

# What's evolving in the Elasticsearch and Lucene



---

Berlin Buzzwords 2019  
Simon Willnauer  
June 18th, 2019



# Queries

Query	Before	After	Improvement
Fuzzy	46 qps		
Phrase	4 qps		
Bool AND	9.3 qps		
Bool OR	3.3 qps		
<b>Term</b>	<b>33 qps</b>		

Query	Before	After	Improvement
Fuzzy	46 qps	59 qps	28%
Phrase	4 qps	7 qps	87%
Bool AND	9.3 qps	23.5 qps	247%
Bool OR	3.3 qps	9.8 qps	292%
<b>Term</b>	<b>33 qps</b>	<b>1,160 qps</b>	<b>3,700%</b>

# MAGIC WAND

# Weak-AND

```
"query" : "elasticsearch and lucene"
```

```
max_score(and) == 1
```

```
max_score(elasticsearch) == 3
```

```
max_score(lucene) == 5
```

# Weak-AND

Min top-10 score	and (1)	elasticsearch (3)	lucene (5)
$\leq 1$	✓	✓	✓
$> 1$ and $\leq 4$	✗	✓	✓
$> 4$ and $\leq 9$	✗	✗	✓
$> 9$	✗	✗	✗

# ~~Weak-AND~~

```
"aggs": { ... }
```

```
"track_total_hits": true
```



# Weak-AND

```
"hits": {  
  "total": 123456789,  
  "hits": [ ... ]  
}
```



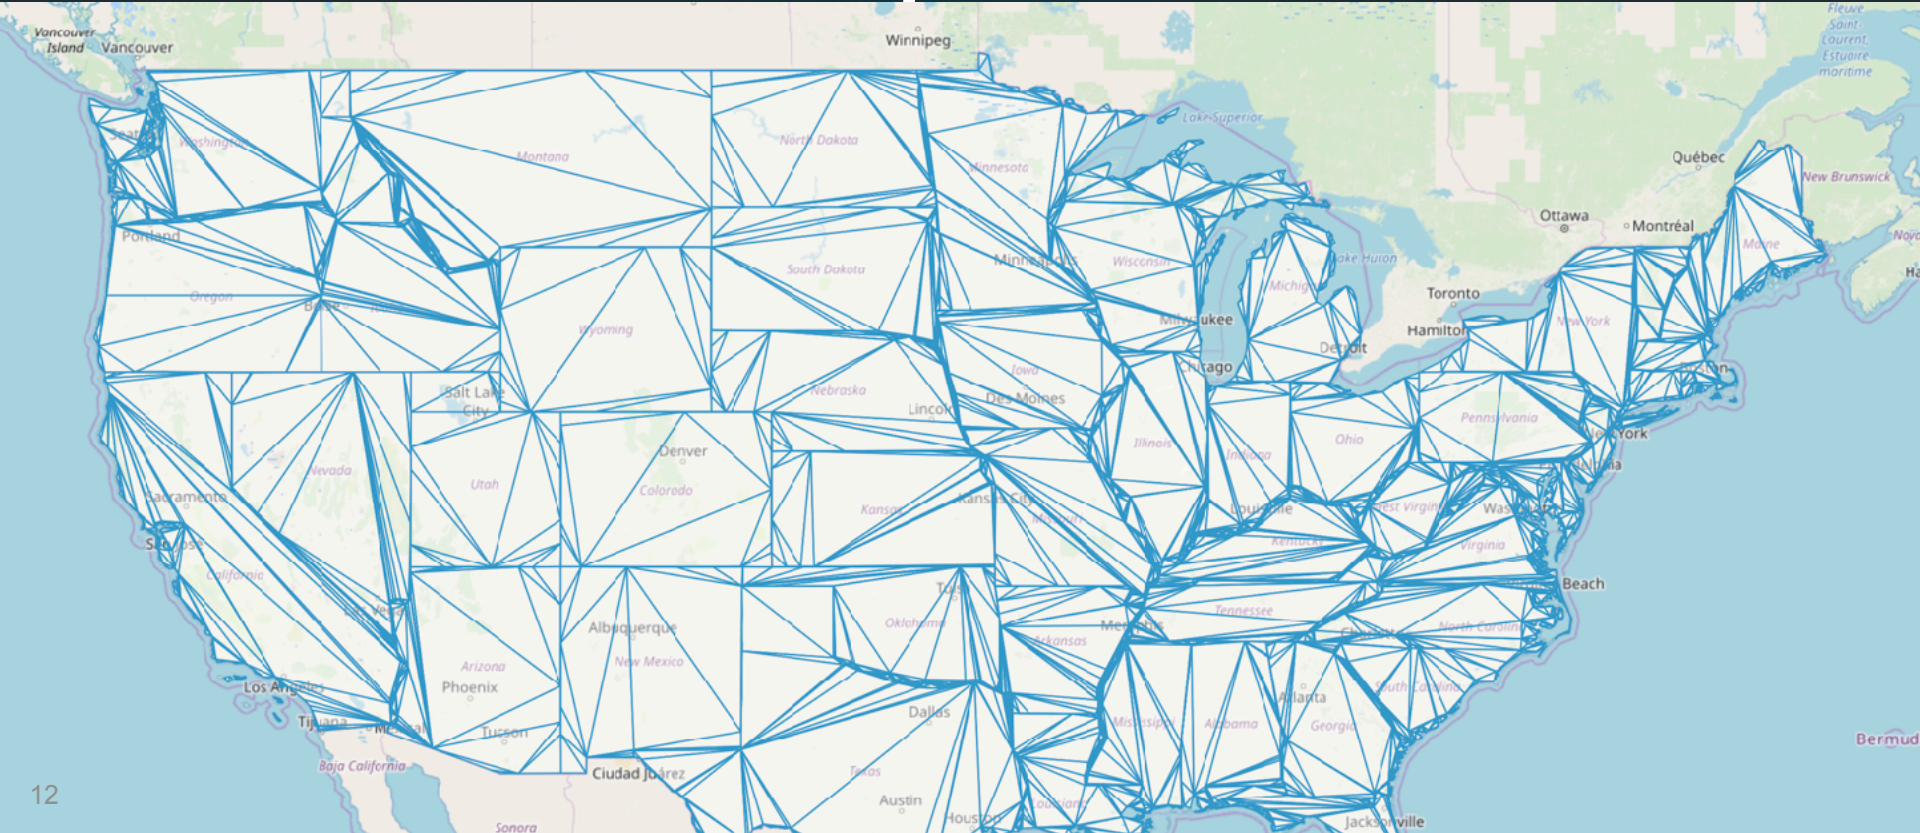
```
"hits": {  
  "total": {  
    "value": 10000,  
    "relation": "gte"  
  },  
  "hits": [ ... ]  
}
```

# Geoshapes

# BKD Trees

- v2.3: 1 dim, for numbers and dates
- v5.0: 2 dim, for geopoints
- v5.2: 2 dim, for number & date ranges
- v6.7: 7 dim, for geoshapes


# BKD Geoshapes



# BKD Geoshapes

- Accurate to 1mm, vs 50m
- Index is 60% smaller
- Indexing 60% faster
- Queries 50% faster
- Plus BKD GeoPoints 80% faster indexing

# Changes API





# Zen 2



minimum\_master\_nodes: 2



minimum\_master\_nodes: 2



minimum\_master\_nodes: 2




minimum\_master\_nodes: 1



~~minimum\_master\_nodes. 1~~

cluster.initial\_master\_nodes

# Adaptive Replica Selection





1P




1R



2P



2R







1P



1R



2P



2R



# Cross Cluster Search

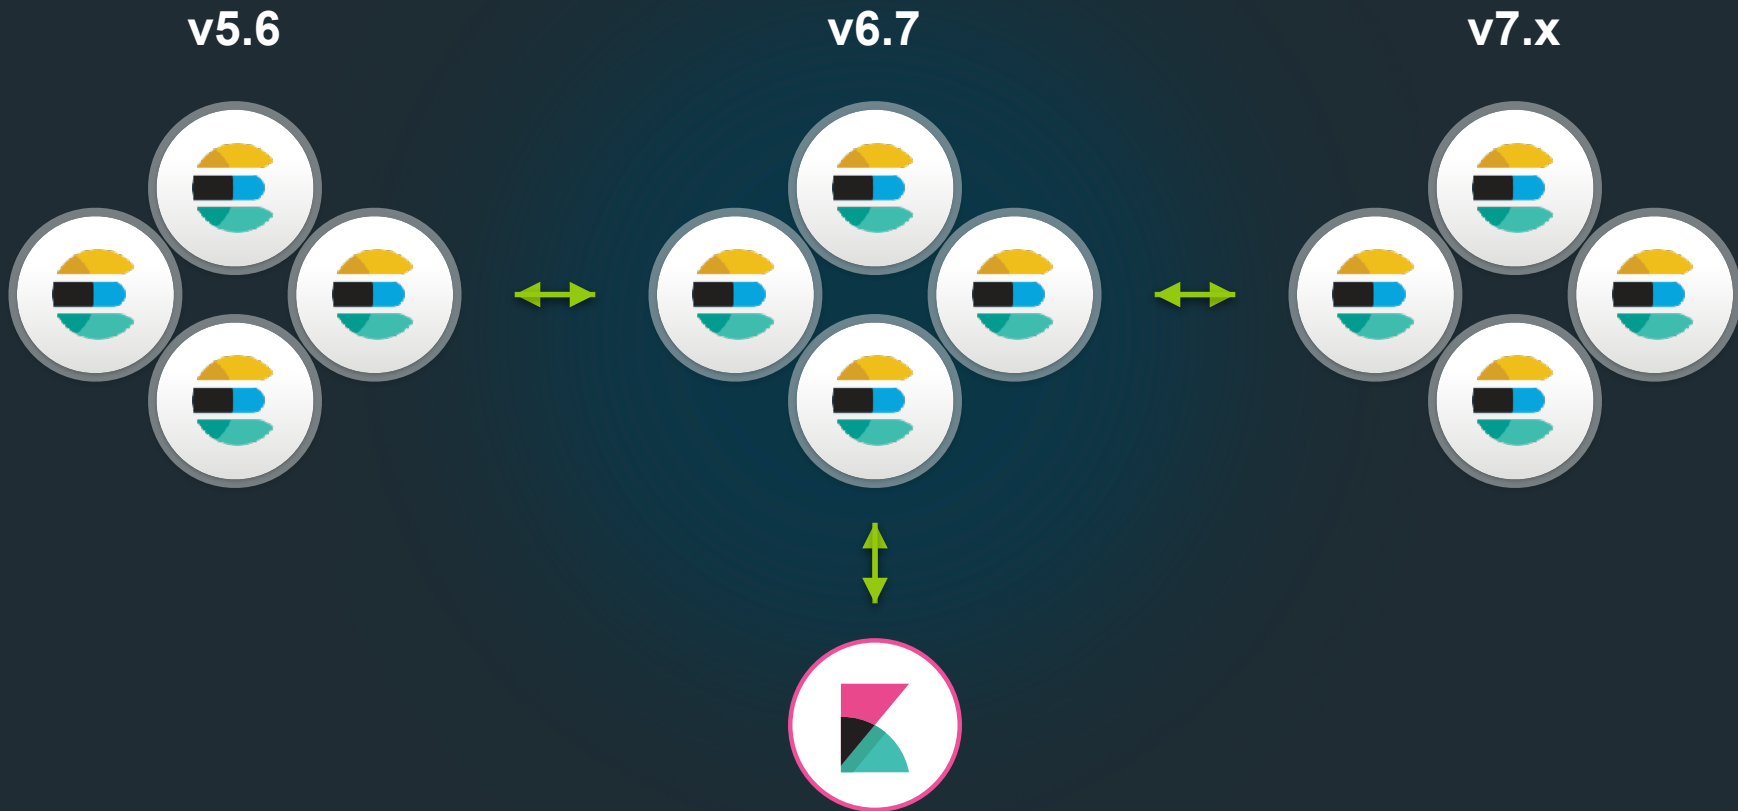
New York

London

Tokyo



# Three Major Versions



# Cross Cluster Replication

New York

London

Tokyo



ldn\_sales



ldn\_sales



New York

London

Tokyo



ny\_sales  
→



tk\_sales  
←



# New York

# London

# Tokyo



ny\_sales



ldn\_sales



tk\_sales



ldn\_sales





# Cross Cluster Replication

[Follower indices](#)

[Auto-follow patterns](#)

A follower index replicates a leader index on a remote cluster.

[+ Create a follower index](#)

<input type="checkbox"/> Name ↑	Status	Remote cluster	Leader index	Actions
<input type="checkbox"/> <a href="#">london_ny_sales</a>	● Active	NewYork	ny_sales	⋮
<input type="checkbox"/> <a href="#">london_tokyo_sales</a>	● Active	Tokyo	tokyo_sales	⋮

Rows per page: 20

# Frozen Indices

Heap

File system cache



Disk



Heap

File system cache



Disk






# Index Lifecycle Management

Hot Nodes

Warm Nodes

Cold Nodes






Hot Nodes



Warm Nodes



Cold Nodes






Hot Nodes



Warm Nodes



Cold Nodes







Hot Nodes



Warm Nodes



Cold Nodes






Hot Nodes



Warm Nodes



Cold Nodes






Hot Nodes



Warm Nodes



Cold Nodes





Hot Nodes



Warm Nodes



Cold Nodes





Hot Nodes



Warm Nodes



Cold Nodes



## Index lifecycle management



### Select or create a policy

An index lifecycle policy is a blueprint for transitioning your data over time. You can create a new policy or edit an existing policy and save it with a new name.

Existing policies

my\_policy5



Create new policy

### Edit policy my\_policy5

Configure the phases of your data and when to transition between them.

#### Hot phase

This phase is required. Your index is being queried and actively written to. You can optimize this phase for write throughput.

Enable rollover

If true, rollover the index when it gets too big or too old. The alias switches to the new index. [Learn more](#)

Maximum index size

3

gigabytes



Maximum age

days



#### Warm phase

# Nanosecond Timestamps

date:

2019/01/01 12:34:56.001



date\_nanos:  
2019/01/01  
12:34:56.000000001

# Full Text Search

# index\_prefixes:

qu, qui, quic, quick

br, bro, brow, brown

fo, fox, foxe, foxes

# index\_phrases:

the\_quick  
quick\_brown  
brown\_fox  
fox\_jumped  
jumped\_over  
over\_the  
the\_lazy  
lazy\_dog

```
match_phrase:  
  "quick brown fox"
```

```
match_phrase_prefix:  
  "quick brown f*"
```

# Intervals Query

the quick brown fox  
jumped over the  
lazy dog

quick...fox

the quick brown fox  
jumped over the  
lazy dog



quick...fox, “lazy dog”

the quick brown fox  
jumped over the  
lazy dog

(quick...fox, "lazy dog")~5

the quick brown fox  
jumped over the  
lazy dog

Thank you, questions?