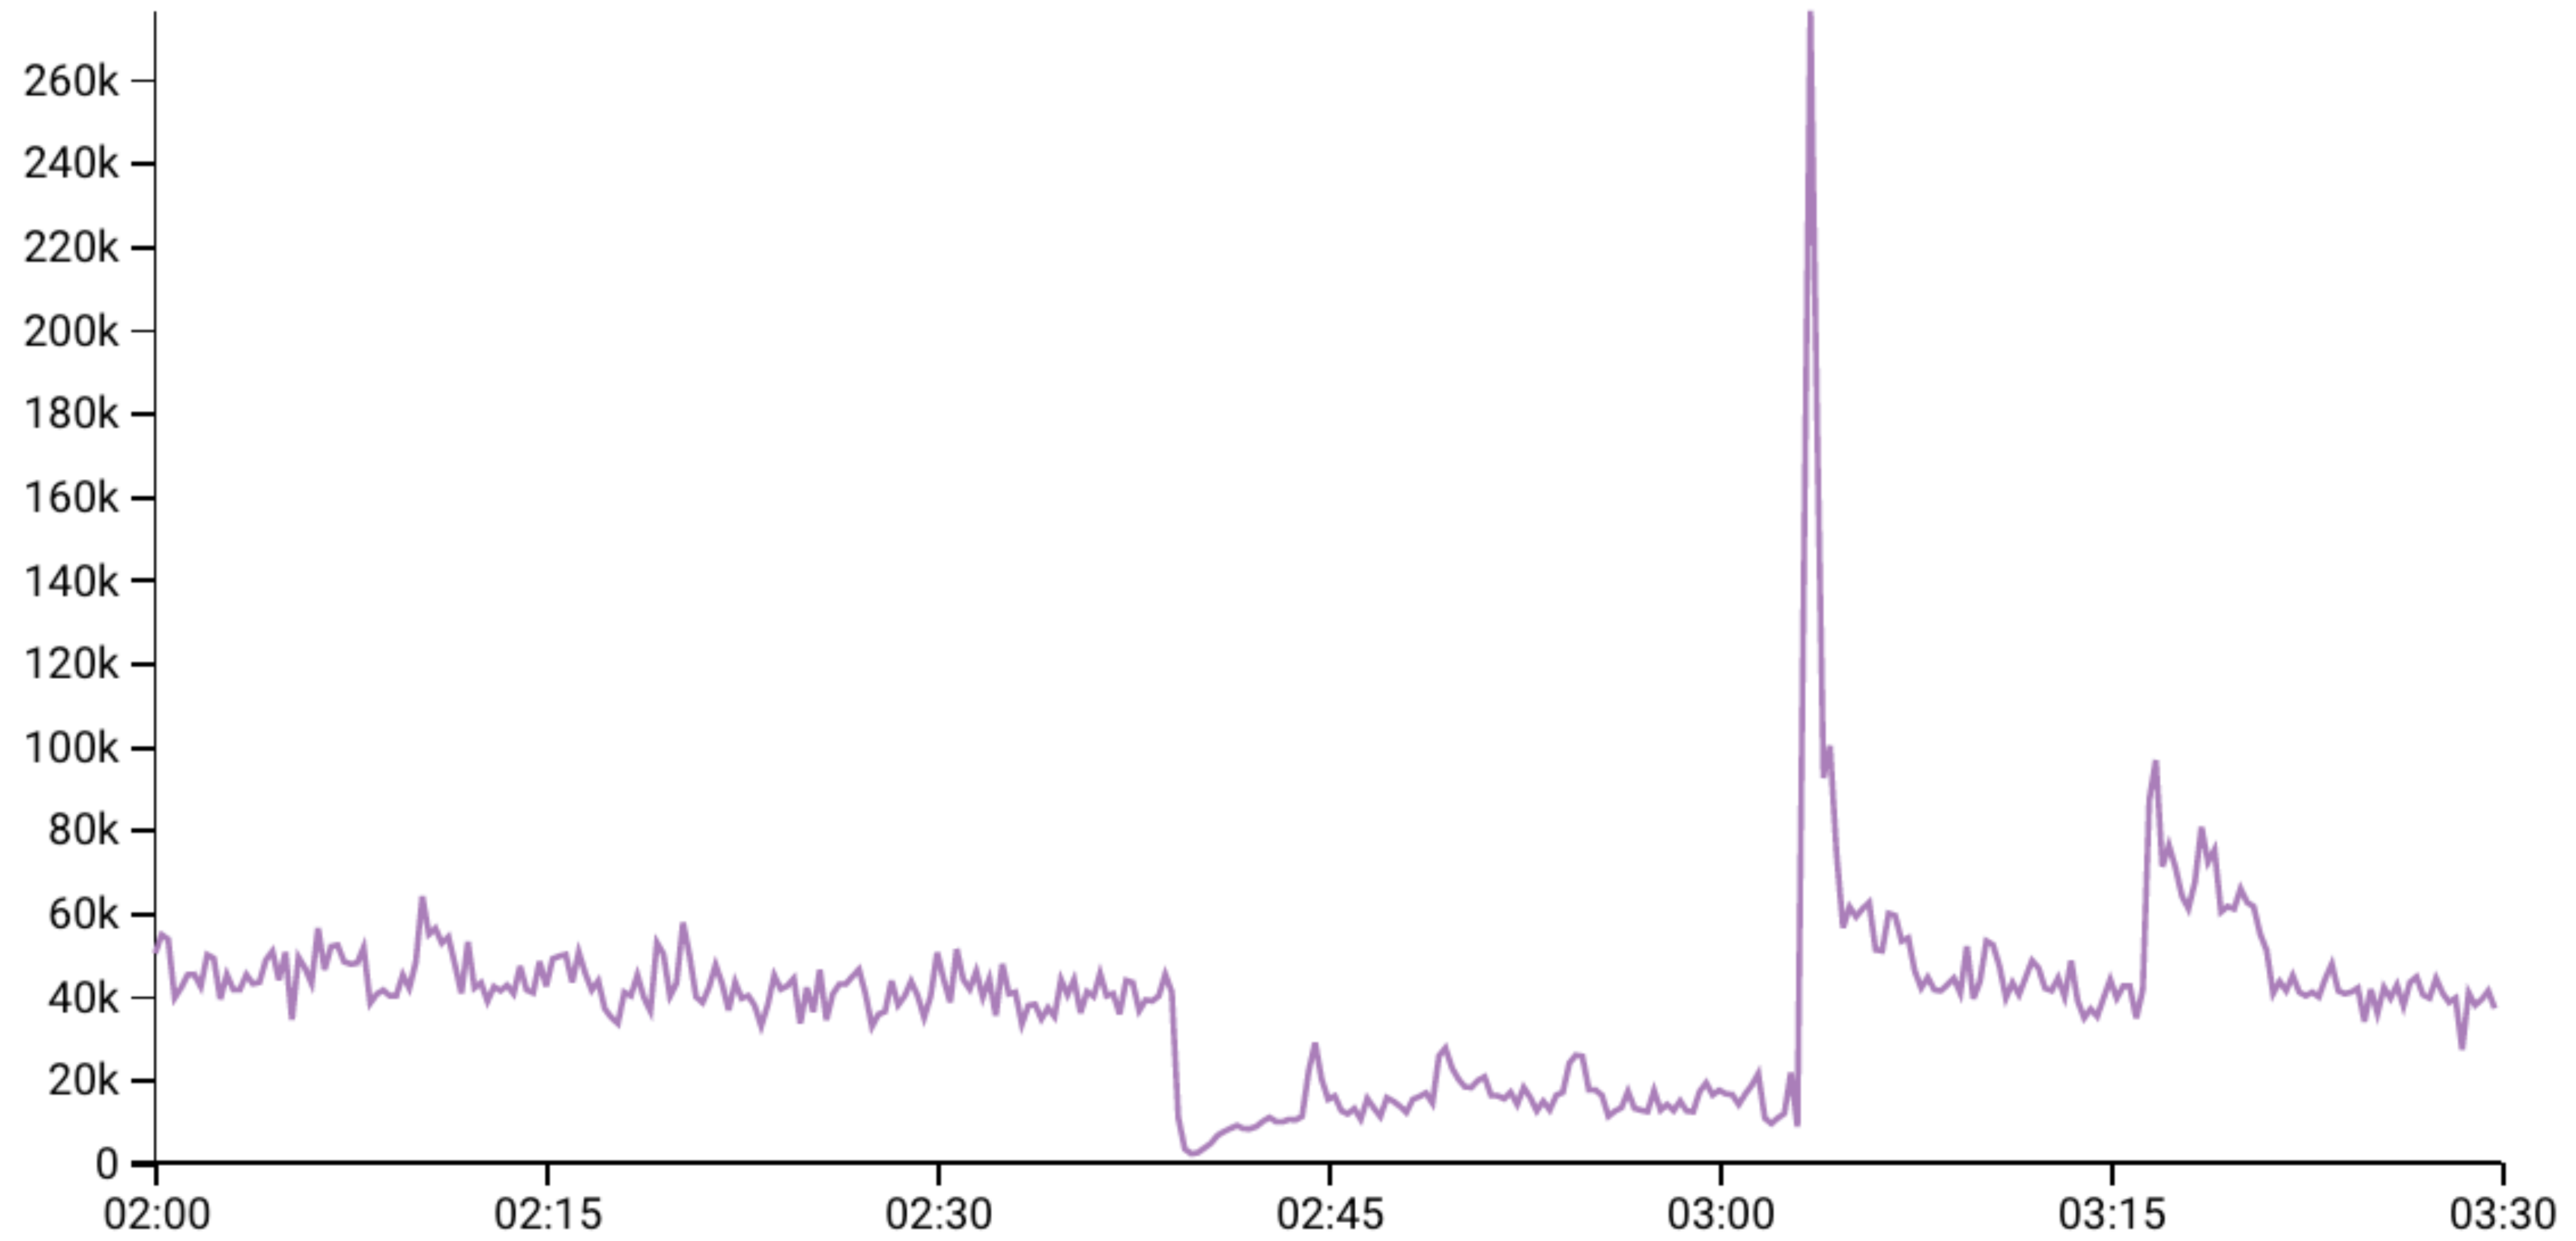
DETOUR



logs vs metrics:  
a false dichotomy

10.2.3.4 - - [1/Jan/1970:18:32:20  
+0000] "GET / HTTP/1.1" 200 5324  
"- " "curl/7.54.0" "-"

**COUNT**



we can derive metrics  
from log streams

```
$ cat access.log  
    | grep ... | awk ...  
    | sort | uniq -c
```

```
{
  time          = "1970-01-01T18:32:20"

  status        = 200
  method        = "GET"
  path          = ...
  host          = "i-123456af"
  client_ip     = "10.2.3.4"
  user_agent    = "curl/7.54.0"

  request_dur_ms = 325
  request_bytes  = 2456
  response_bytes = 5324
}
```

structured logs

summary events

canonical log lines

arbitrarily wide data blobs



events



a metric is an aggregation  
of events

why do we aggregate?

1. Python

38.113.234.181 - - [20/May/2010:11:27:36 +0100] "GET /post/274953/ HTTP/1.1" 404 15 "-" "Voyager/1.0"

67.195.114.50 - - [20/May/2010:11:31:16 +0100] "GET /sources/14/ HTTP/1.0" 200 14262 "-" "Mozilla/5.0 (compatible; Yahoo! Slurp/3.0; http://help.yahoo.com/help/us/ysearch/slurp)"

61.135.216.104 - - [20/May/2010:11:32:56 +0100] "GET /feeds/latest/ HTTP/1.1" 200 48364 "-" "Mozilla/5.0 (compatible; YoudaoFeedFetcher/1.0;http://www.youdao.com/help/reader/faq/topic006/;1 subscribers;)"

125.22.2.42 - - [20/May/2010:11:33:09 +0100] "GET /feeds/latest/ HTTP/1.1" 200 48364 "-" "Mozilla/4.0 (compatible; MSIE 7.0; Windows NT 5.1; .NET CLR 1.1.4322; .NET CLR 2.0.50727; MS-RTC LM 8; .NET CLR 3.0.04506.648; .NET CLR 3.5.21022; MSOffice 12)"

77.88.26.26 - - [20/May/2010:11:33:50 +0100] "GET /posts/2234/ HTTP/1.1" 200 33674 "-" "Yandex/1.01.001 (compatible; Win16; I)"

24.62.152.212 - - [20/May/2010:11:34:07 +0100] "GET /media/style.css HTTP/1.1" 200 4847 "http://www.example.com/" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_6\_3; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

24.62.152.212 - - [20/May/2010:11:34:07 +0100] "GET /media/exmpl.png HTTP/1.1" 200 28479 "http://www.example.com/" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_6\_3; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

24.62.152.212 - - [20/May/2010:11:34:05 +0100] "GET / HTTP/1.1" 200 110836 "-" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_6\_3; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

24.62.152.212 - - [20/May/2010:11:34:08 +0100] "GET /media/img/m-inact.gif HTTP/1.1" 200 2571 "http://www.example.com/" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_6\_3; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

24.62.152.212 - - [20/May/2010:11:34:08 +0100] "GET /media/img/side-container.gif HTTP/1.1" 200 1415 "http://www.example.com/" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_6\_3; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

24.62.152.212 - - [20/May/2010:11:34:08 +0100] "GET /media/img/m-act.gif HTTP/1.1" 200 143 "http://www.example.com/" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_6\_3; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

166.205.4.62 - - [20/May/2010:10:39:16 +0000] "GET /media/style.css HTTP/1.1" 304 - "http://example.com/" "Mozilla/5.0 (iPad; U; CPU OS 3\_2 like Mac OS X; en-us) AppleWebKit/531.21.10 (KHTML, like Gecko) Version/4.0.4 Mobile/7B367 Safari/531.21.10"

166.205.4.62 - - [20/May/2010:11:39:16 +0100] "GET /media/exmpl.png HTTP/1.1" 304 - "http://example.com/" "Mozilla/5.0 (iPad; U; CPU OS 3\_2 like Mac OS X; en-us) AppleWebKit/531.21.10 (KHTML, like Gecko) Version/4.0.4 Mobile/7B367 Safari/531.21.10"

166.205.4.62 - - [20/May/2010:11:39:13 +0100] "GET / HTTP/1.1" 200 26130 "-" "Mozilla/5.0 (iPad; U; CPU OS 3\_2 like Mac OS X; en-us) AppleWebKit/531.21.10 (KHTML, like Gecko) Version/4.0.4 Mobile/7B367 Safari/531.21.10"

166.205.4.62 - - [20/May/2010:10:39:17 +0000] "GET /media/img/m-act.gif HTTP/1.1" 304 - "http://example.com/" "Mozilla/5.0 (iPad; U; CPU OS 3\_2 like Mac OS X; en-us) AppleWebKit/531.21.10 (KHTML, like Gecko) Version/4.0.4 Mobile/7B367 Safari/531.21.10"

166.205.4.62 - - [20/May/2010:11:39:18 +0100] "GET /media/img/m-inact.gif HTTP/1.1" 304 - "http://example.com/" "Mozilla/5.0 (iPad; U; CPU OS 3\_2 like Mac OS X; en-us) AppleWebKit/531.21.10 (KHTML, like Gecko) Version/4.0.4 Mobile/7B367 Safari/531.21.10"

166.205.4.62 - - [20/May/2010:10:39:20 +0000] "GET /media/img/side-container.gif HTTP/1.1" 304 - "http://example.com/" "Mozilla/5.0 (iPad; U; CPU OS 3\_2 like Mac OS X; en-us) AppleWebKit/531.21.10 (KHTML, like Gecko) Version/4.0.4 Mobile/7B367 Safari/531.21.10"

67.195.114.50 - - [20/May/2010:11:40:09 +0100] "GET /post/249681/ HTTP/1.0" 404 15 "-" "Mozilla/5.0 (compatible; Yahoo! Slurp/3.0; http://help.yahoo.com/help/us/ysearch/slurp)"

67.195.114.50 - - [20/May/2010:11:40:15 +0100] "GET /post/256414/ HTTP/1.0" 404 15 "-" "Mozilla/5.0 (compatible; Yahoo! Slurp/3.0; http://help.yahoo.com/help/us/ysearch/slurp)"

220.181.7.76 - - [20/May/2010:11:42:35 +0100] "GET / HTTP/1.1" 200 26130 "-" "Baiduspider+(http://www.baidu.com/search/spider.htm)"

67.195.114.50 - - [20/May/2010:11:42:58 +0100] "GET /post/274910/ HTTP/1.0" 404 15 "-" "Mozilla/5.0 (compatible; Yahoo! Slurp/3.0; http://help.yahoo.com/help/us/ysearch/slurp)"

68.59.242.134 - - [20/May/2010:10:43:14 +0000] "GET /media/style.css HTTP/1.1" 304 - "http://example.com/" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

68.59.242.134 - - [20/May/2010:11:43:14 +0100] "GET / HTTP/1.1" 200 26130 "-" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

68.59.242.134 - - [20/May/2010:10:43:16 +0000] "GET /media/img/favicon.ico HTTP/1.1" 200 1406 "-" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

68.59.242.134 - - [20/May/2010:10:44:14 +0000] "GET /media/style.css HTTP/1.1" 304 - "http://example.com/posts/2/" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

68.59.242.134 - - [20/May/2010:11:44:12 +0100] "GET /posts/2/ HTTP/1.1" 200 34408 "http://example.com/" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

67.195.114.50 - - [20/May/2010:11:45:37 +0100] "GET /post/259342/ HTTP/1.0" 404 15 "-" "Mozilla/5.0 (compatible; Yahoo! Slurp/3.0; http://help.yahoo.com/help/us/ysearch/slurp)"

209.85.228.82 - - [20/May/2010:11:46:18 +0100] "GET /feeds/latest/ HTTP/1.1" 200 48364 "-" "FeedBurner/1.0 (http://www.FeedBurner.com)"

209.85.228.82 - - [20/May/2010:11:46:23 +0100] "GET /feeds/latest/ HTTP/1.1" 200 48364 "-" "FeedBurner/1.0 (http://www.FeedBurner.com)"

72.14.199.102 - - [20/May/2010:11:46:33 +0100] "GET /feeds/latest/ HTTP/1.1" 200 48364 "-" "Feedfetcher-Google; (+http://www.google.com/feedfetcher.html; 9 subscribers; feed-id=18260778798542229819)"

196.203.53.144 - - [20/May/2010:11:46:37 +0100] "GET /feeds/latest/ HTTP/1.0" 200 48364 "-" "Mozilla/5.0 (Windows; U; Windows NT 5.0; fr; rv:1.9.2.3) Gecko/20100401 Firefox/3.6.3"

68.59.242.134 - - [20/May/2010:10:46:43 +0000] "GET /media/style.css HTTP/1.1" 304 - "http://example.com/posts/3/" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

68.59.242.134 - - [20/May/2010:11:46:41 +0100] "GET /posts/3/ HTTP/1.1" 200 25865 "http://example.com/posts/2/" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

68.59.242.134 - - [20/May/2010:11:48:13 +0100] "GET /media/style.css HTTP/1.1" 304 - "http://example.com/posts/4/" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

68.59.242.134 - - [20/May/2010:11:48:12 +0100] "GET /posts/4/ HTTP/1.1" 200 25930 "http://example.com/posts/3/" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

66.249.65.42 - - [20/May/2010:11:48:14 +0100] "GET /posts/4/ HTTP/1.1" 200 25930 "-" "Mediapartners-Google"

66.249.65.40 - - [20/May/2010:11:48:37 +0100] "GET /post/274703/ HTTP/1.1" 404 15 "-" "Mozilla/5.0 (compatible; Googlebot/2.1; +http://www.google.com/bot.html)"

66.249.65.40 - - [20/May/2010:11:49:13 +0100] "GET /post/274704/ HTTP/1.1" 404 15 "-" "Mozilla/5.0 (compatible; Googlebot/2.1; +http://www.google.com/bot.html)"

68.59.242.134 - - [20/May/2010:10:49:54 +0000] "GET /media/style.css HTTP/1.1" 304 - "http://example.com/posts/5/" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

68.59.242.134 - - [20/May/2010:11:49:52 +0100] "GET /posts/5/ HTTP/1.1" 200 29611 "http://example.com/posts/4/" "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; InfoPath.2)"

66.249.65.42 - - [20/May/2010:11:49:54 +0100] "GET /posts/5/ HTTP/1.1" 200 29611 "-" "Mediapartners-Google"

66.249.65.40 - - [20/May/2010:11:50:18 +0100] "GET /post/274687/ HTTP/1.1" 404 15 "-" "Mozilla/5.0 (compatible; Googlebot/2.1; +http://www.google.com/bot.html)"

66.249.65.40 - - [20/May/2010:11:51:23 +0100] "GET /post/274716/ HTTP/1.1" 404 15 "-" "Mozilla/5.0 (compatible; Googlebot/2.1; +http://www.google.com/bot.html)"

66.249.65.40 - - [20/May/2010:11:52:29 +0100] "GET /post/274712/ HTTP/1.1" 404 15 "-" "Mozilla/5.0 (compatible; Googlebot/2.1; +http://www.google.com/bot.html)"

75.57.176.28 - - [20/May/2010:11:52:42 +0100] "GET /media/style.css HTTP/1.1" 200 4847 "http://www.example.com/" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_5\_8; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

75.57.176.28 - - [20/May/2010:11:52:42 +0100] "GET / HTTP/1.1" 200 26130 "-" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_5\_8; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

75.57.176.28 - - [20/May/2010:11:52:43 +0100] "GET /media/img/m-inact.gif HTTP/1.1" 200 2571 "http://www.example.com/" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_5\_8; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

75.57.176.28 - - [20/May/2010:11:52:43 +0100] "GET /media/img/side-container.gif HTTP/1.1" 200 1415 "http://www.example.com/" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_5\_8; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

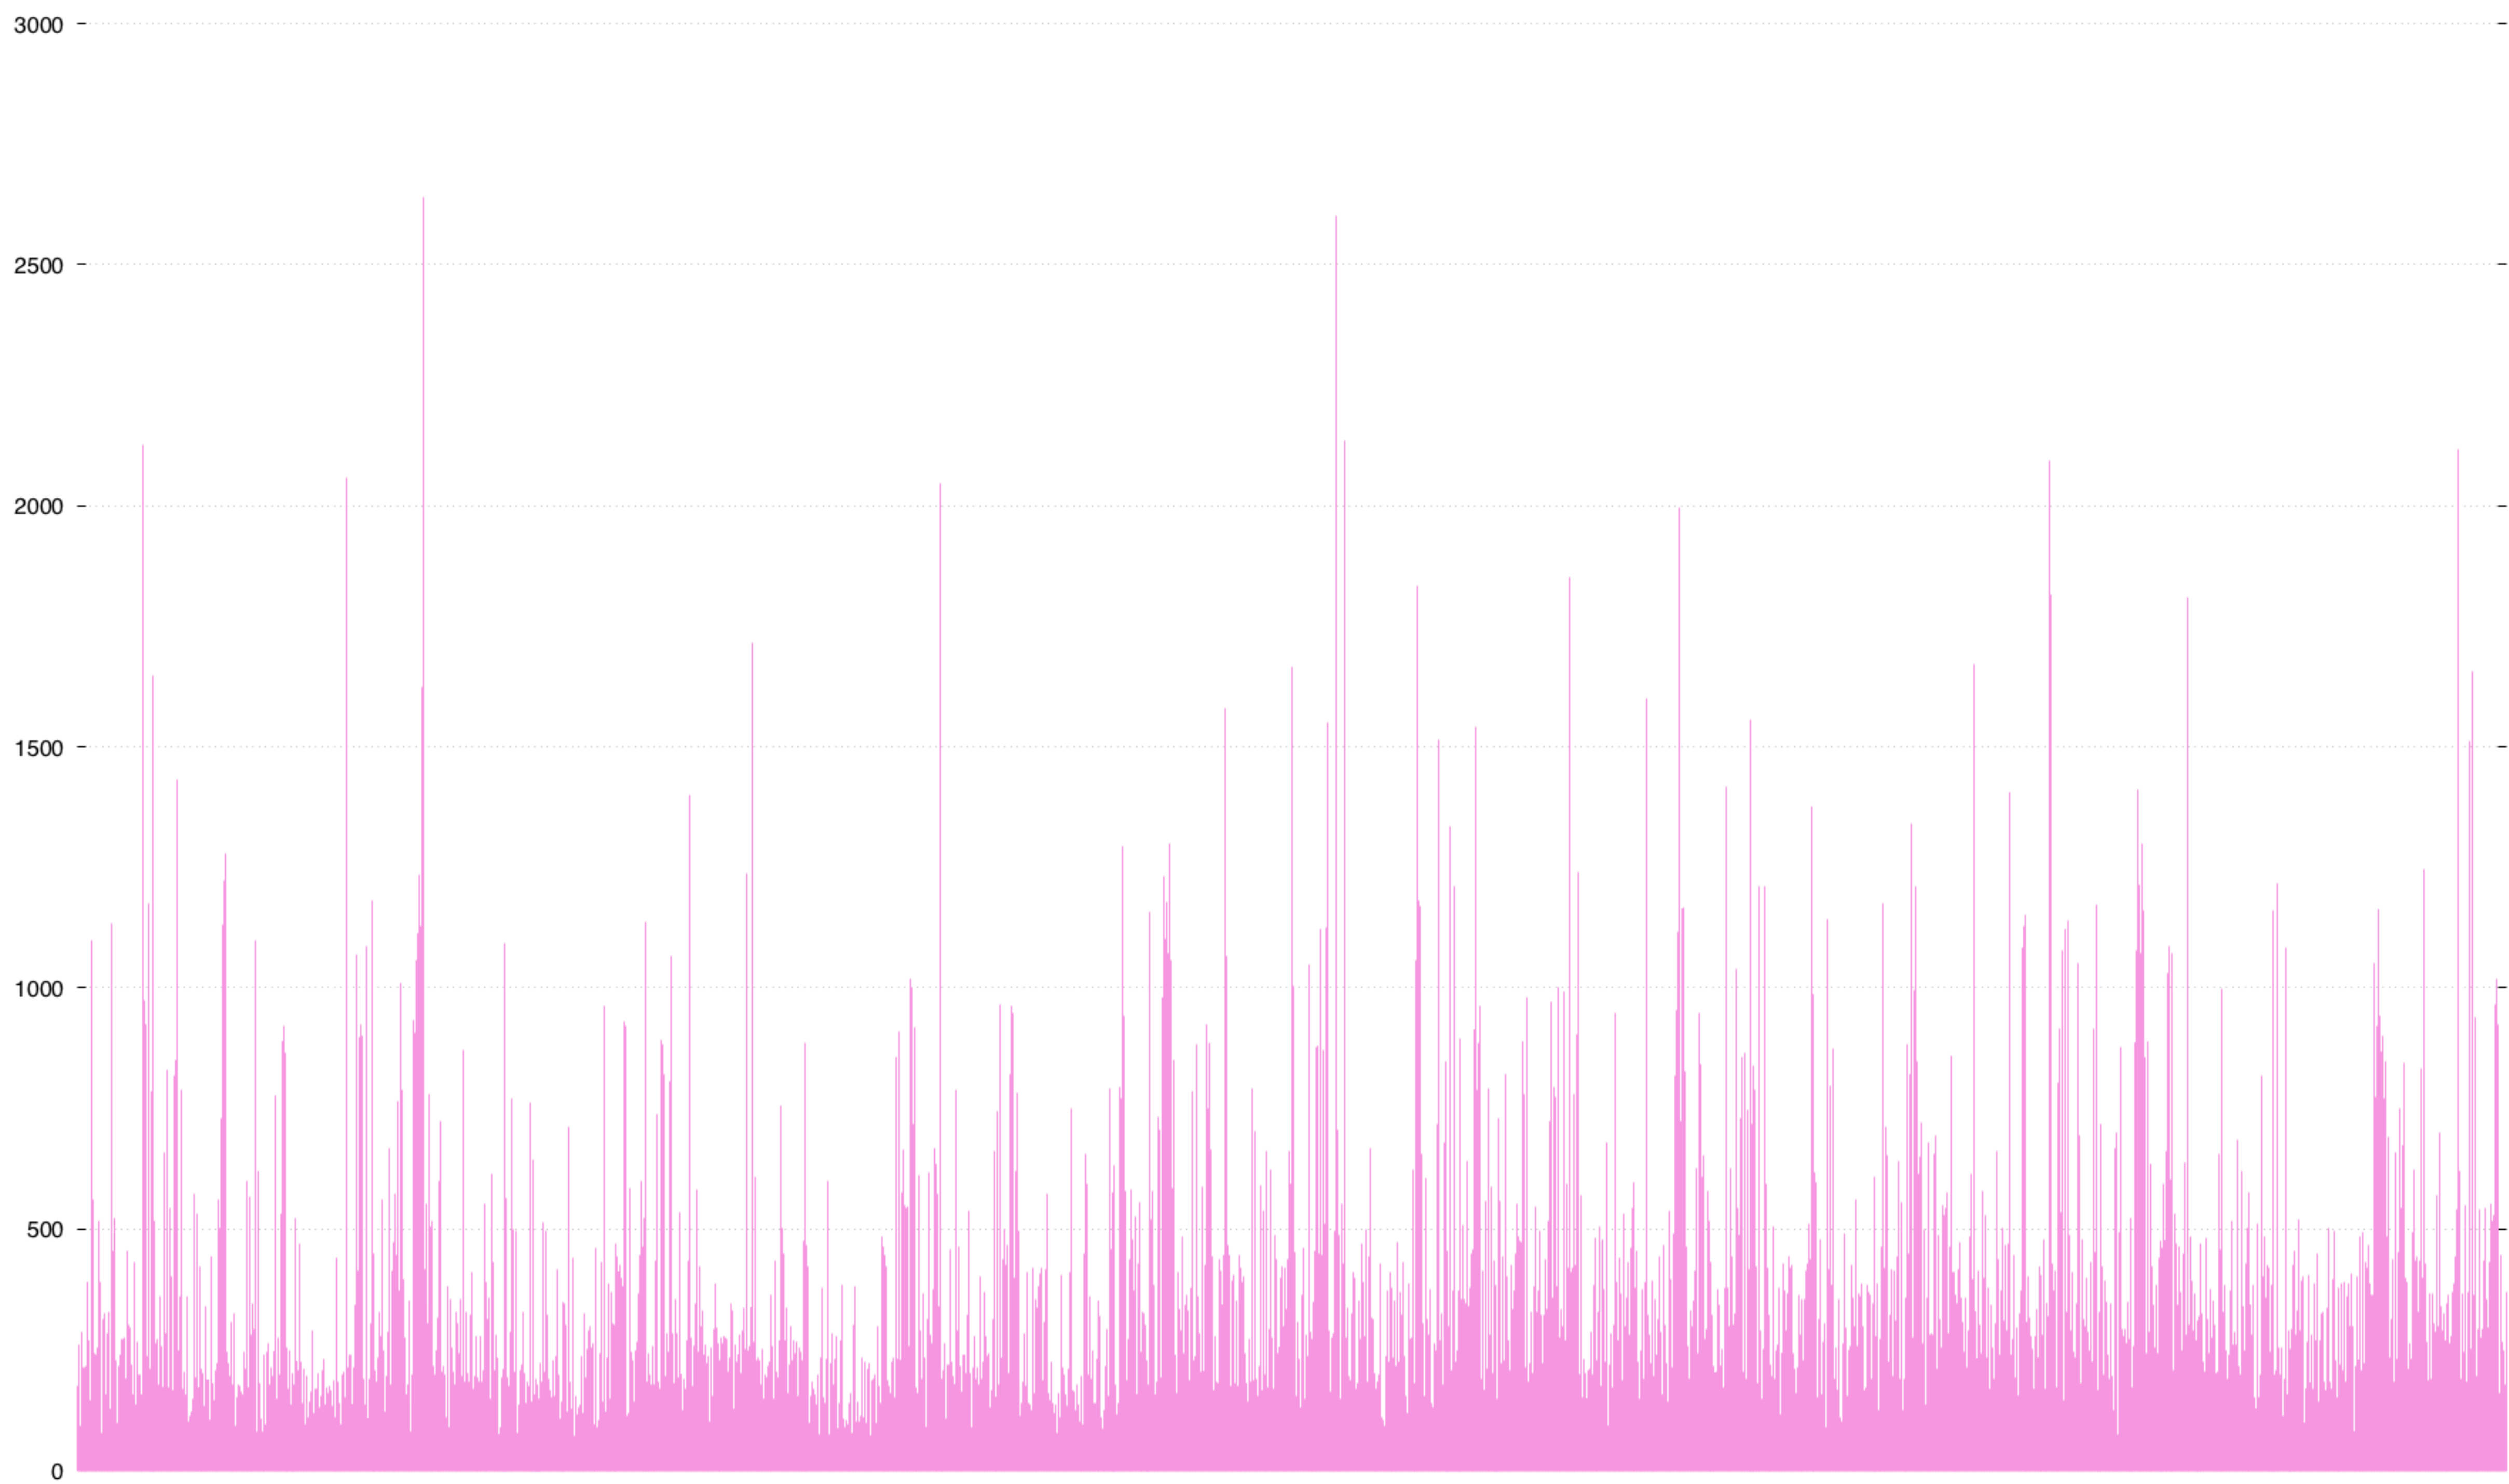
75.57.176.28 - - [20/May/2010:10:52:42 +0000] "GET /media/exmpl.png HTTP/1.1" 200 28479 "http://www.example.com/" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_5\_8; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

75.57.176.28 - - [20/May/2010:11:52:43 +0100] "GET /media/img/m-act.gif HTTP/1.1" 200 143 "http://www.example.com/" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10\_5\_8; en-us) AppleWebKit/531.22.7 (KHTML, like Gecko) Version/4.0.5 Safari/531.22.7"

66.249.65.40 - - [20/May/2010:11:53:34 +0100] "GET /post/274702/ HTTP/1.1" 404 15 "-" "Mozilla/5.0 (compatible; Googlebot/2.1; +http://www.google.com/bot.html)"

67.195.114.50 - - [20/May/2010:11:54:36 +0100] "GET /post/256204/ HTTP/1.0" 404 15 "-" "Mozilla/5.0 (compatible; Yahoo! Slurp/3.0; http://help.yahoo.com/help/us/ysearch/slurp)"

66.249.65.40 - - [20/May/2010:11:54:39 +0100] "GET /post/274873/ HTTP/1.1" 404 15 "-" "Mozilla/5.0 (compatible; Googlebot/2.1; +http://www.google.com/bot.html)"



count

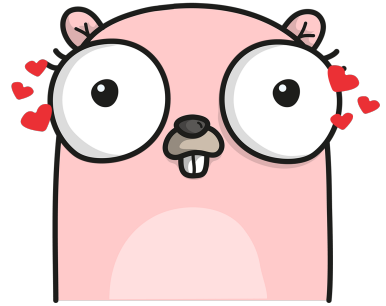
p50

p99

max

histogram

app

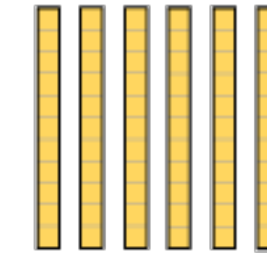


events

{ k: v }



column store



analytical  
queries

you are here



SELECT ...  
GROUP BY



users



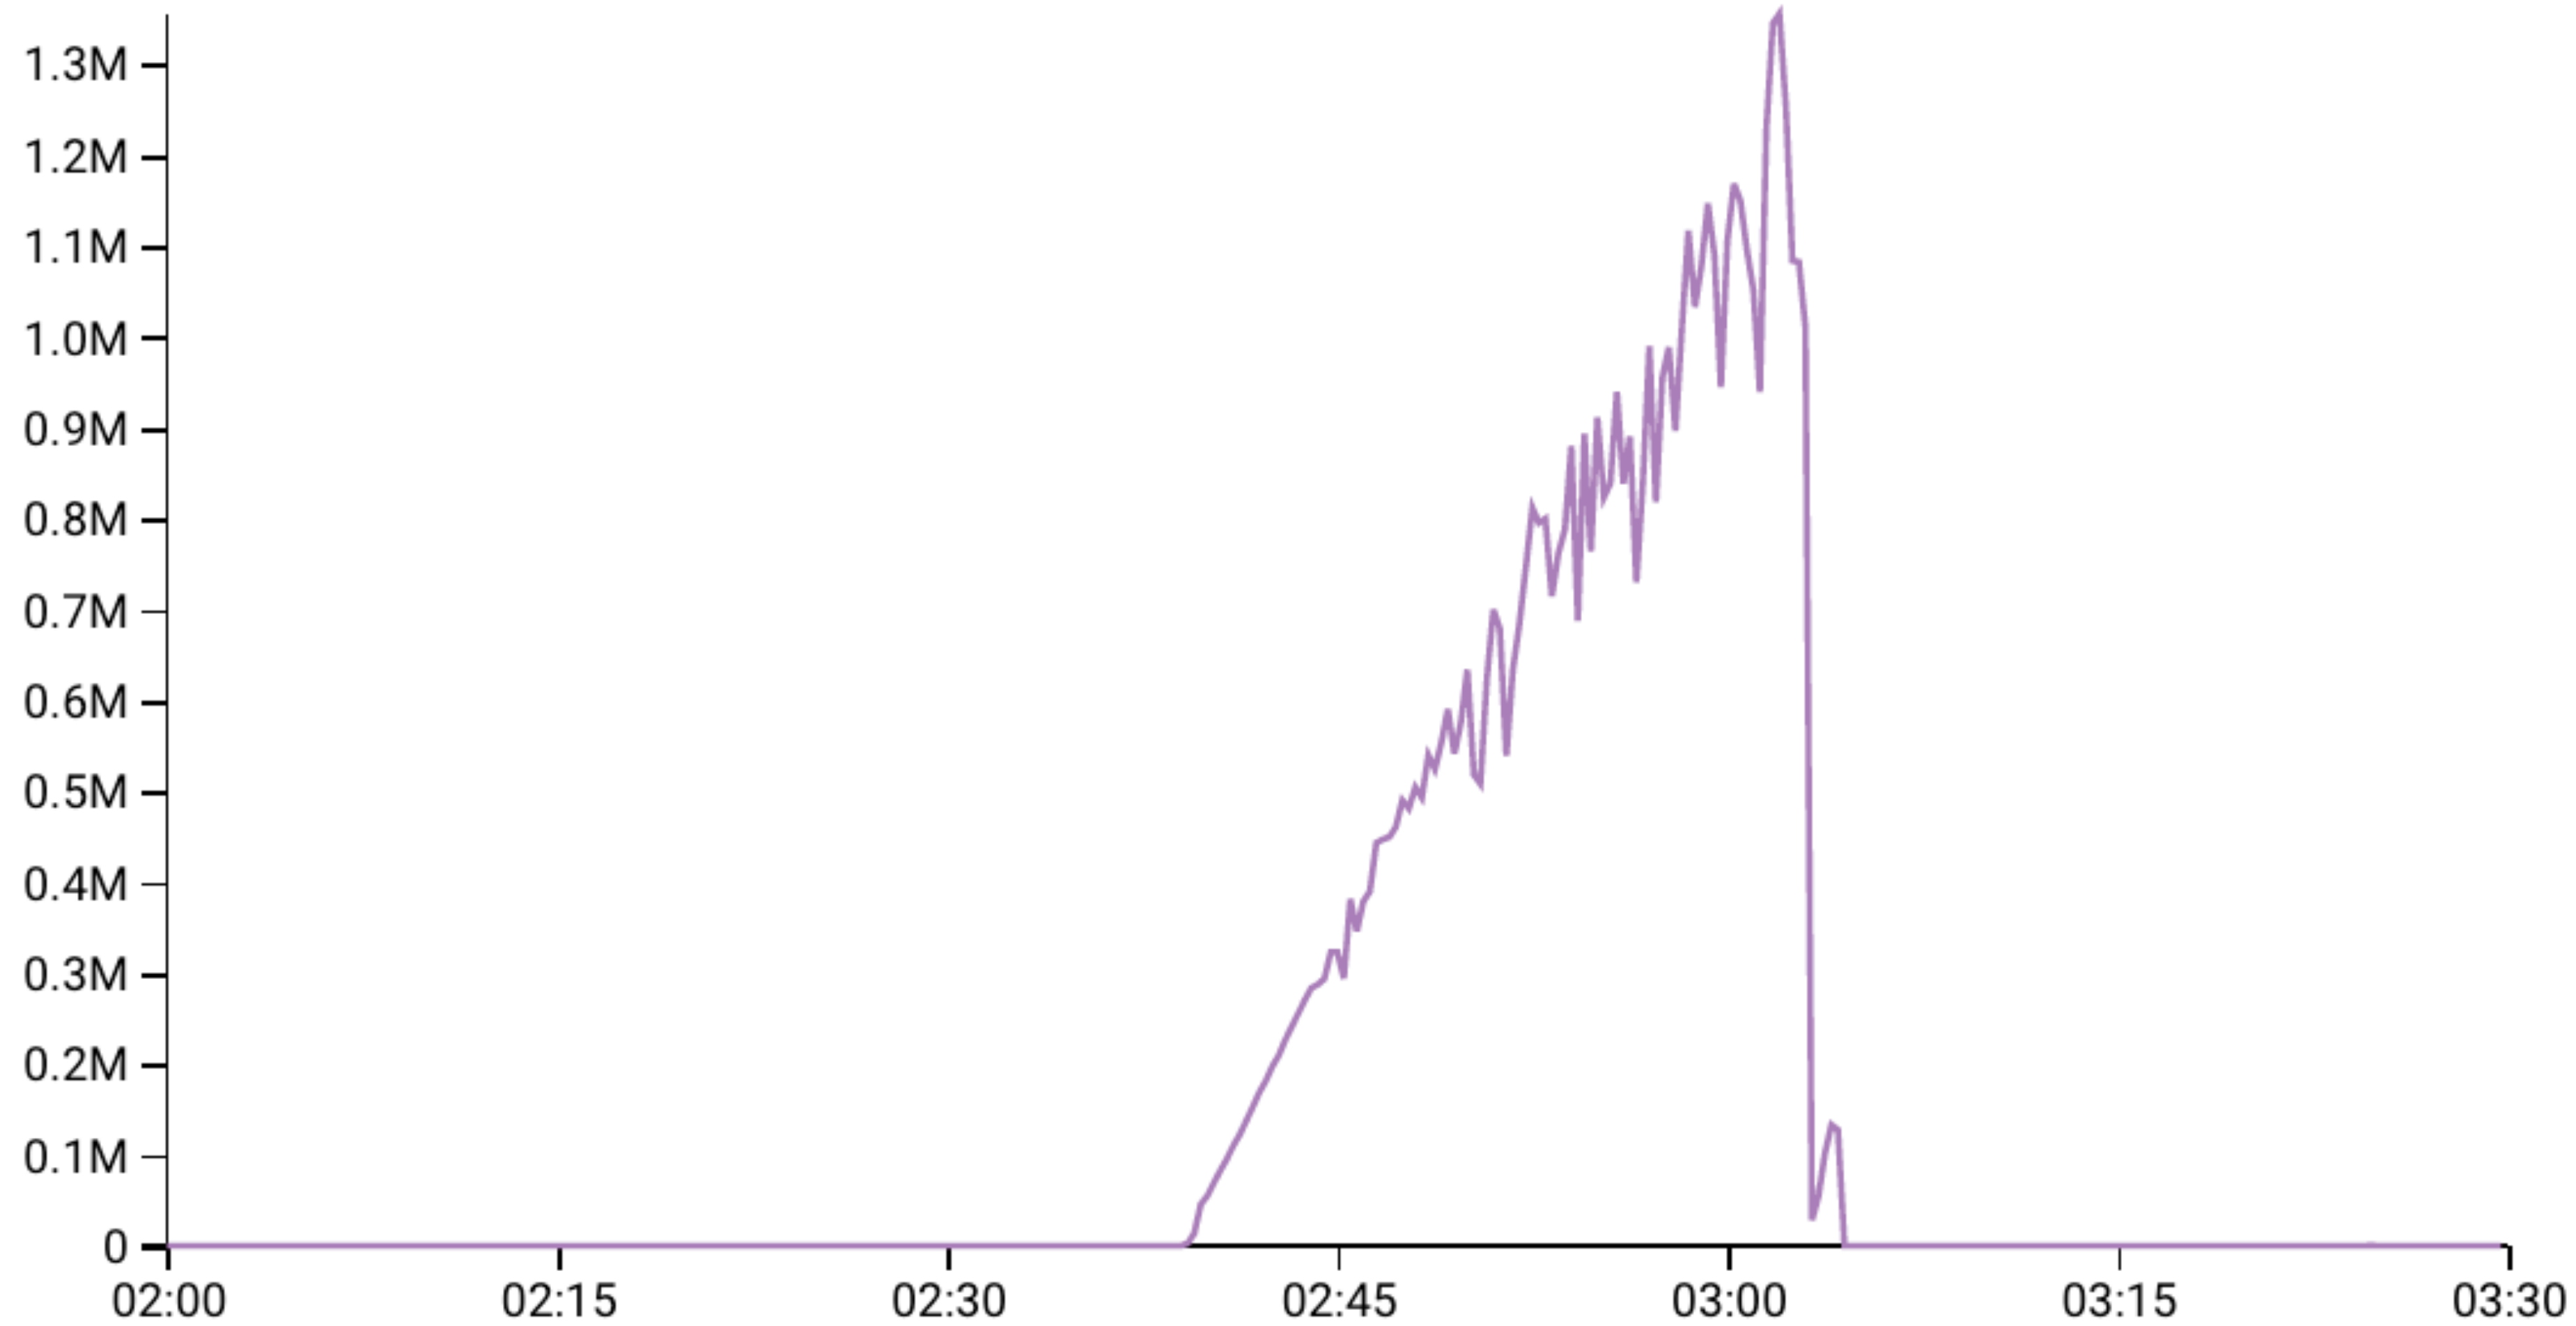
prometheus and the  
problem with metrics





"it's slow"

**P99(duration\_ms)**




**p99(request\_latency)  
> 1000ms**

300 requests were slow  
... which ones?!

group by

most monitoring questions are

✨ top-k 🏆

 top traffic by IP address

 top resource usage by customer

 top latency by country

 top error count by host

 top request size by client

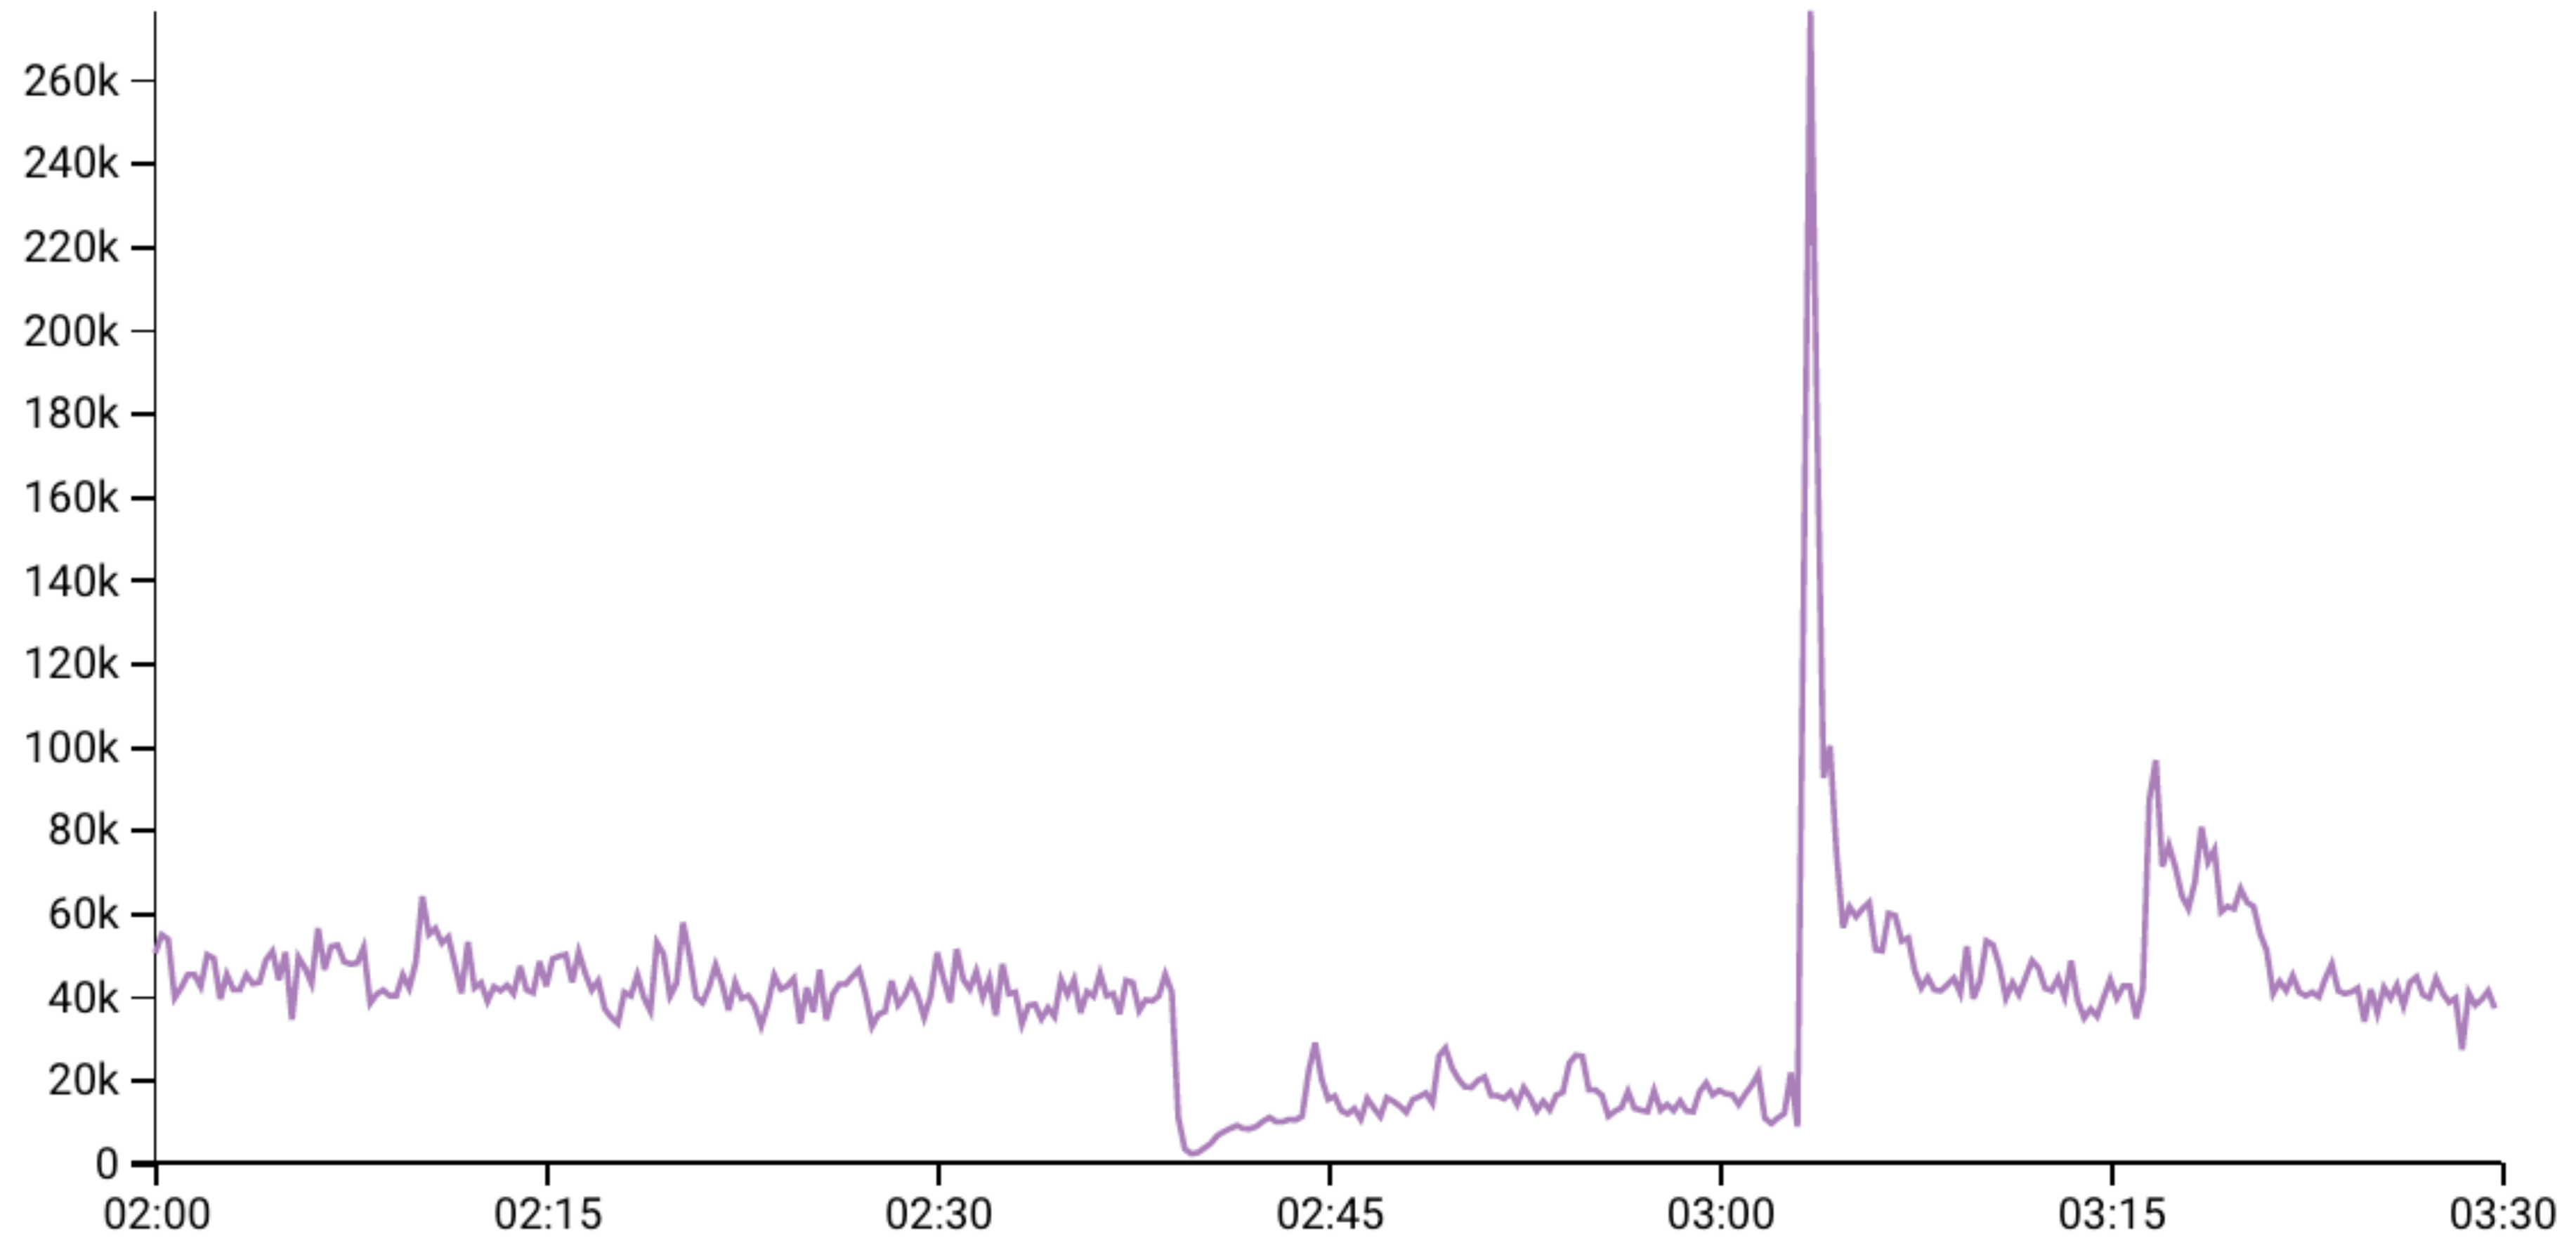
how many users  
are impacted?

```
SELECT user_id, COUNT(*)  
FROM requests  
WHERE request_latency >= 1000  
GROUP BY user_id
```

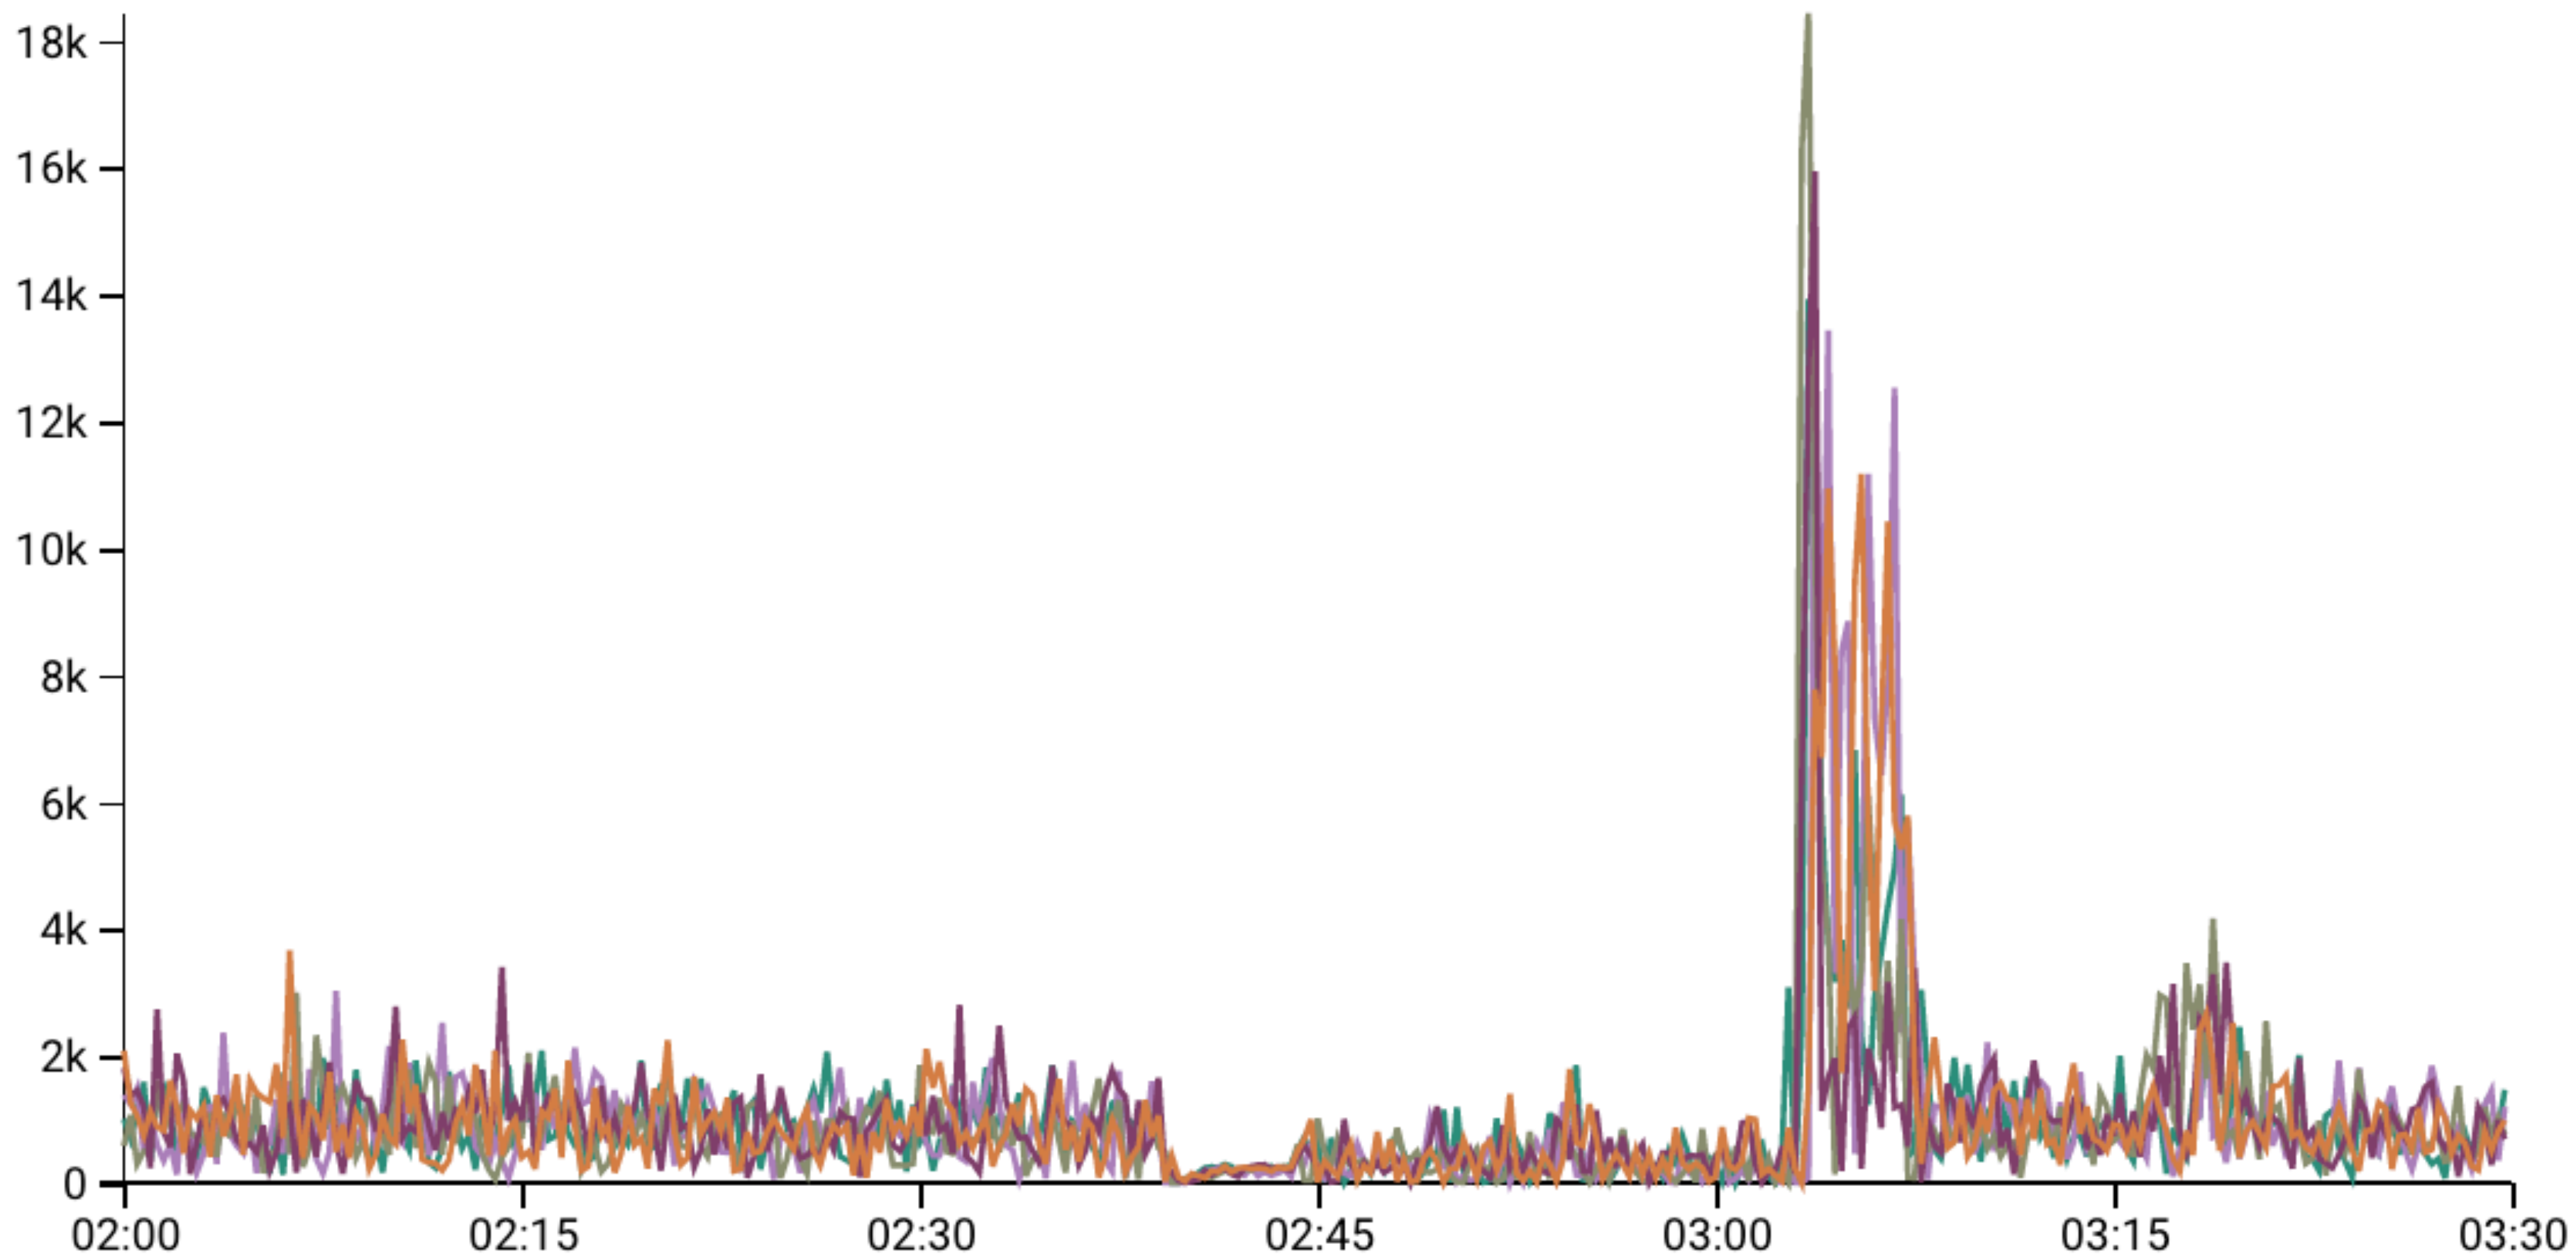
metrics will not  
tell you this

✨✨ cardinality 🌙

**COUNT**



COUNT



`http_requests_total{status=200}`

10

`http_requests_total{status=201}`

`http_requests_total{status=301}`

`http_requests_total{status=304}`

`...`

`http_requests_total{status=503}`

user\_id 10k

**100k**

ip address space =  $2^{32}$

4 billion possible values

```
kubectl get pods 100
```

build\_id

100

the curse of  
dimensionality

```
{
  status          = 200
  method          = "GET"
  path            = ...
  host            = "i-123456af"
  zone            = "eu-central-1a"
  client_ip       = "10.2.3.4"
  user_agent      = "curl/7.54.0"
  client_country  = "de"

  user_id         = 30032
  partition_id    = 31
  build_id        = "9045e1"
  customer_plan   = "platinum"
  endpoint        = "tweet_detail"
}
```

```
{
  status      = 200      10
  method      = "GET"    5
  path        = ...      300
  host        = "i-123456" 20
  zone        = "eu-central" 5
  client_ip    = "10.2.3.4" 1k
  user_agent   = "curl/7.5" 300
  client_country = "de"   20

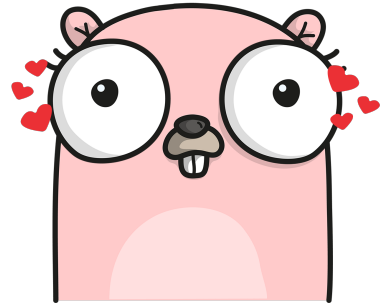
  user_id      = 30032    1k
  partition_id = 31       32
  build_id     = "9045e1" 10
  customer_plan = "platinum" 3
  endpoint     = "tweet_delete" 20
}
```

$$\begin{array}{ccccccccc}
 10 & \times & 5 & \times & 300 & \times & 20 & \times & 5 & \times \\
 1k & \times & 300 & \times & 20 & \times & 1k & \times & 32 & \times \\
 10 & \times & 3 & \times & 20 & & & & & 
 \end{array}$$

$$= 172'800'000'000 \\
 000'000'000 \text{ 💣}$$

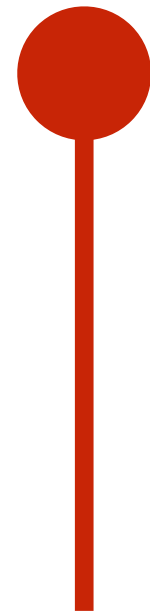


app



events

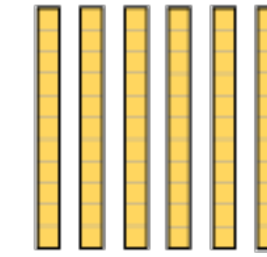
{ k: v }



you are here

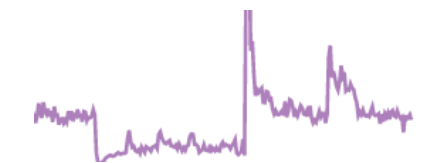


column store



analytical  
queries

SELECT ...  
GROUP BY



users



recording events



**Liz Fong-Jones (方禮真)** ✓

@lizthegrey

```
void do_rpc() {  
    ...  
    record_event({  
        x: x,  
        y:y,  
        status: status,  
        user: user,  
        version: version,  
        ...  
    })  
}
```

```
{
  time          = "1970-01-01T18:32:20"

  status        = 200
  method        = "GET"
  path          = ...
  host          = "i-123456af"
  region        = "eu-central-1"
  zone          = "eu-central-1a"
  client_ip     = "10.2.3.4"
  user_agent    = "curl/7.54.0"
  client_country = "de"
  kernel        = "5.0.0-1018-aws"

  user_id       = 30032
  tweet_id      = 2297111098
  partition_id  = 31
  build_id      = "9045e1"
  request_id    = "f2a3bdc4"
  customer_plan = "platinum"
  feature_blub  = true
  cache         = "miss"
  endpoint      = "tweet_detail"

  request_dur_ms = 325
  db_dur_ms      = 5
  db_pool_dur_ms = 3
  db_query_count = 63
  cache_dur_ms   = 2
  svc_a_dur_ms   = 32
  svc_b_dur_ms   = 90
  request_bytes  = 2456
  response_bytes = 5324
}
```

```
{
  time          = "1970-01-01T18:32:20"

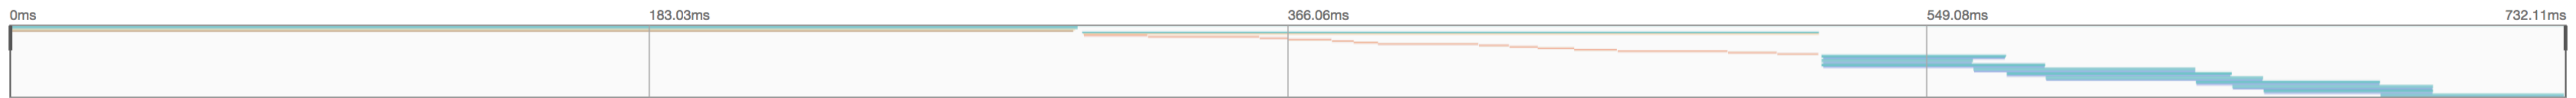
  status        = 200
  method        = "GET"
  path          = ...
  host          = "i-123456af"
  region        = "eu-central-1"
  zone          = "eu-central-1a"
  client_ip     = "10.2.3.4"
  user_agent    = "curl/7.54.0"
  client_country = "de"
  kernel       = "5.0.0-1018-aws"
}
```

```
{  
  user_id      = 30032  
  tweet_id     = 2297111098  
  partition_id = 31  
  build_id     = "9045e1"  
  request_id   = "f2a3bdc4"  
  customer_plan = "platinum"  
  feature_blub = true  
  cache        = "miss"  
  endpoint     = "tweet_detail"  
}
```

```
{  
  request_dur_ms = 325  
  db_dur_ms      = 5  
  db_pool_dur_ms = 3  
  db_query_count = 63  
  cache_dur_ms   = 2  
  svc_a_dur_ms   = 32  
  svc_b_dur_ms   = 90  
  request_bytes  = 2456  
  response_bytes = 5324  
}
```

✓ frontend: HTTP GET /dispatch

Trace Start: **July 20, 2018 2:48 PM** | Duration: **732.11ms** | Services: **6** | Depth: **5** | Total Spans: **51**



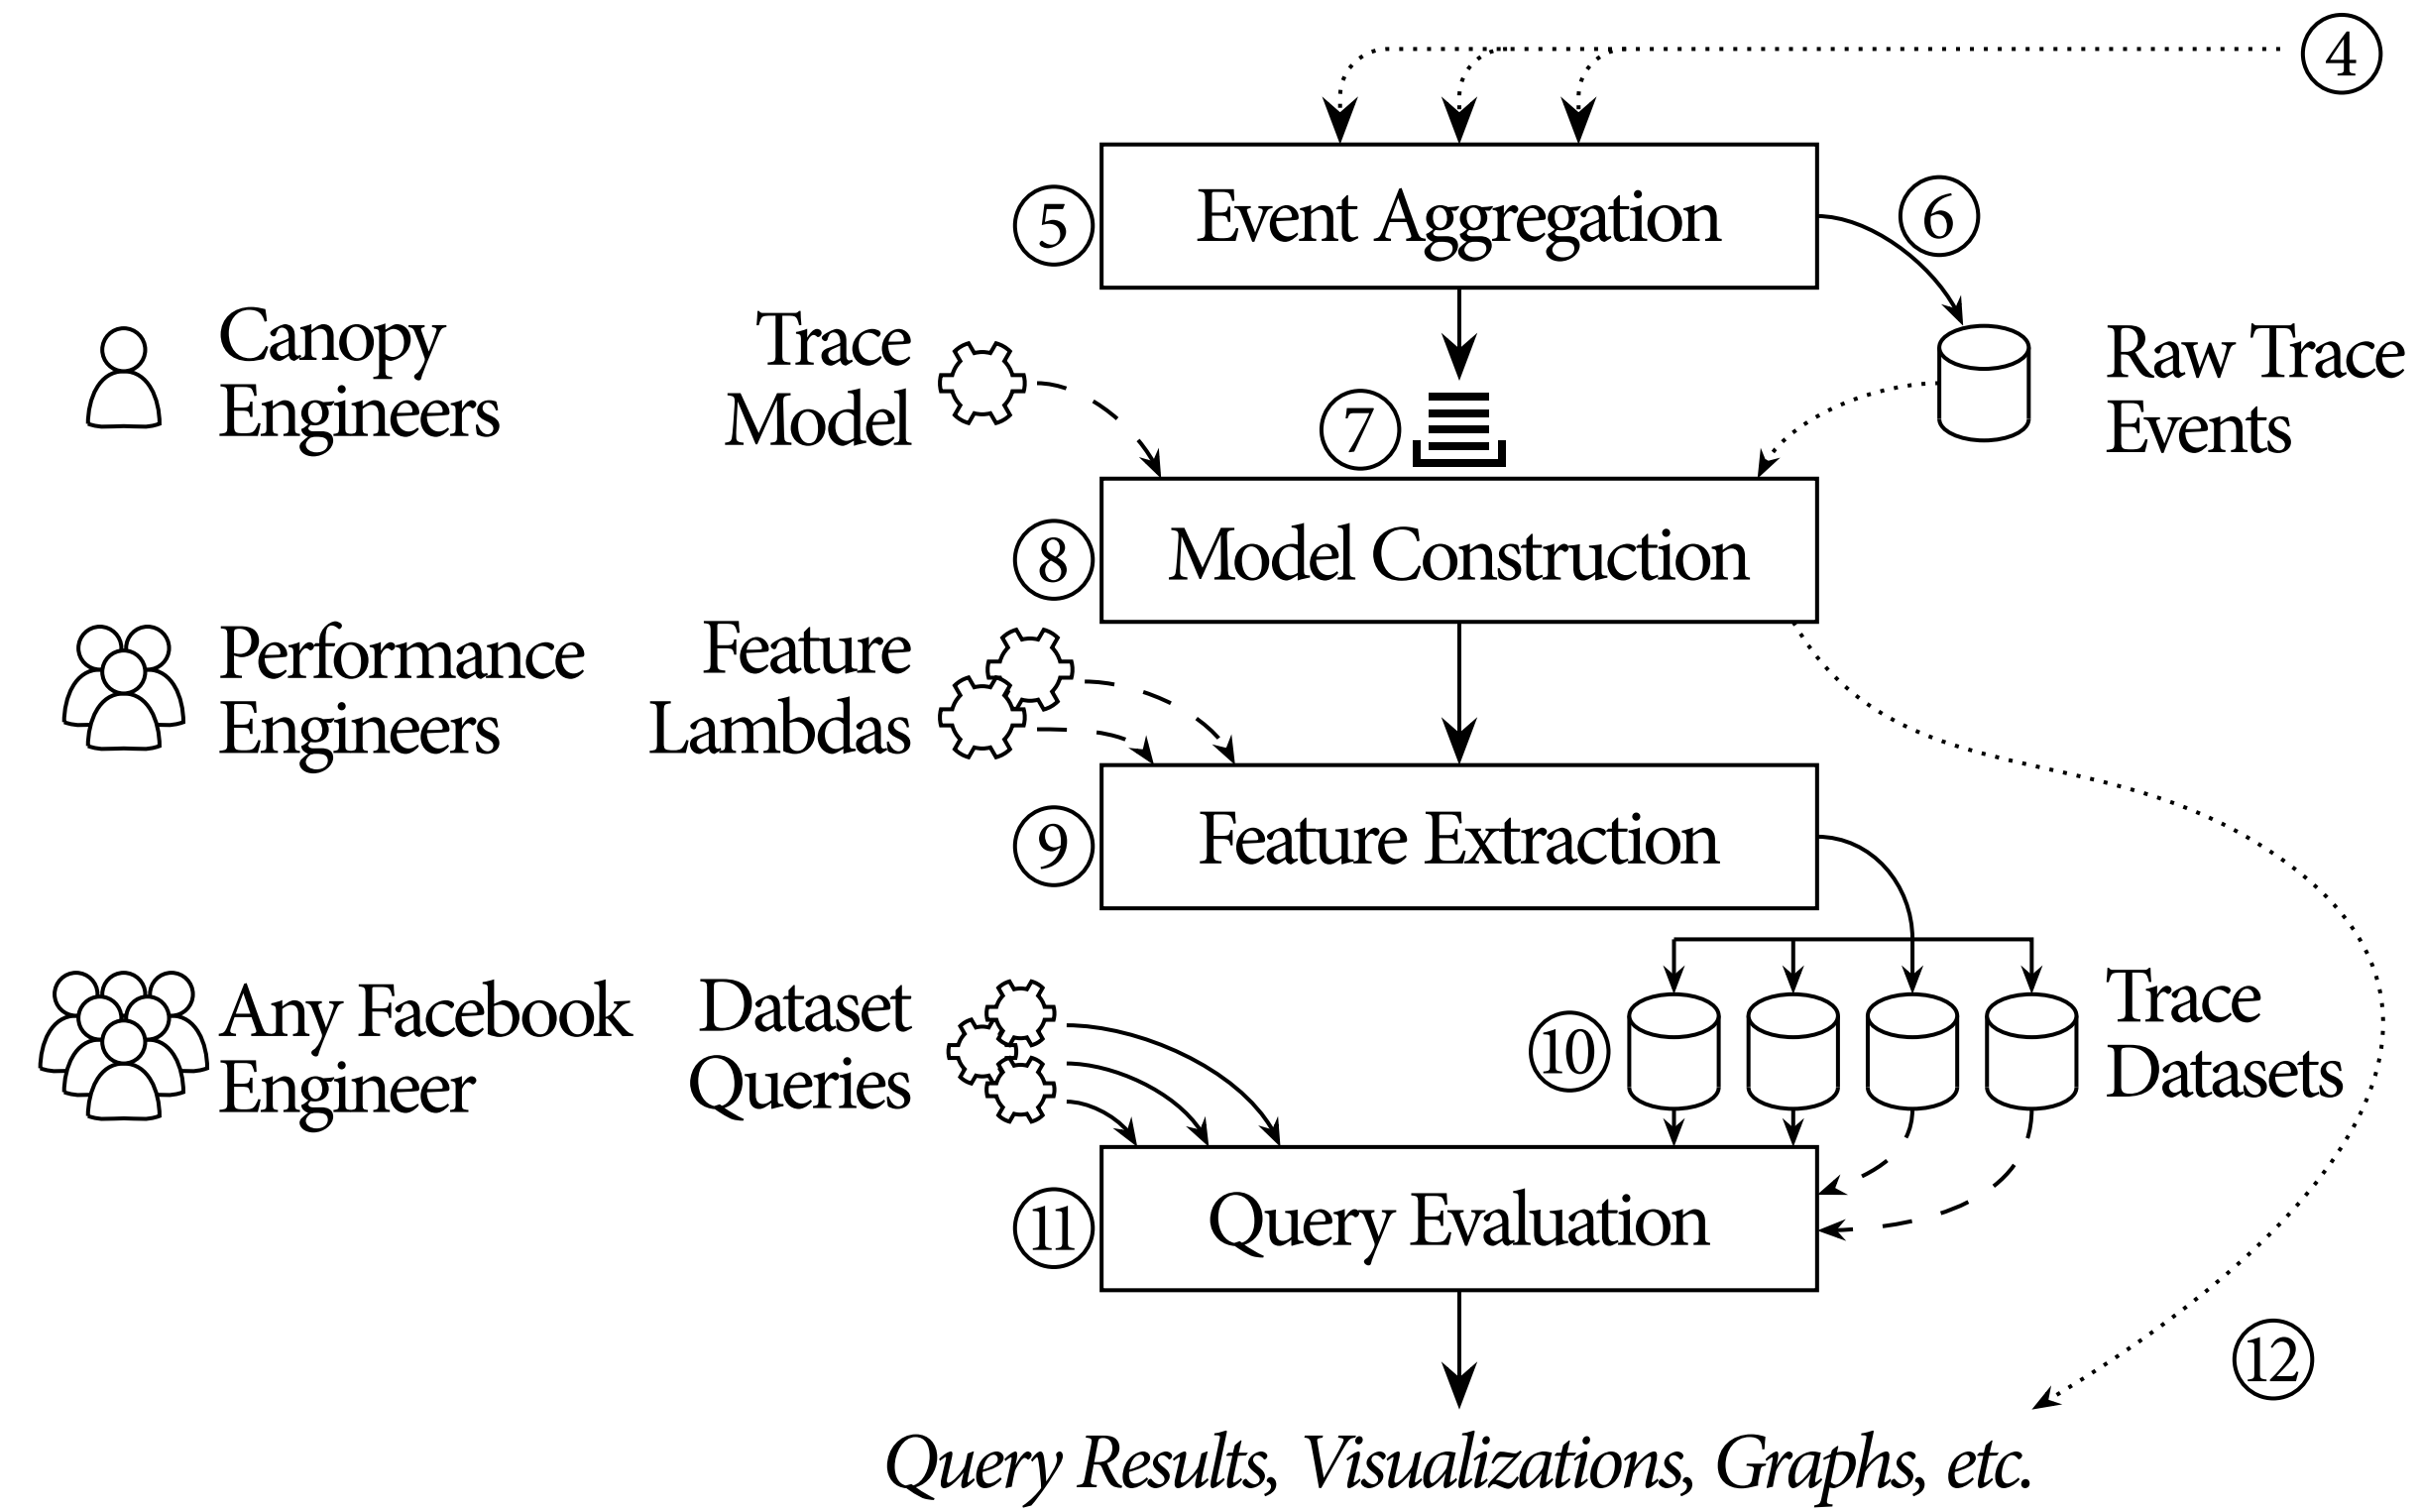
**Service & Operation**

0ms 183.03ms 366.06ms 549.08ms 732.11ms

Service & Operation	Start Time (ms)	End Time (ms)	Duration (ms)
frontend HTTP GET /dispatch	0	183.03	183.03
frontend HTTP GET: /customer	183.03	305.89	122.86
frontend HTTP GET	305.89	305.88	-0.01
customer HTTP GET /customer	305.88	304.46	-1.42
mysql SQL SELECT	304.46	304	-0.46
frontend Driver::findNearest	304	211.15	-92.85
driver Driver::findNearest	211.15	210.78	-0.37
redis FindDriverIDs	210.78	18.23	-192.55
redis GetDriver	18.23	32	13.77
redis GetDriver	32	8.17	-23.83
redis GetDriver	8.17	12.34	4.17
redis GetDriver	12.34	6.24	-6.10
redis GetDriver	6.24	6.96	0.72
redis GetDriver	6.96	28.81	21.85
redis GetDriver	28.81	8.77	-20.04
redis GetDriver	8.77	8.12	-0.65
redis GetDriver	8.12	10.49	2.37
redis GetDriver	10.49	12.4	1.91
redis GetDriver	12.4	31.63	19.23
redis GetDriver	31.63	14.02	-17.61
redis GetDriver	14.02	11.72	-2.30
frontend HTTP GET: /route	11.72	52.92	41.20
frontend HTTP GET	52.92	52.89	-0.03
route HTTP GET /route	52.89	52.08	-0.81
frontend HTTP GET: /route	52.08	43.45	-8.63
frontend HTTP GET	43.45	43.44	-0.01
route HTTP GET /route	43.44	42.48	-0.96
frontend HTTP GET: /route	42.48	64.18	21.70
frontend HTTP GET	64.18	64.15	-0.03
route HTTP GET /route	64.15	63.48	-0.67
frontend HTTP GET: /route	63.48	63.67	0.19
frontend HTTP GET	63.67	63.65	-0.02
route HTTP GET /route	63.65	63.36	-0.29

traces vs events:  
a false dichotomy

we can derive events  
from traces



# Canopy: An End-to-End Performance Tracing And Analysis System

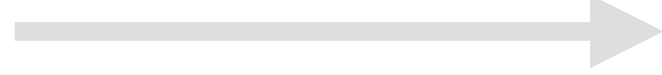
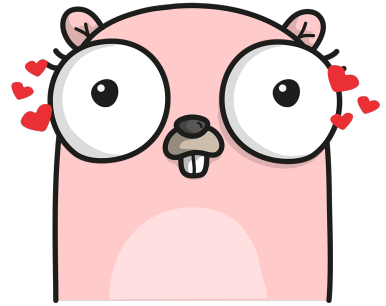
Jonathan Kaldor<sup>†</sup>   Jonathan Mace<sup>\*</sup>   Michał Bejda<sup>†</sup>   Edison Gao<sup>†</sup>   Wiktor Kuropatwa<sup>†</sup>  
Joe O'Neill<sup>†</sup>   Kian Win Ong<sup>†</sup>   Bill Schaller<sup>†</sup>   Pingjia Shan<sup>†</sup>   Brendan Viscomi<sup>†</sup>  
Vinod Venkataraman<sup>†</sup>   Kaushik Veeraraghavan<sup>†</sup>   Yee Jiun Song<sup>†</sup>

<sup>†</sup>Facebook

<sup>\*</sup>Brown University

stick those events in kafka

app

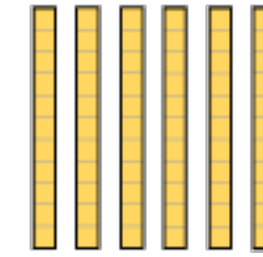


events

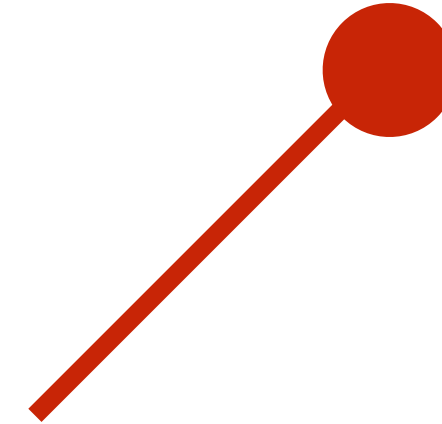
{ k: v }



column store



you are here



analytical  
queries

SELECT ...  
GROUP BY



users



columnar storage  
changed my life

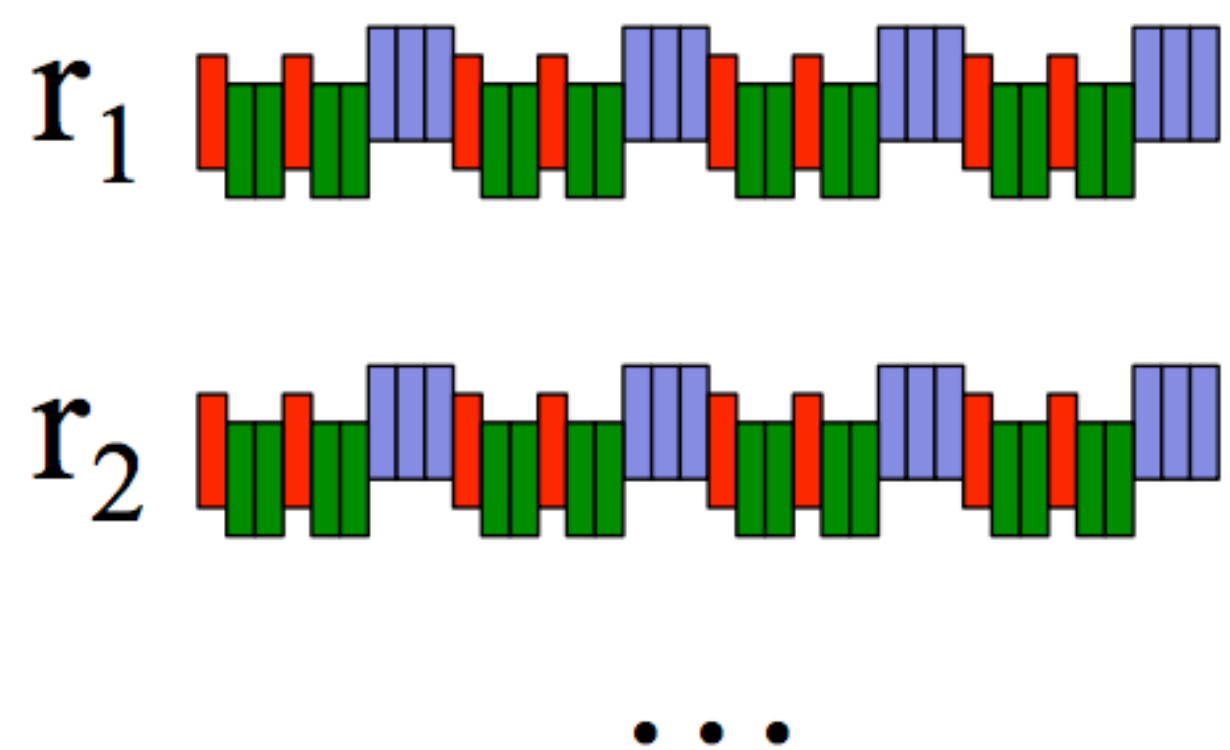
<b>Table 2.3:</b> Latency numbers that every WSC engineer should know. (Updated version of table from [Dea09].)	
Operation	Time
L1 cache reference	1.5 ns
L2 cache reference	5 ns
Branch misprediction	6 ns
Uncontended mutex lock/unlock	20 ns
L3 cache reference	25 ns
Main memory reference	100 ns
Decompress 1 KB with Snappy [Sna]	500 ns
“Far memory”/Fast NVM reference	1,000 ns (1us)
Compress 1 KB with Snappy [Sna]	2,000 ns (2us)
Read 1 MB sequentially from memory	12,000 ns (12 us)
SSD Random Read	100,000 ns (100 us)
Read 1 MB bytes sequentially from SSD	500,000 ns (500 us)
Read 1 MB sequentially from 10Gbps network	1,000,000 ns (1 ms)
Read 1 MB sequentially from disk	10,000,000 ns (10 ms)
Disk seek	10,000,000 ns (10 ms)
Send packet California→Netherlands→California	150,000,000 ns (150 ms)

- 1TB Hitachi Deskstar 7K1000
  - disk seek time = 14ms
  - transfer rate = 69MB/s
- 62.5 billion rows (= 1TB / 16 bytes)
- 28 years (= 62.5 billion rows \* 14 ms/row /  $32 \times 10^9$  ms/year)

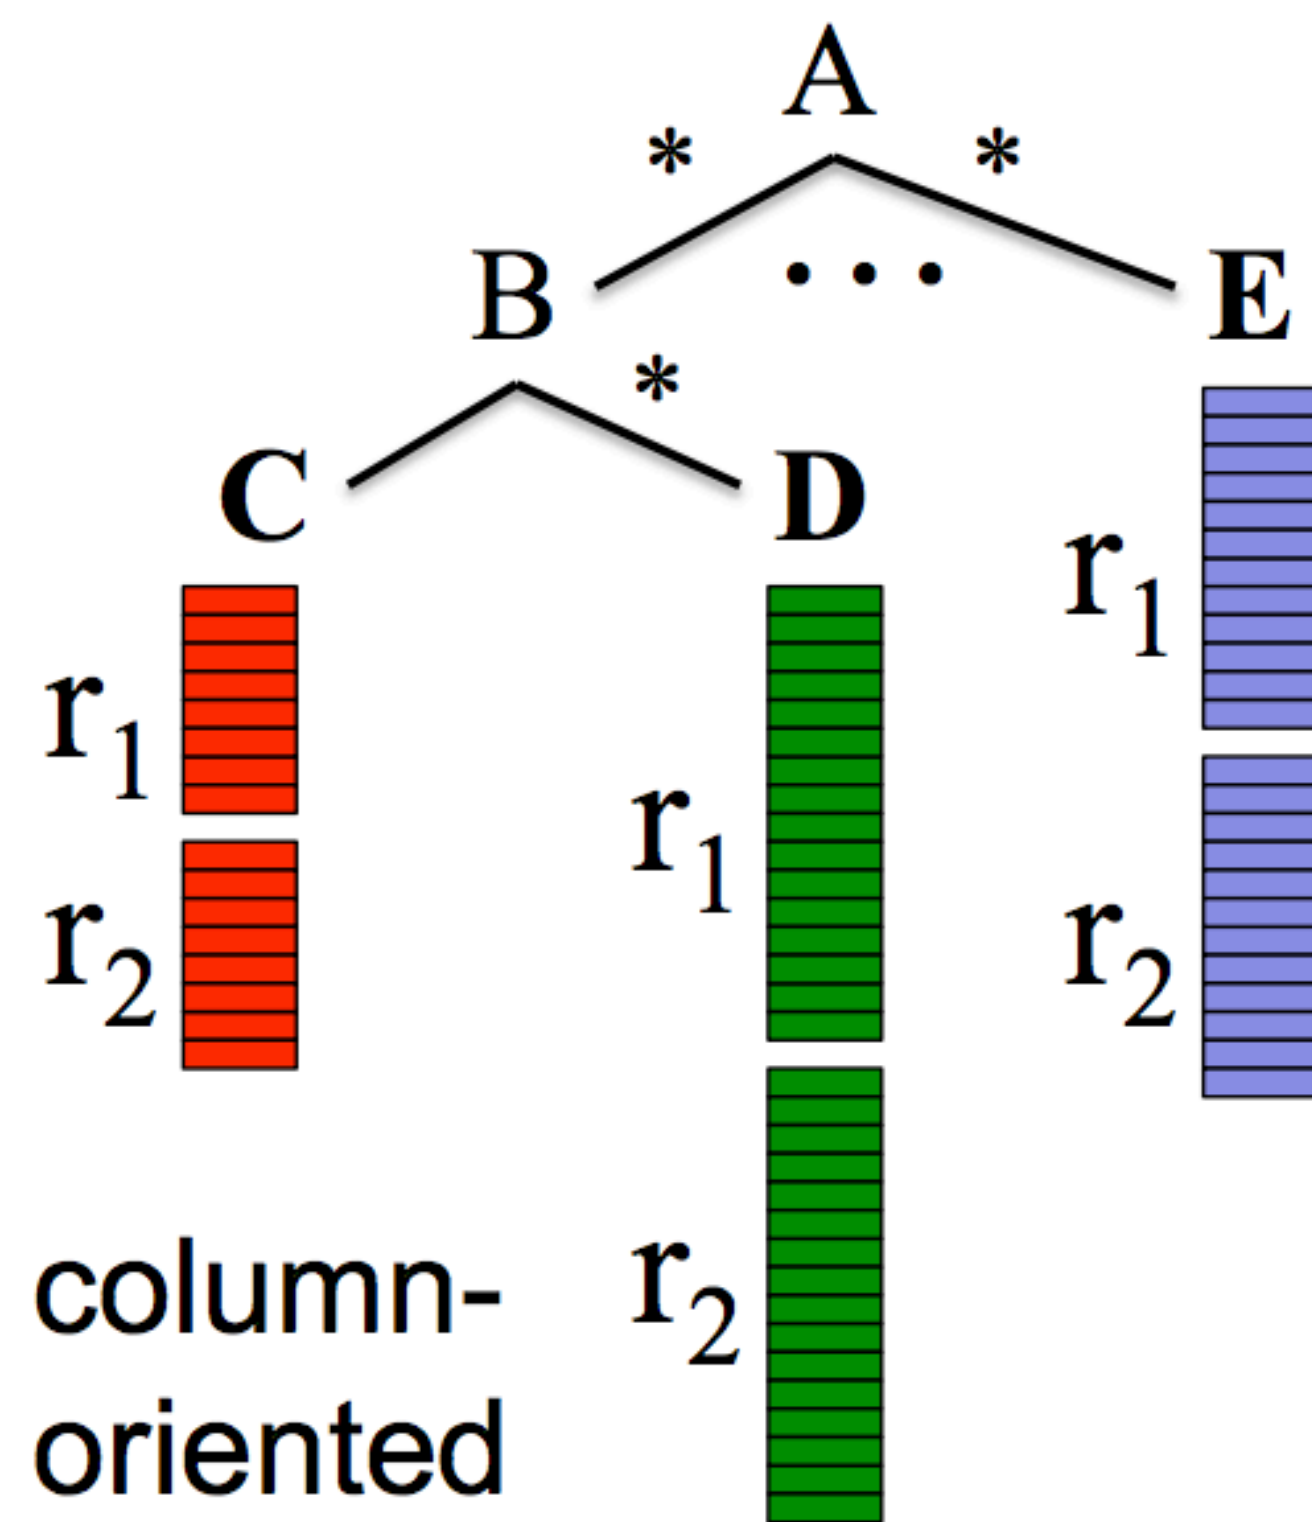
- 1TB Hitachi Deskstar 7K1000
- transfer rate = 69MB/s
- 4 hours (= 1.000.000MB / 69MB/s / 3600 s/hour)

- SSD
- transfer rate = 1GB/s
- 15 minutes (= 1.000GB / 1GB/s / 60 s/min)

10GB



record-oriented



column-oriented

10 GB / 8 bytes per data point

= 1.3 billion

events

status

200

200

200

200

404

200

200

200

404

200

status

4 \* 200

404

3 \* 200

404

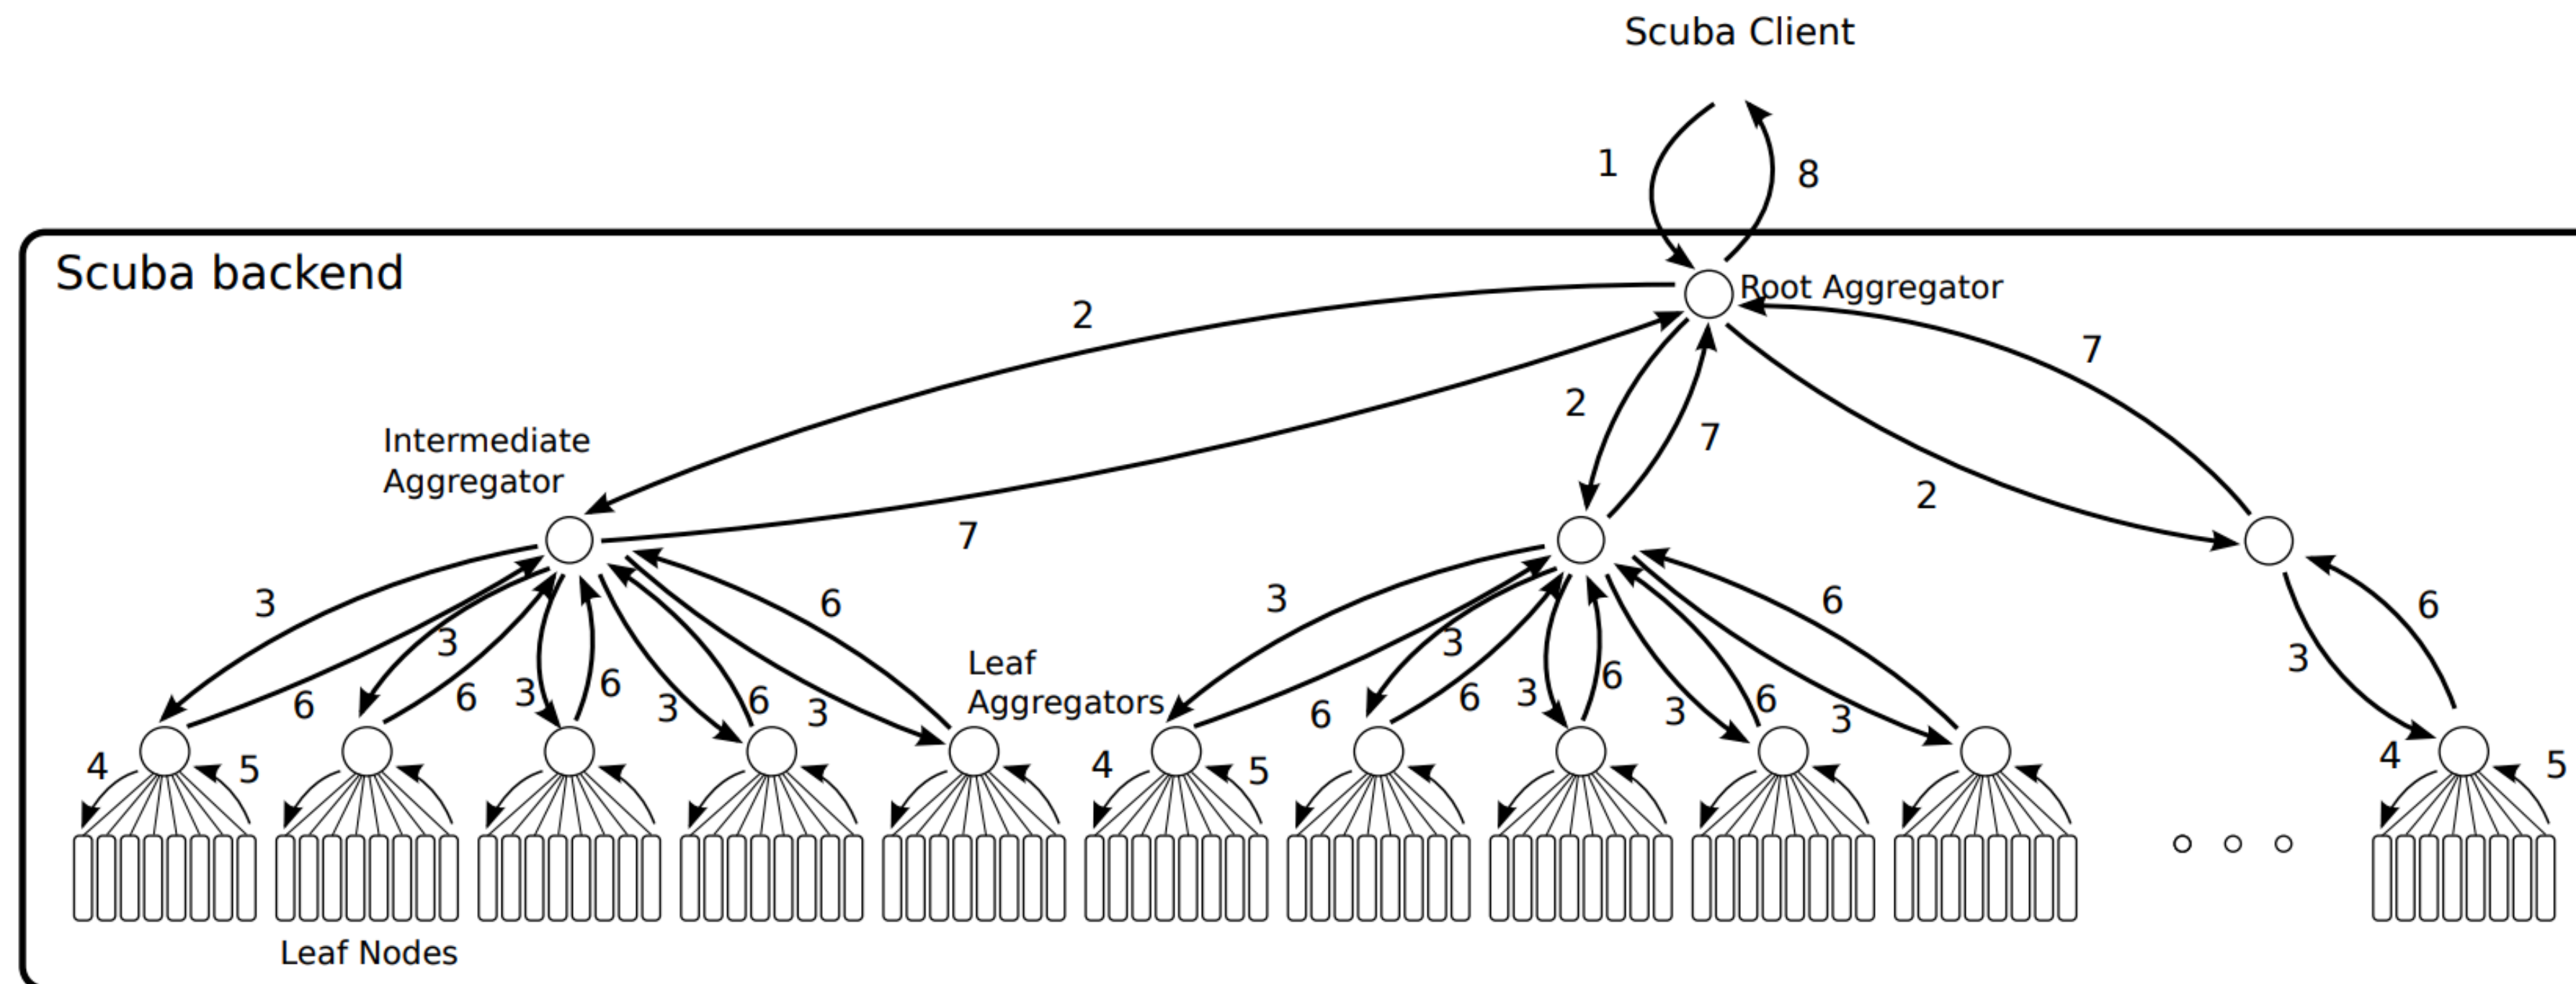
200

time-based partitioning

dynamic sampling

it's lossy, but that's fine

vectorized processing



sequential scans



columnar layout



time-based partitioning



compression / sampling



vectorized processing

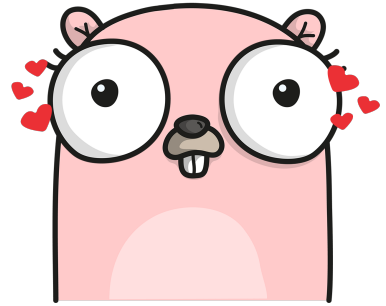


sharding



putting it all  
together

app

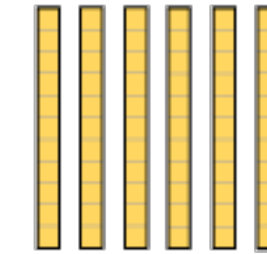


events

{ k: v }



column store



analytical  
queries

SELECT ...  
GROUP BY



users



we need more of this  
in the **m♡nit♡or♡ing** space!

```
SELECT user_id, COUNT(*)  
FROM requests  
WHERE status >= 500  
GROUP BY user_id  
ORDER BY COUNT(*) DESC  
LIMIT 10
```

✨✨ top-k 🏆

✨✨ cardinality 🌙

✨✨ events 🎆

# **Dremel: Interactive Analysis of Web-Scale Datasets**

Sergey Melnik, Andrey Gubarev, Jing Jing Long, Geoffrey Romer,  
Shiva Shivakumar, Matt Tolton, Theo Vassilakis  
Google, Inc.

{melnik, andrey, jlong, gromer, shiva, mtolton, theov}@google.com

# **Scuba: Diving into Data at Facebook**

Lior Abraham<sup>\*</sup>  
Vinayak Borkar  
Daniel Merl  
Subbu Subramanian

John Allen  
Bhuwan Chopra  
Josh Metzler  
Janet L. Wiener

Oleksandr Barykin  
Ciprian Gerea  
David Reiss  
Okay Zed

Facebook, Inc. Menlo Park, CA

- **Dremel: Interactive Analysis of Web-Scale Datasets** from Google, 2010
- **Scuba: Diving into Data at Facebook** from Facebook, 2016
- **Canopy: An End-to-End Performance Tracing And Analysis System** from Facebook, 2017
- **Look at Your Data** by John Rauser, Velocity 2011
- **Observability for Emerging Infra** by Charity Majors, Strange Loop 2017
- **Why We Built Our Own Distributed Column Store** by Sam Stokes, Strange Loop 2017
- **The Design and Implementation of Modern Column-Oriented Database Systems** by Abadi et al, 2013
- **Designing Data-Intensive Applications** by Martin Kleppmann, 2017
- **Monitoring in the time of Cloud Native** by Cindy Sridharan, 2017
- **Logs vs. metrics: a false dichotomy** by Nick Stenning, 2019
- **Using Canonical Log Lines for Online Visibility** by Brandur Leach, 2016
- **The Datacenter as a Computer: Designing Warehouse-Scale Machines** by Barroso et al, 2018



hi@igor.io

@igonwhilefalse