
SHRIEK!

Extreme Distributions in Postgres

Stefan Litsche


Zalando
SOLR
SHOP
MOBILE
LOUNGE
3rd PARTY
SERVICE 2
SERVICE 1
SERVICE 3
SERVICE 5
SERVICE 6
TOOLING
POSTGRE SQL
BI
LOGISTICS

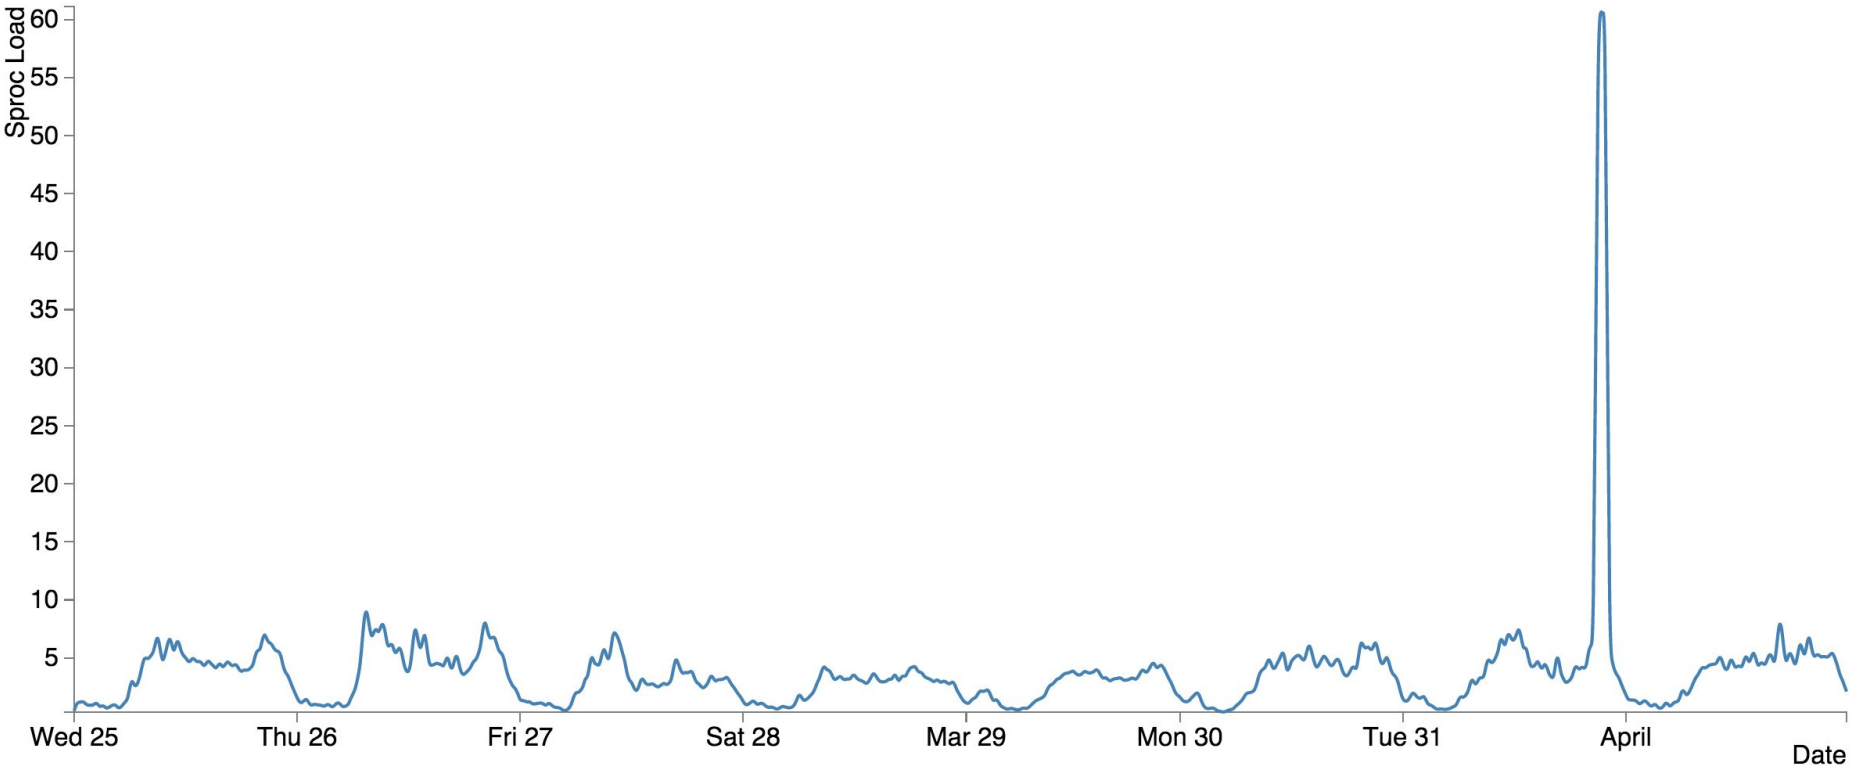
Sproc Load
60
55
50
45
40
35
30
25
20
15
10
5
Wed 25
Thu 26
Fri 27
Sat 28
Mar 29
Mon 30
Tue 31
April
Date

Phases of Statement Execution

- Parse
- Analyze
- Plan
- Execute

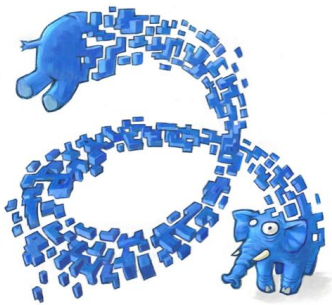

Planning execution

- Cost Model
- Cardinality Model
- Object Statistics

Cost Model

- random_page_cost
- seq_page_cost
- cpu_tuple_cost
- cpu_index_tuple_cost
- cpu_operator_cost

Cardinality Model

Size of intermediate Results

Based on object statistics

Uses probabilistic approach

Selectivity

column1 = a

S(a) = 1 / # keys in column1

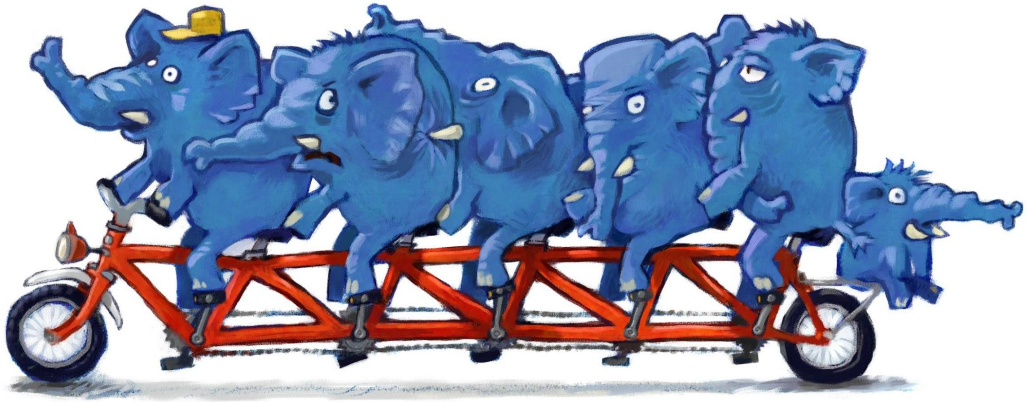

Independence of Predicates

c1 = a AND c2 = b

$P(a \cap b) = P(a) * P(b)$

- Dependencies in many domains
- Redundant Attributes

brand = “Honda” AND model = “Accord”

```
P  = 0.1 * 0.01
   = 0.001
```

Object Statistics

Distribution of data

Importance mentioned always

Auto analyze enabled by default

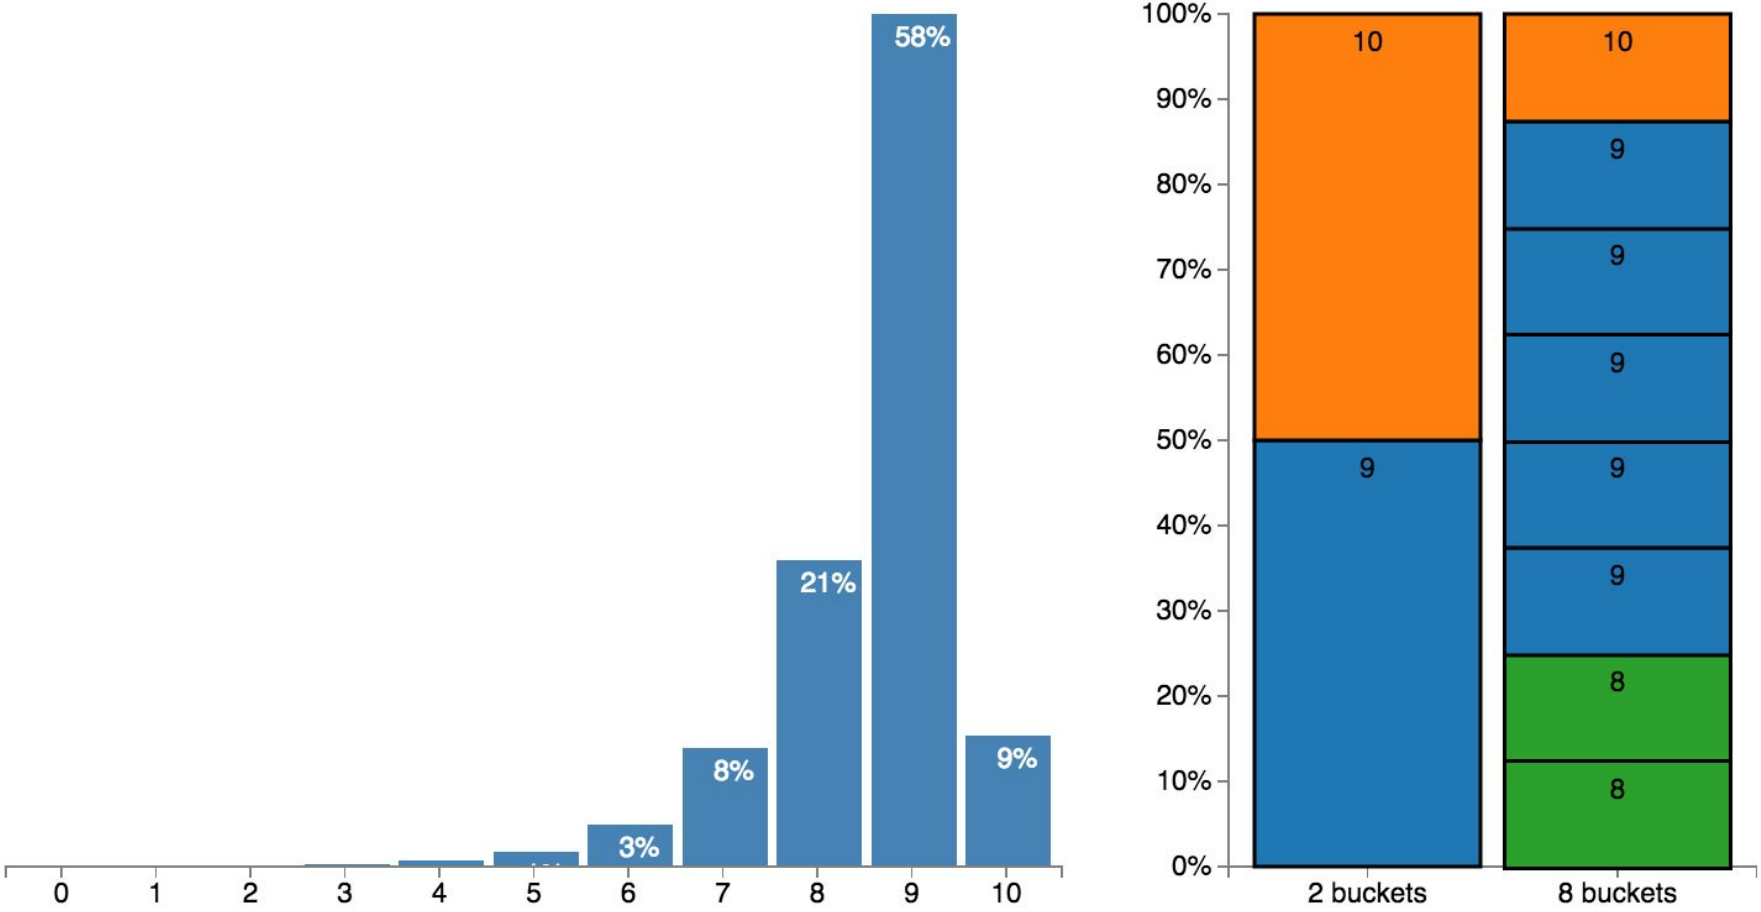

Mapping of Histogram to Buckets

View `pg_stats`

- uniform distributions
  - `histogram_bounds`
  - `for <, <= ...`
  - `not =`

- not uniform distributions
  - `most_common_freqs`
  - `for =, <, <= ...`

default_statistics_target

- number of rows inspected
- number of MCV
- number of buckets in histogram bounds

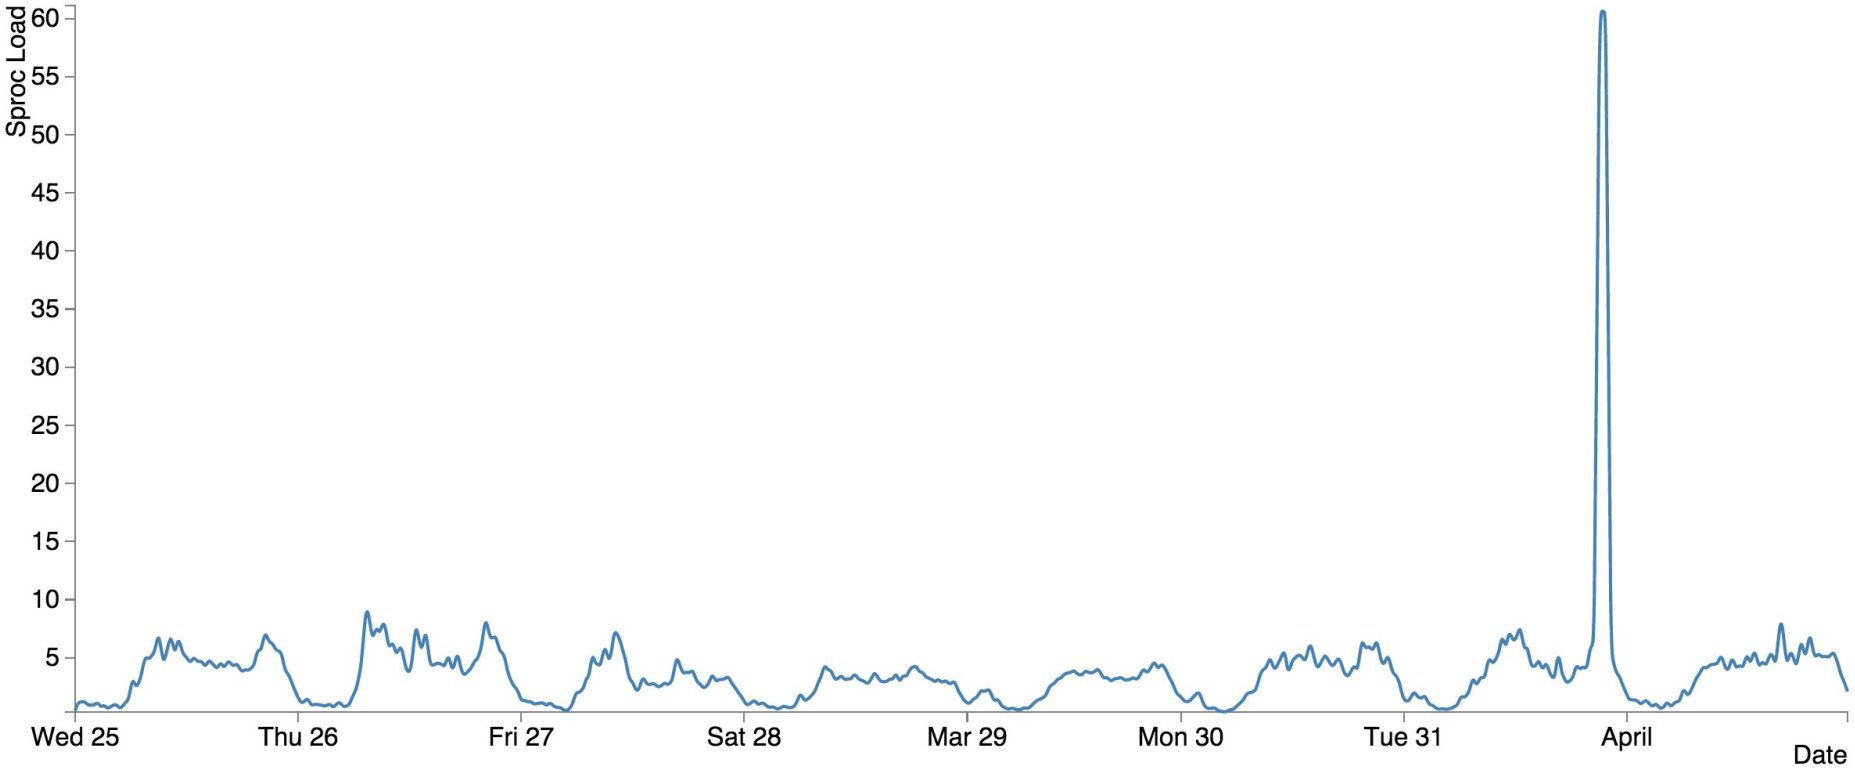
Sproc Load
60
55
50
45
40
35
30
25
20
15
10
5
Wed 25
Thu 26
Fri 27
Sat 28
Mar 29
Mon 30
Tue 31
April
Date

Forensics Before analyze

```
-[ RECORD 1 ]----------+-------------------------------------------------
tablename              | sli_testdata
attname                | log4
inherited              | f
null_frac              | 0
avg_width              | 4
n_distinct             | 4
most_common_vals       | {DONE_A}
most_common_freqs      | {0.900733}
histogram_bounds       | {NEW,DONE_B,DONE_B}
correlation            | 0.845195
```

Forensics After analyze

```
-[ RECORD 1 ]----------+-------------------------------------------------
tablename              | sli_testdata
attname                | log4
inherited              | f
null_frac              | 0
avg_width              | 4
n_distinct             | 4
most_common_vals       | {DONE_A,DONE_B,PROCESSING,NEW}
most_common_freqs      | {0.899572,0.09947,0.0009475,1.08333e-05}
histogram_bounds       |
correlation            | 0.845195
```

Compare Selectivity NEW

S_{before} = (1-sum(mvf))/(n_distinct-n_mcv)

= (1 - 0.9)/ (4 - 1)

= 0.033

S_{after} = 0.000,01

Test

- Sample Table with
  - 4 different status values
  - different Freq of `NEW`
- Run 100 times `analyze`
- Collect `pg_stats` values

Default Statistics Target

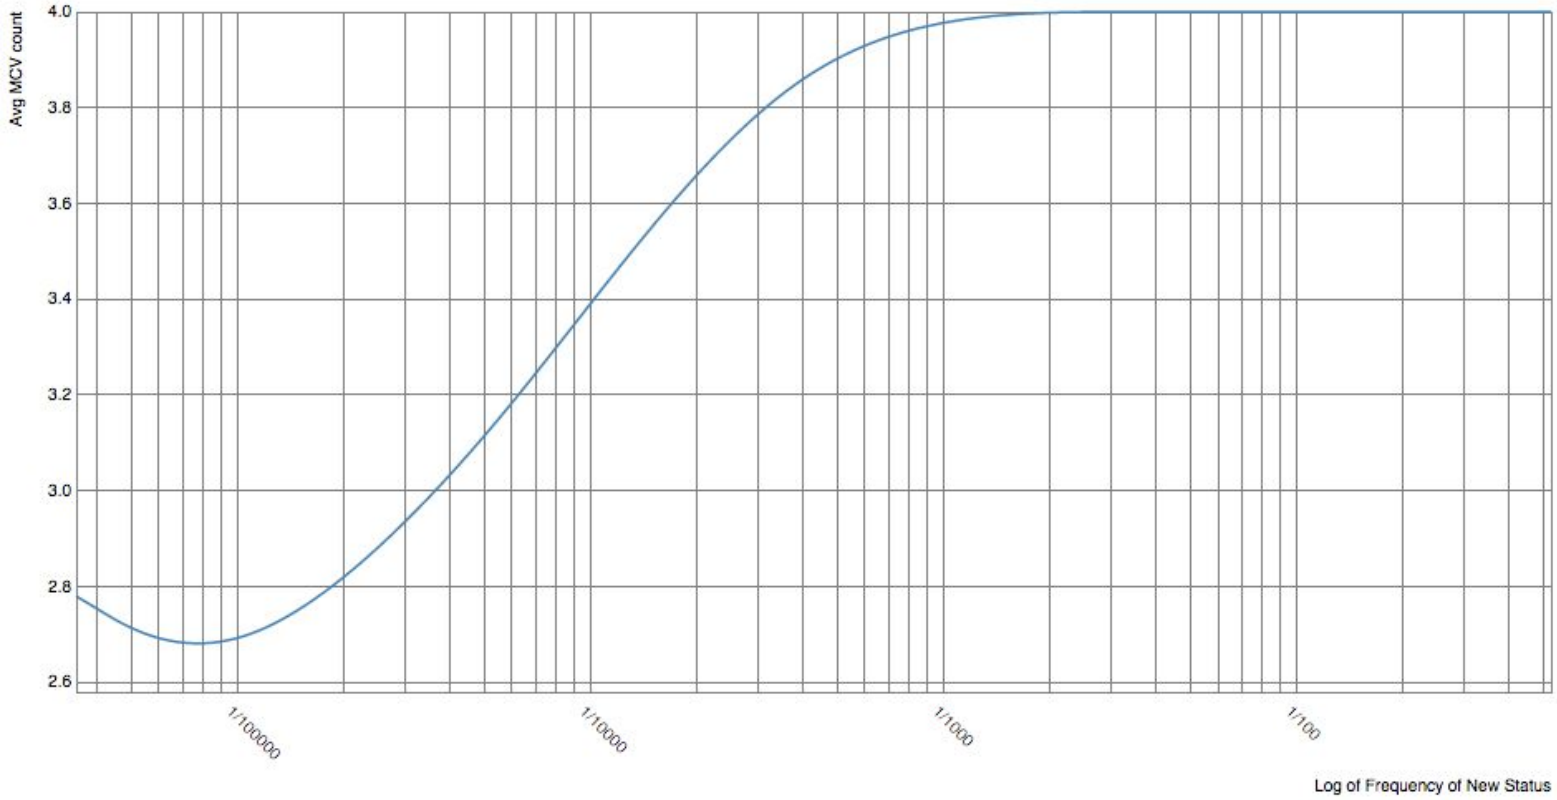

analyze.c

# of NEW in Sample	Most Common Values	Histogram Bounds
0	{DONE_A,DONE_B,PROCESSING}	{}
1	{DONE_A,DONE_B}	{PROCESSING,NEW}
>1	{DONE_A,DONE_B,PROCESSING,NEW}	{}

Statistics Target

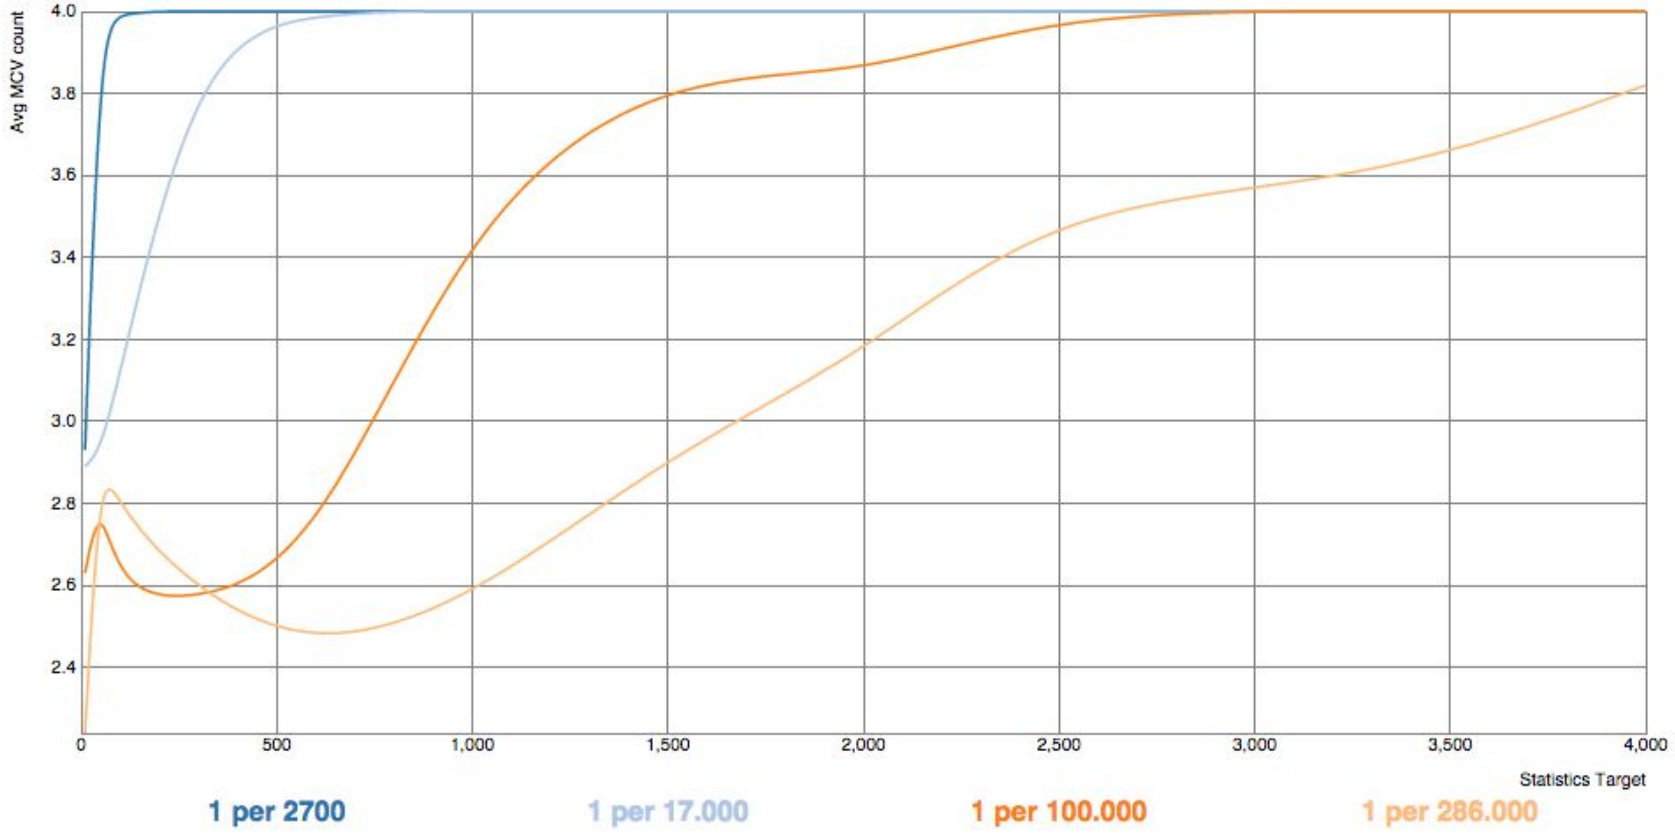

Sample Size

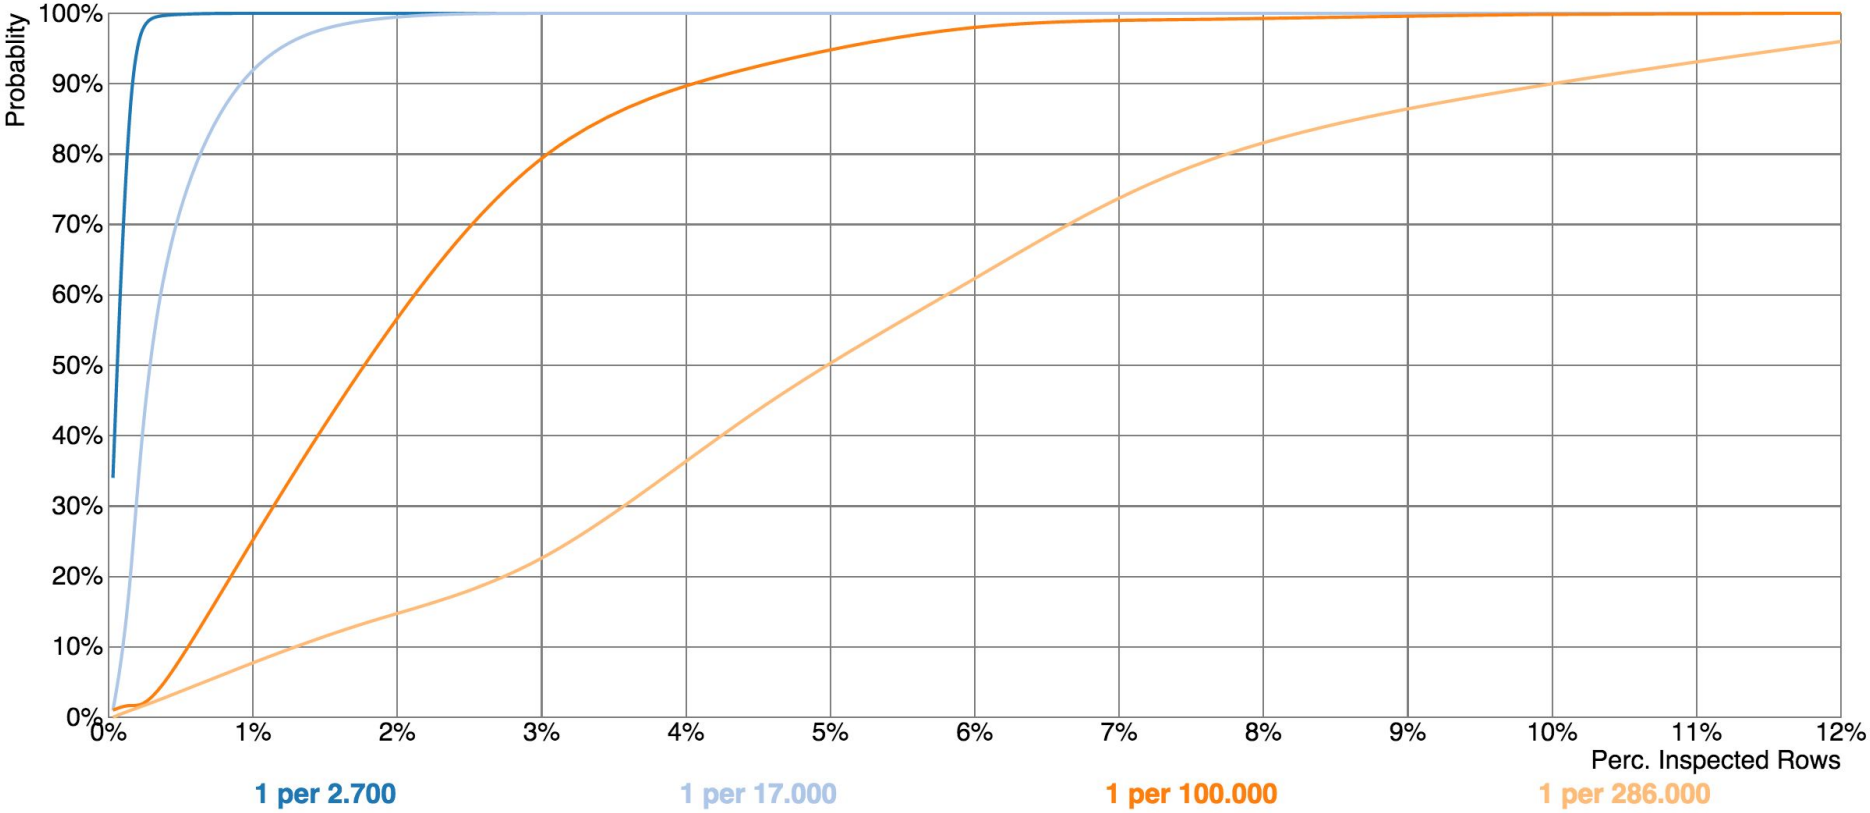
Probablity
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
0%
1%
2%
3%
4%
5%
6%
7%
8%
9%
10%
11%
12%
Perc. Inspected Rows
1 per 2.700
1 per 17.000
1 per 100.000
1 per 286.000

On Which Column

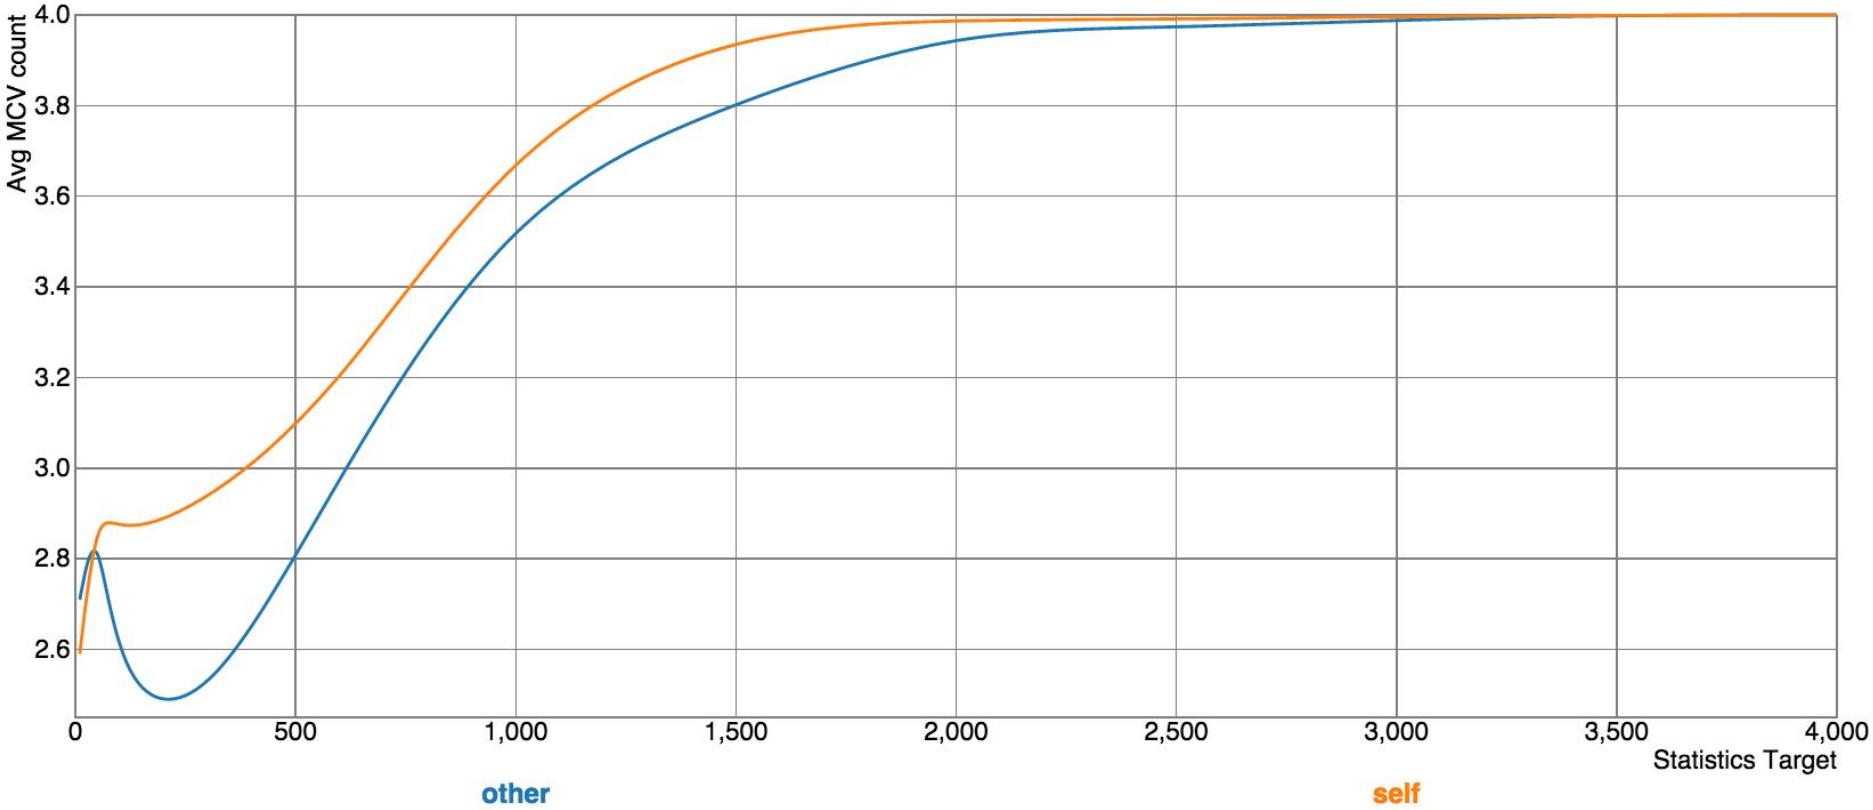

Rowcount

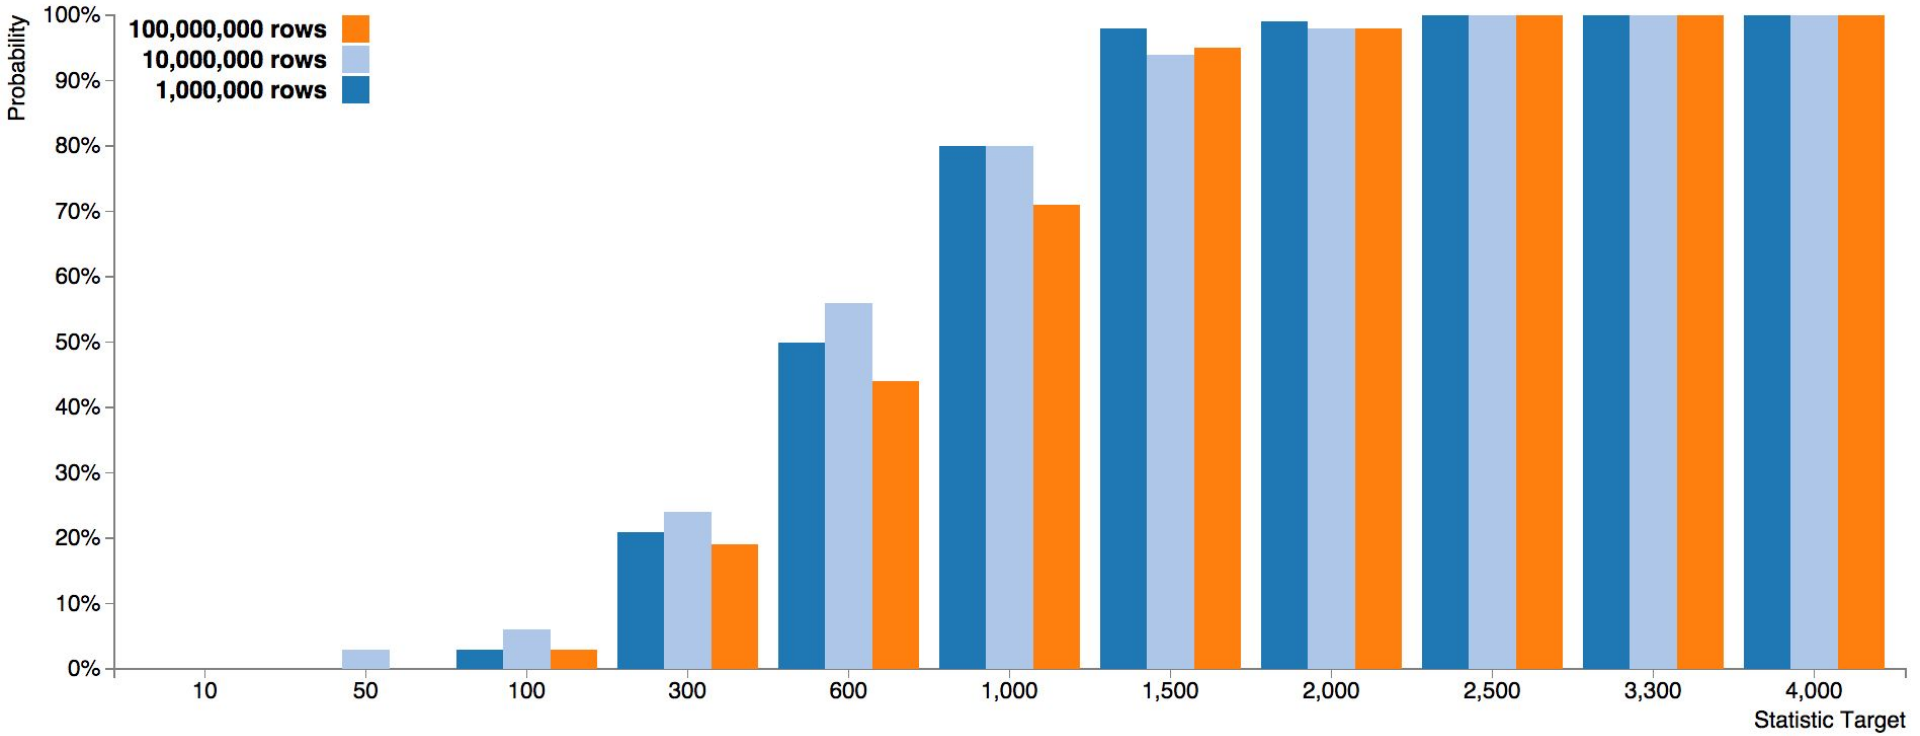
Probability
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
100,000,000 rows
10,000,000 rows
1,000,000 rows
10
50
100
300
600
1,000
1,500
2,000
2,500
3,300
4,000
Statistic Target

Summary

- Monitor
- Bump Statistics Target
- Archive
- Alternatives

Questions?