



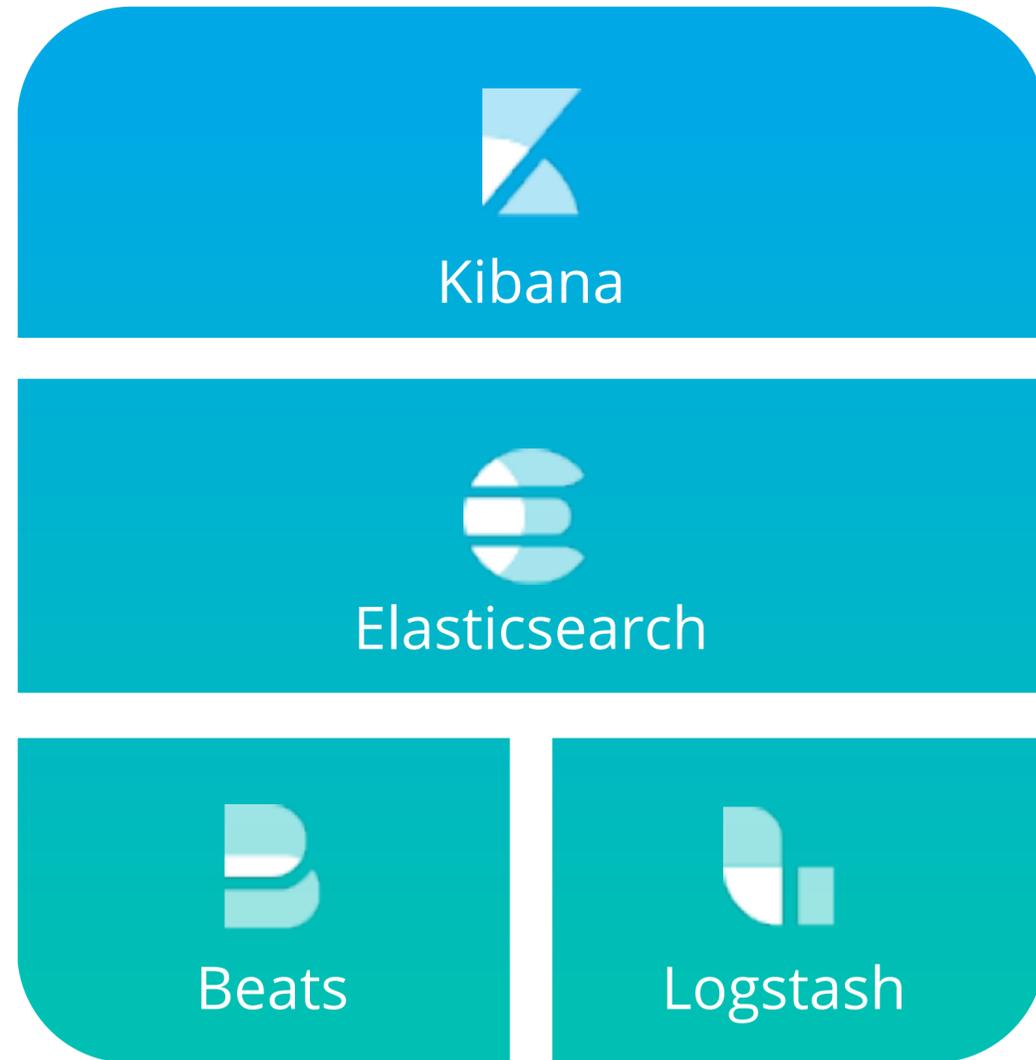
elastic  
中文社区

IT大咖说  
知识分享平台

# Elastic Stack 5.x and beyond ...

Medcl

# Elastic Stack & X-Pack



**Security**



**Alerting**



**Monitoring**



**Reporting**



**Graph**



**Machine Learning**



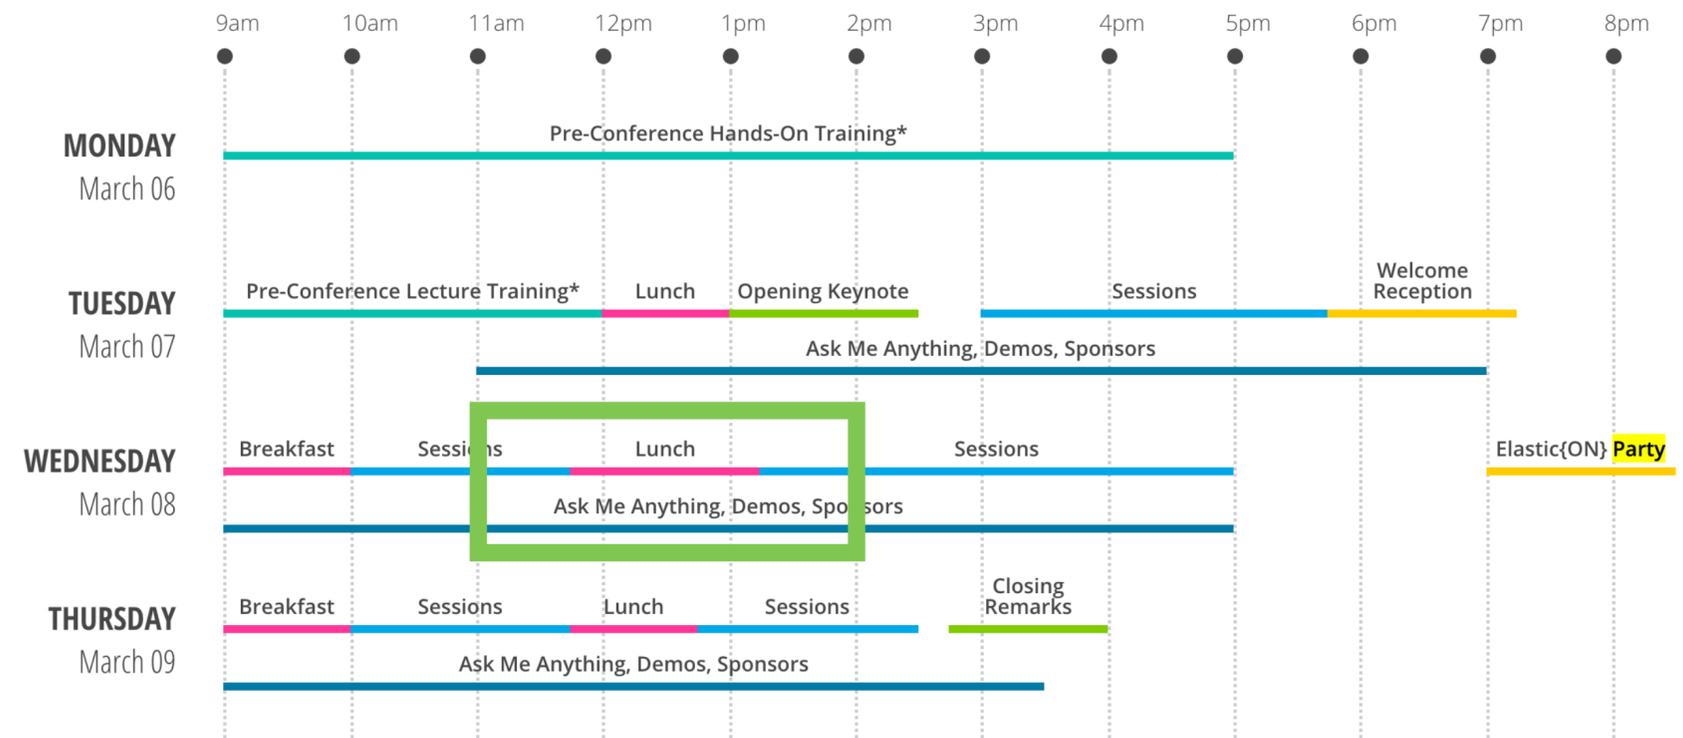
# Elasticsearch since 5.0

# Numeric & Date Range Fields (5.2)

## Mapping Improvements

- New types for date/number ranges (5.2)  
(*date\_range, int\_range, float\_range*)

## What's happening Wednesday 11am - 2pm



# Keyword Normalizer (5.2)

## Mapping Improvements

```
{  
  "city": {  
    "type": "text"  
    "fields": {  
      "city.keyword": {  
        "type": "keyword"  
      }  
    }  
  }  
}
```

← No Analysis

San Francisco  
SAN FRANCISCO  
san francisco  
San francisc0

Normalizer → san francisco

# Terms Aggregation Partitioning (5.2)

## Returning ALL the Terms, in Manageable Chunks

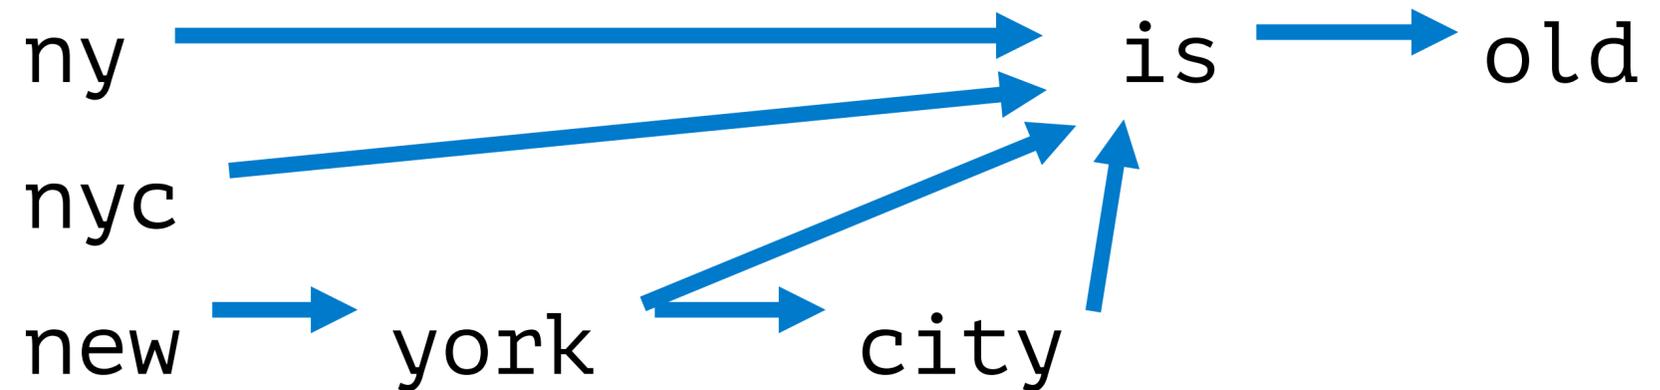
- frequent request
- return all responses from a terms aggs
- Terms can now be broken into partitions and partitions are returned by number

```
{
  "size": 0,
  "aggs": {
    "expired_sessions": {
      "terms": {
        "field": "account_id",
        "include": {
          "partition": 0,
          "num_partitions": 20
        },
        "size": 10000,
        "order": {
          "last_access": "asc"
        }
      },
      "aggs": {
        "last_access": {
          "max": {
            "field": "access_date"
          }
        }
      }
    }
  }
}
```

# Synonym Graph Token Filter (5.2)

## Search & Aggregation Improvements

- Improved querying for multi-word synonyms `SynonymGraphFilter`



# Cluster Allocation Explain API (5.2)

## Operational Optimizations - Understand Shard Allocation

`/_cluster/allocation/explain`

- Diagnose unassigned shards
- clear human readable descriptions when things fail

# Cross Cluster Search (5.3)

Tribe node is dead. Long live Cross-cluster search.

- Minimal viable solution to supersede tribe
- Addresses many of the challenges with tribe node
- Reduces the problem domain to query execution
- Cluster related information is reduced to a namespace

# Field Collapsing (5.3)

One method to rule them all...

- Simple (almost) no setup!
- Great for query-time group/category de-dup

```
GET /twitter/tweet/_search
{
  "query": {
    "match": {
      "message": "elasticsearch"
    }
  },
  "collapse" : {
    "field" : "user", ①
    "inner_hits": {
      "name": "last_tweets", ②
      "size": 5, ③
      "sort": [{ "date": "asc" }] ④
    },
    "max_concurrent_group_searches": 4 ⑤
  },
  "sort": ["likes"]
}
```

# Elasticsearch Keystore

If you like it, you should put it in a keystore.

- Sensitive settings should not be protected by filesystem permissions only.
- Commands feel familiar:
  - bin/elasticsearch-keystore create
  - bin/elasticsearch-keystore list
  - bin/elasticsearch-keystore add the.setting.name.to.set
  - bin/elasticsearch-keystore remove the.setting.name.to.remove
- Just the framework/start: sensitive settings to be pulled in

# And many more ...

- Batched Reduction of Search Results
- Smarter query caching
- Faster geo, range, and nested queries
- Unified highlighter
- Cancellable searches
- More Painless improvements
- Index partitioning/routing



# Kibana since 5.0



# New visualizations

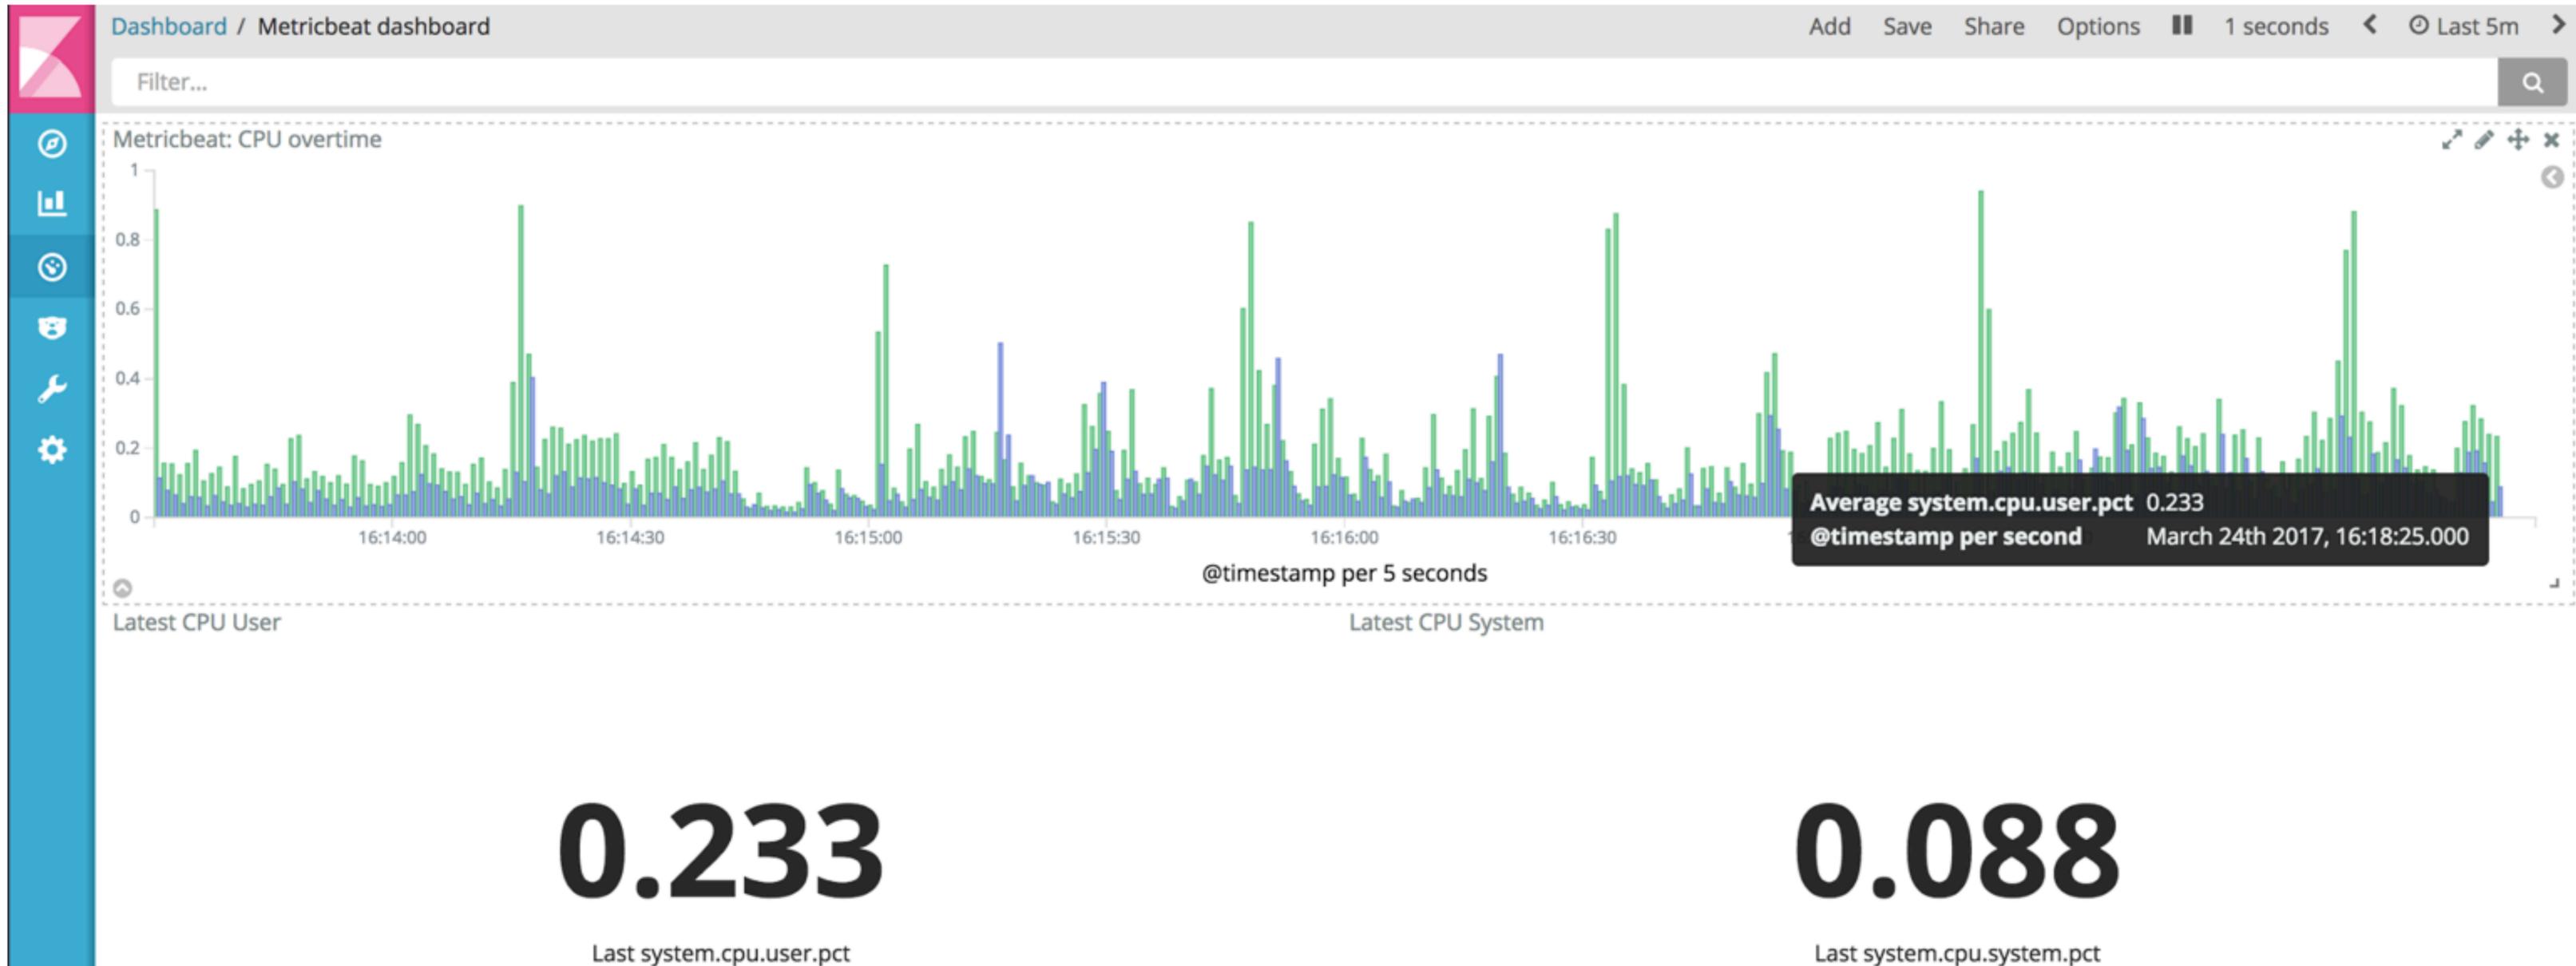
## Heatmap (5.2)

Heatmap - Downloads by Size

ZW	2,303	769	0	357	1,774
ZM	3,618		0	138	1,679
ZA	3,569	2,945	1,973	275	1,771
YE	0	6,457	0	868	2,479
XK			3,643		
VN	2,783	494	0	95	1,511
VE	4,908	6,820	0	56	0
UZ	0	1,864	1,779	278	1,710
UY	2,647		1,612		2,738
US	0	90	0	0	0
	png	php	jpg	gif	css

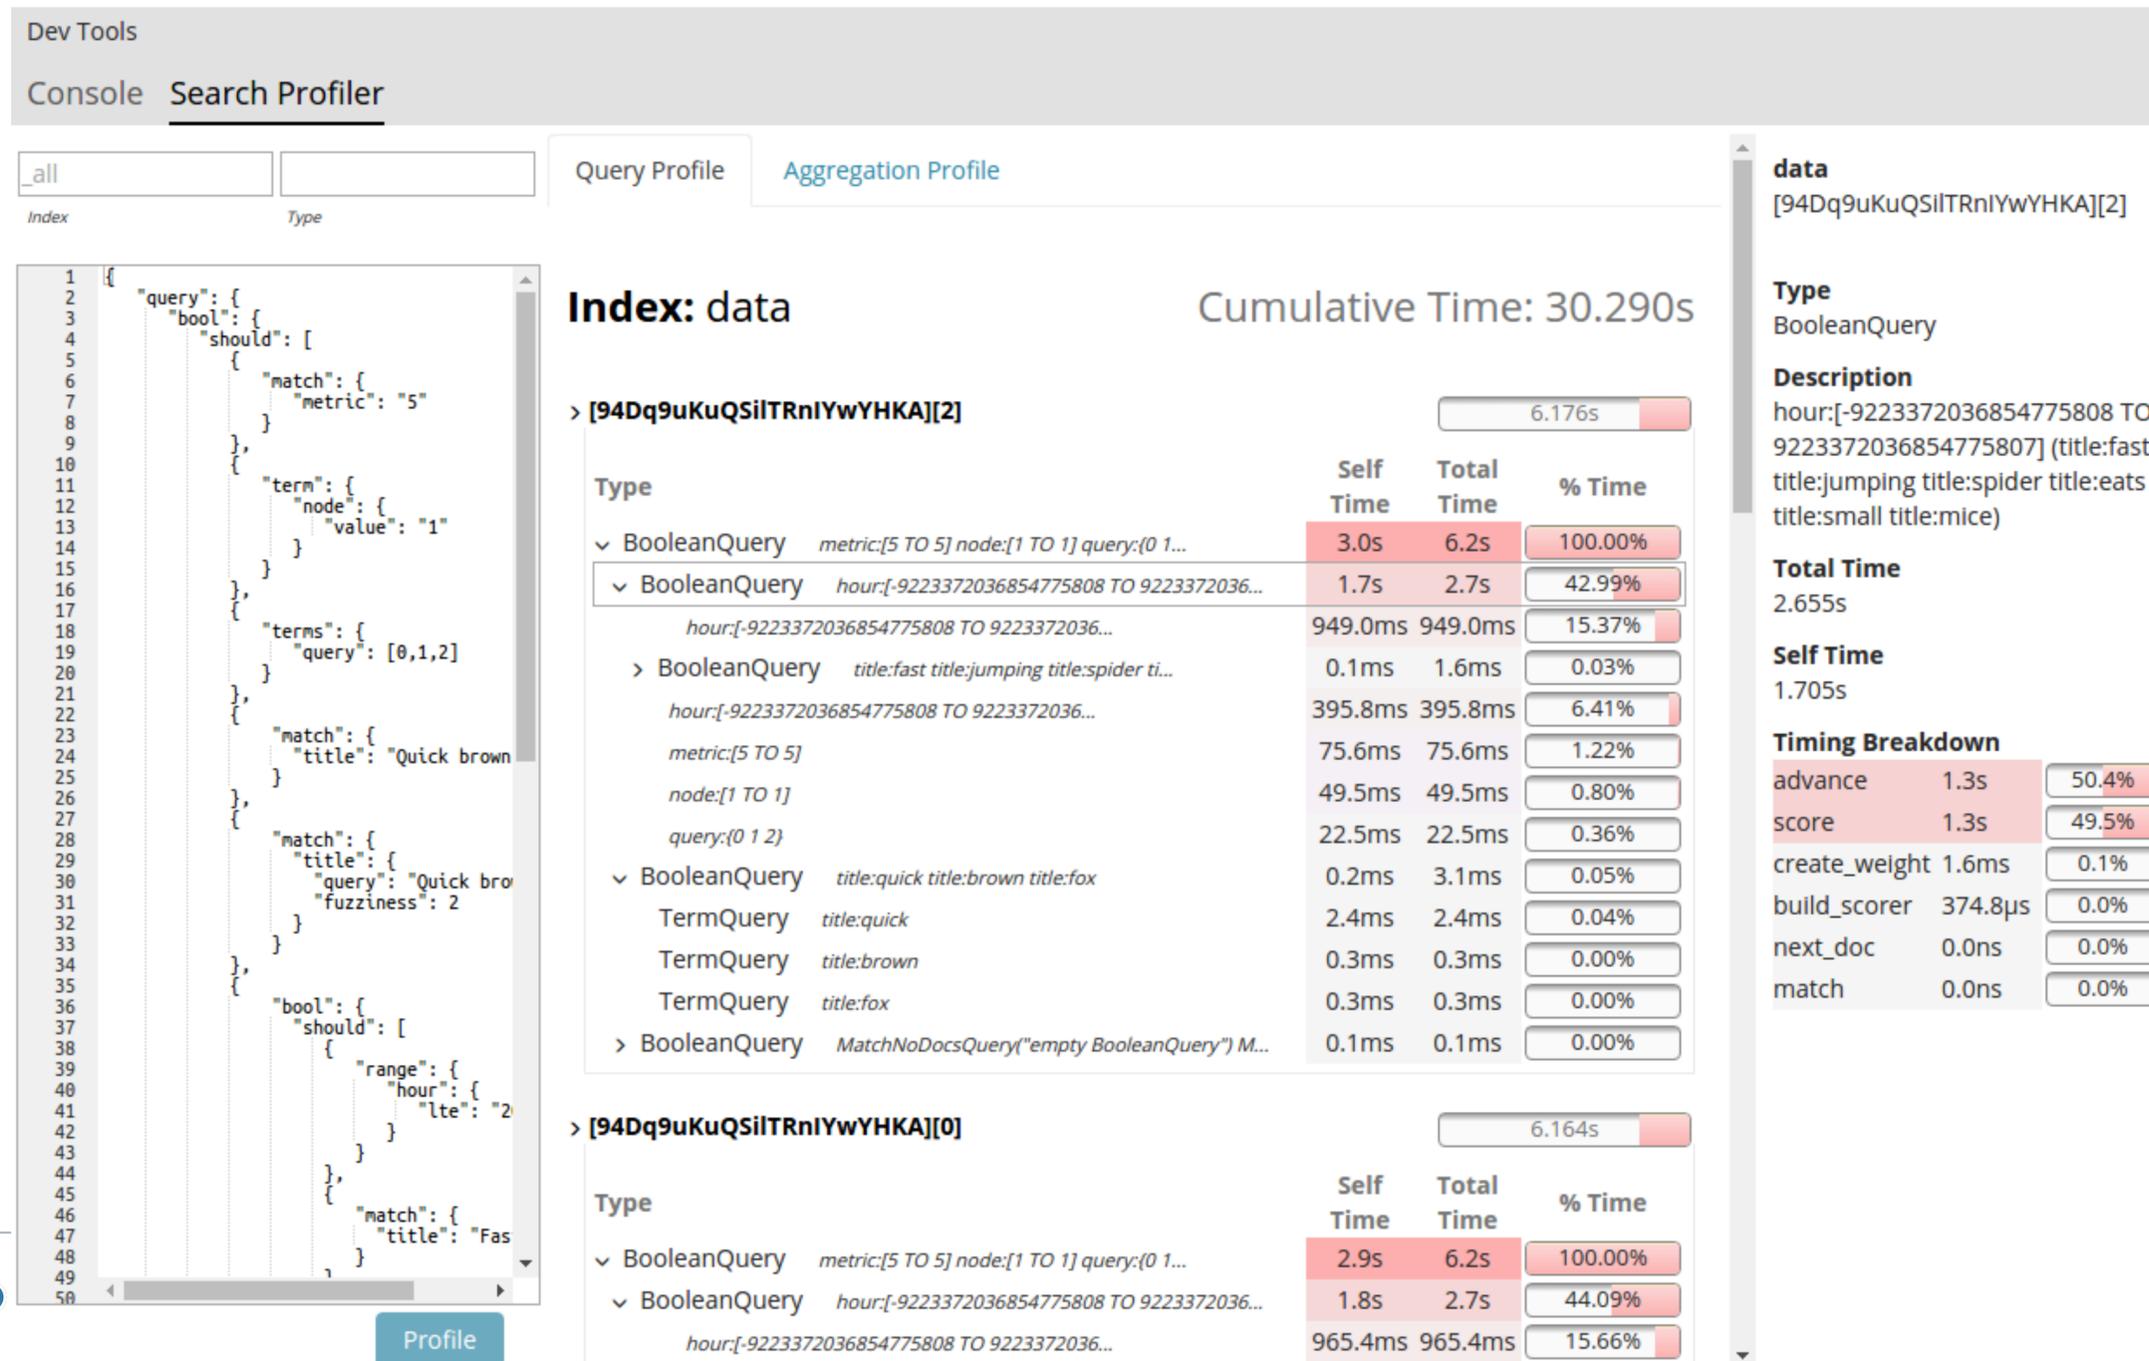
# New Analytics

## Top Hits Aggregation (5.3) - Visualize the 'latest' metric



# Profile your Search Queries

## Search Profiler (5.1) - Detect and visualize bottlenecks in your query



Dev Tools  
Console Search Profiler

Index: data  
Type: BooleanQuery

Cumulative Time: 30.290s

```
1 {
2   "query": {
3     "bool": {
4       "should": [
5         {
6           "match": {
7             "metric": "5"
8           }
9         },
10        {
11          "term": {
12            "node": {
13              "value": "1"
14            }
15          }
16        },
17        {
18          "terms": {
19            "query": [0,1,2]
20          }
21        },
22        {
23          "match": {
24            "title": "Quick brown
25          }
26        },
27        {
28          "match": {
29            "title": {
30              "query": "Quick bro
31              "fuzziness": 2
32            }
33          }
34        },
35        {
36          "bool": {
37            "should": [
38              {
39                "range": {
40                  "hour": {
41                    "lte": "2
42                  }
43                }
44              },
45              {
46                "match": {
47                  "title": "Fas
48                }
49              }
50            ]
51          }
52        }
53      ]
54    }
55  }
56 }
```

Type	Self Time	Total Time	% Time
> [94Dq9uKuQSiITRniYwYHKA][2]	6.176s		
BooleanQuery metric:[5 TO 5] node:[1 TO 1] query:{0 1...	3.0s	6.2s	100.00%
BooleanQuery hour:[-9223372036854775808 TO 9223372036...	1.7s	2.7s	42.99%
hour:[-9223372036854775808 TO 9223372036...	949.0ms	949.0ms	15.37%
BooleanQuery title:fast title:jumping title:spider ti...	0.1ms	1.6ms	0.03%
hour:[-9223372036854775808 TO 9223372036...	395.8ms	395.8ms	6.41%
metric:[5 TO 5]	75.6ms	75.6ms	1.22%
node:[1 TO 1]	49.5ms	49.5ms	0.80%
query:{0 1 2}	22.5ms	22.5ms	0.36%
BooleanQuery title:quick title:brown title:fox	0.2ms	3.1ms	0.05%
TermQuery title:quick	2.4ms	2.4ms	0.04%
TermQuery title:brown	0.3ms	0.3ms	0.00%
TermQuery title:fox	0.3ms	0.3ms	0.00%
BooleanQuery MatchNoDocsQuery("empty BooleanQuery") M...	0.1ms	0.1ms	0.00%

Type	Self Time	Total Time	% Time
> [94Dq9uKuQSiITRniYwYHKA][0]	6.164s		
BooleanQuery metric:[5 TO 5] node:[1 TO 1] query:{0 1...	2.9s	6.2s	100.00%
BooleanQuery hour:[-9223372036854775808 TO 9223372036...	1.8s	2.7s	44.09%
hour:[-9223372036854775808 TO 9223372036...	965.4ms	965.4ms	15.66%

**data**  
[94Dq9uKuQSiITRniYwYHKA][2]

**Type**  
BooleanQuery

**Description**  
hour:[-9223372036854775808 TO 9223372036854775807] (title:fast title:jumping title:spider title:eats title:small title:mice)

**Total Time**  
2.655s

**Self Time**  
1.705s

**Timing Breakdown**

Component	Time	%
advance	1.3s	50.4%
score	1.3s	49.5%
create_weight	1.6ms	0.1%
build_scorer	374.8µs	0.0%
next_doc	0.0ns	0.0%
match	0.0ns	0.0%

Profile

\* requires X-Pack (Basic)



# Internationalization Support

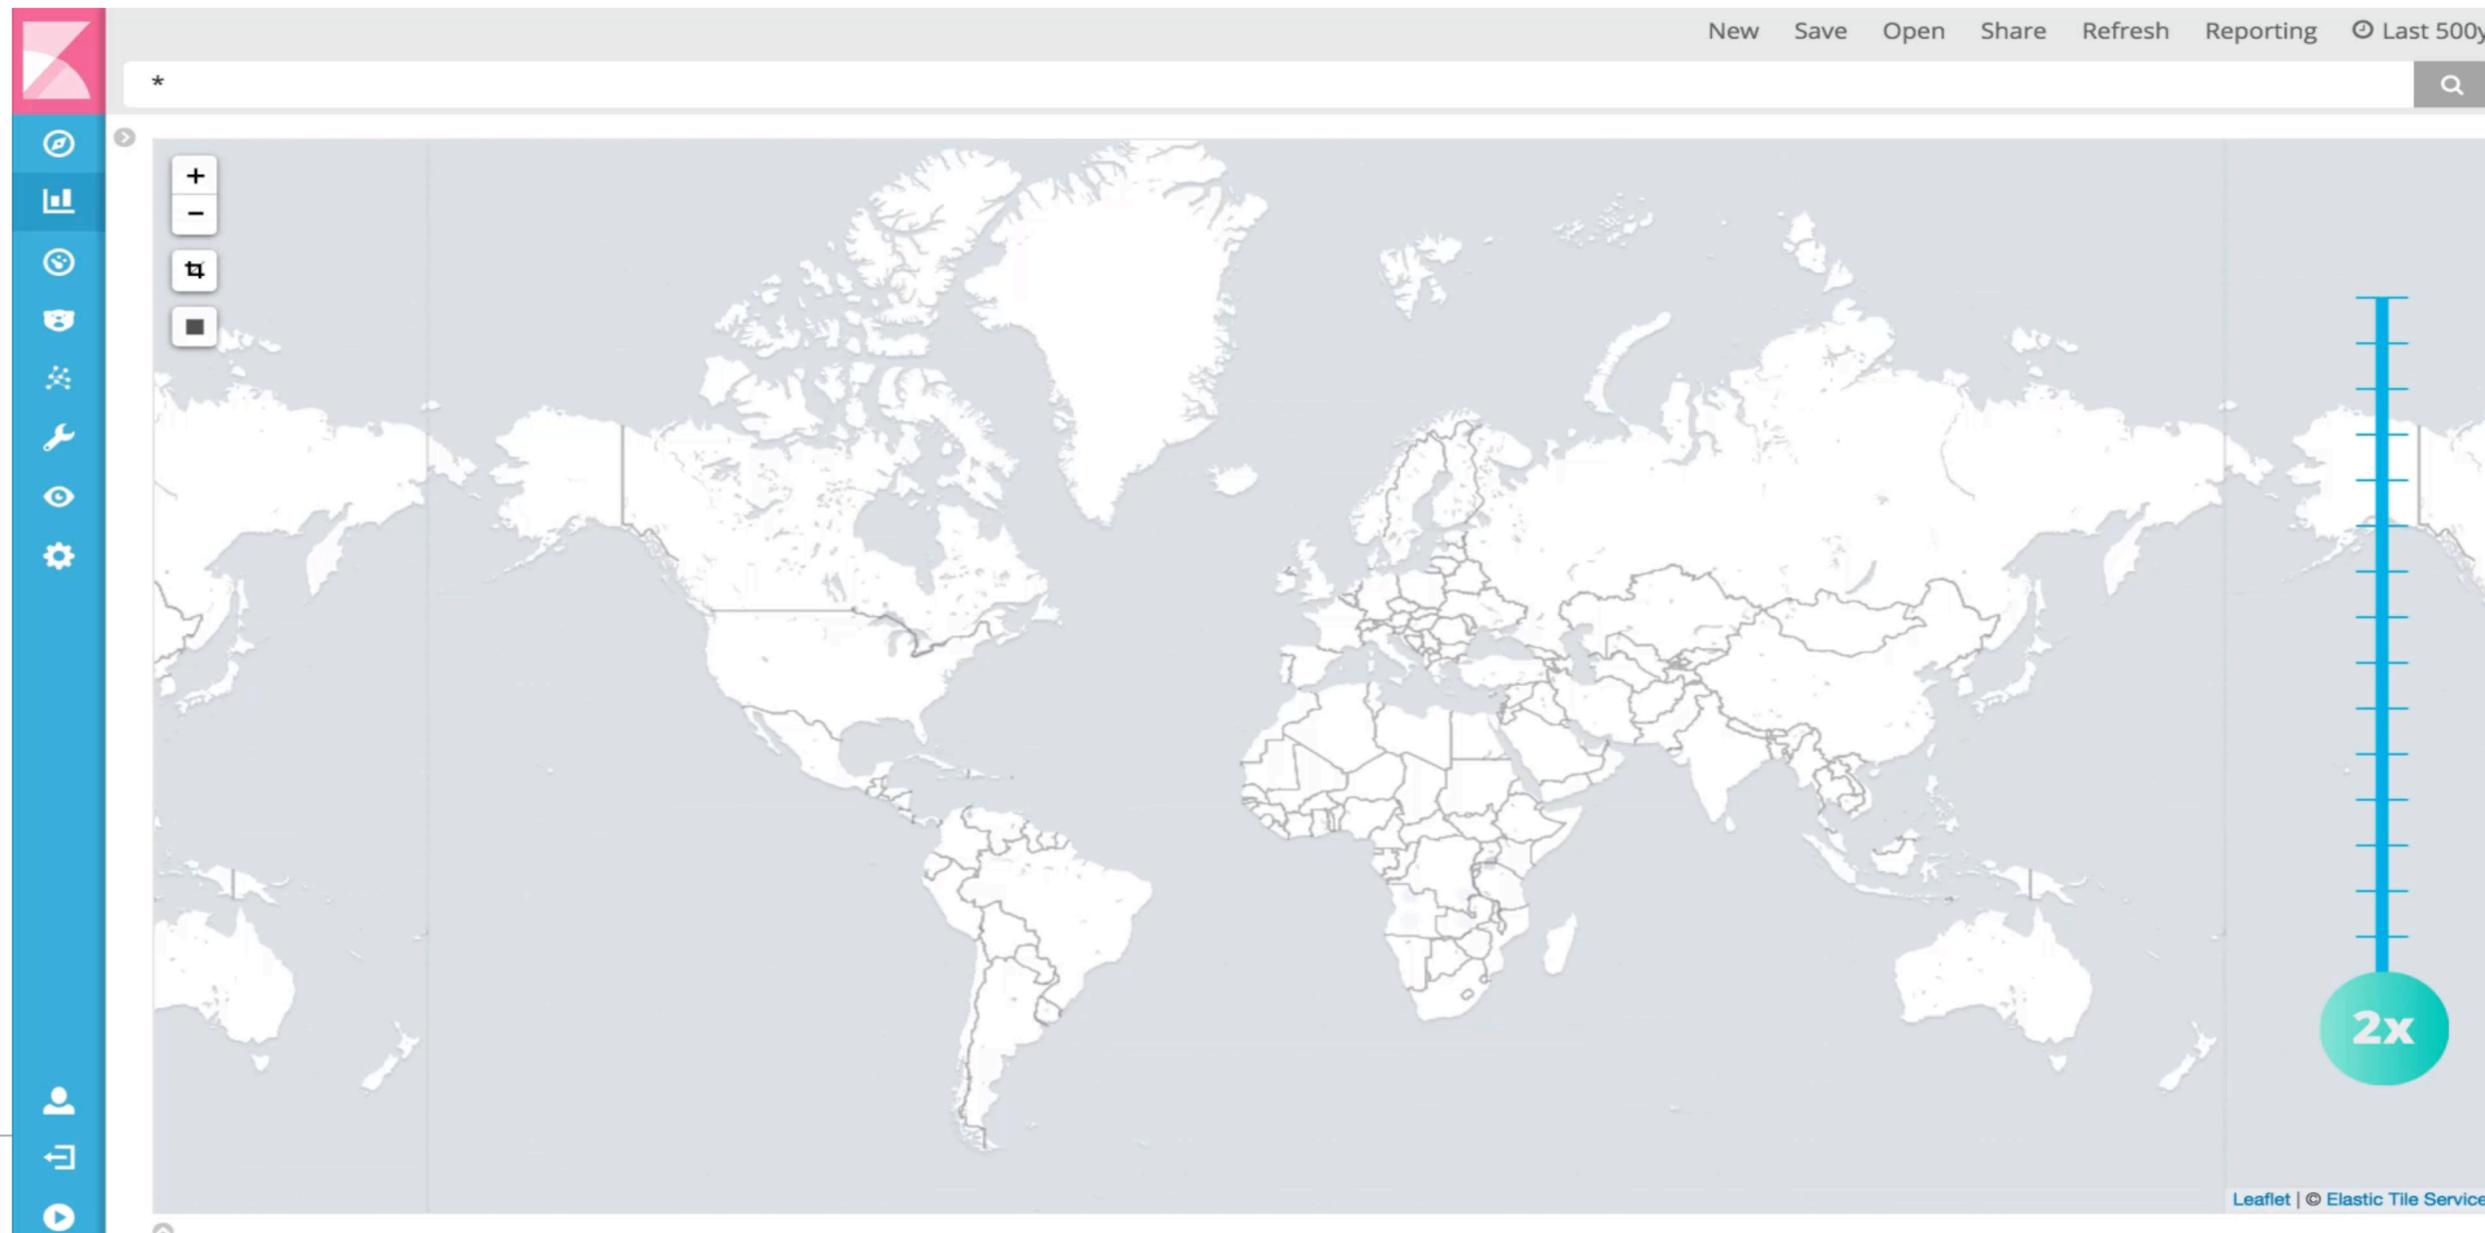
## I18N, phase 1 complete (5.2)

- Adheres to browser preference for language
- Translations as plugins
- Thanks IBM!



# Uses Elastic Tile Service

Offers up to 18 level zoom



X-Pack (Basic) required  
for zoom > 10

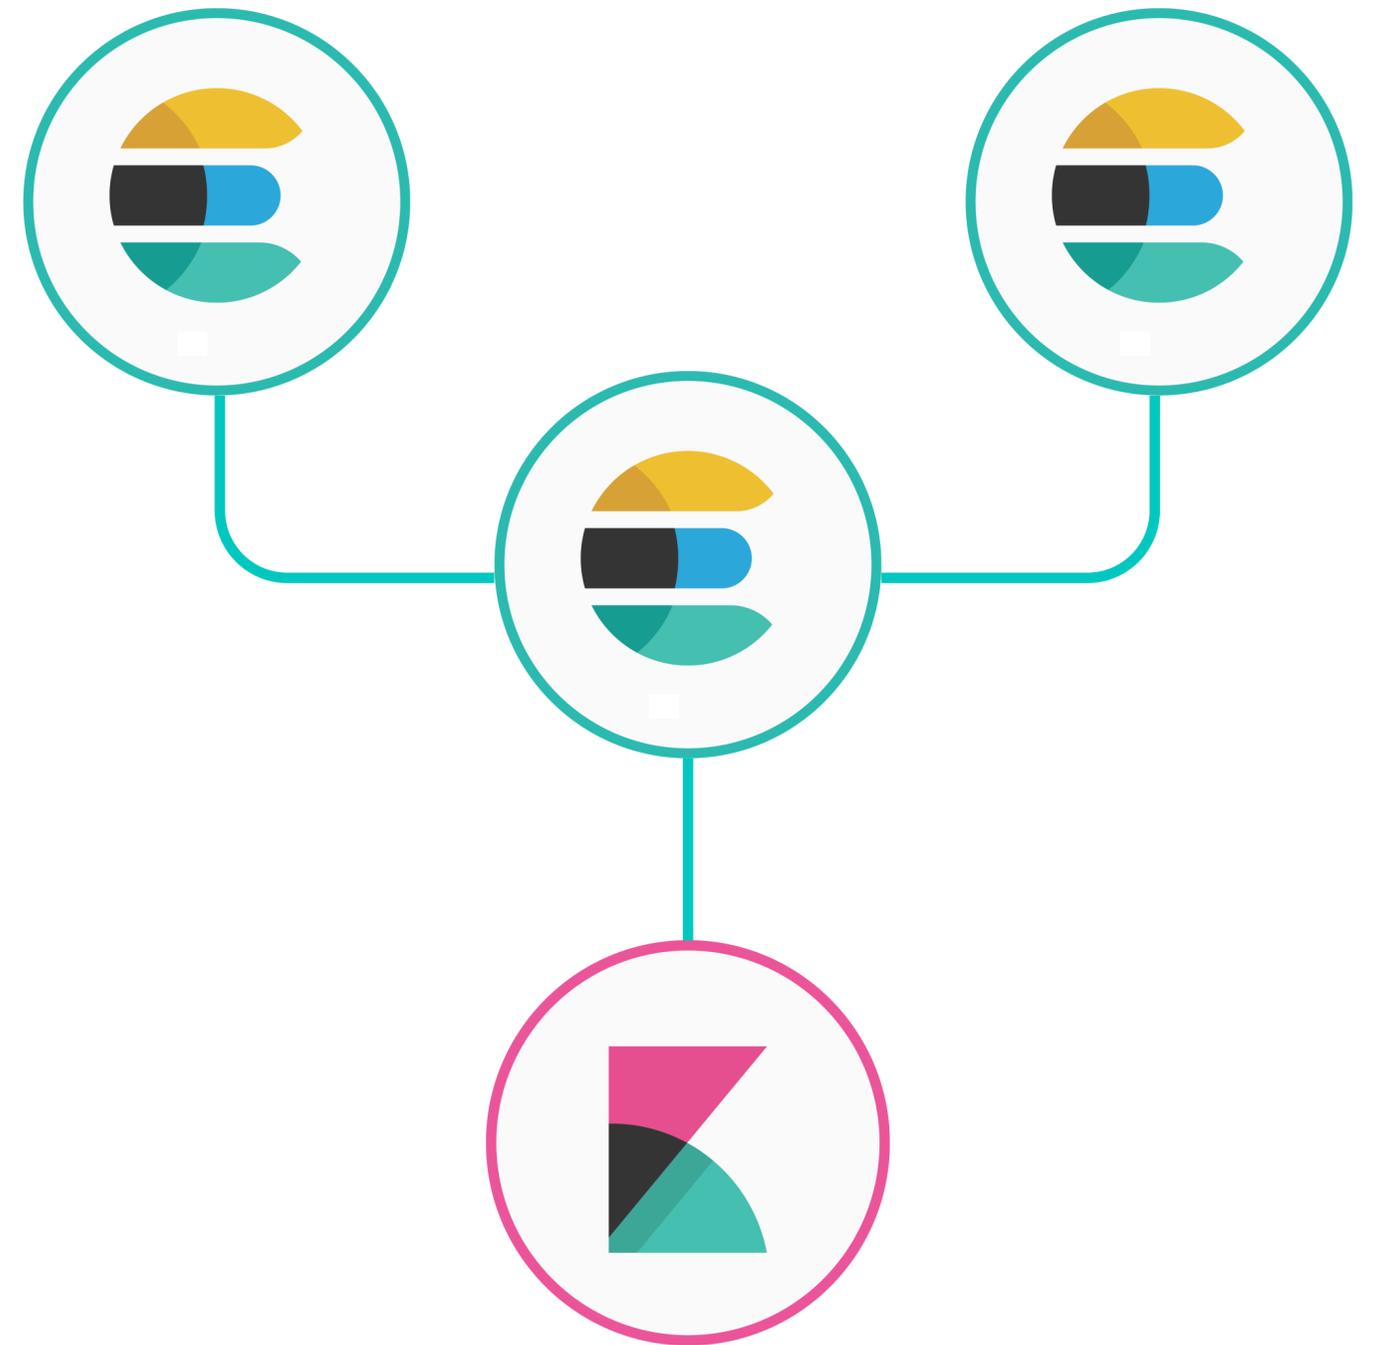
# Added support for Tribe

## Tribe Support

As simple as:

`elasticsearch.tribe.url`

No jumping through hoops.



# View/Edit mode for Dashboards



elastic  
中文社区



**kibana** Dashboard / Editing Stock Market Save Cancel Add Options July 20th 2009, 06:14:32.633 to September 13th 2010, 11:16:56.633

Save dashboard

Stock Market

Save as a new dashboard

Store time with dashboard

Save

Filter...

**Average open, high, low close**

● Average open  
● Average high  
● Average low  
● Average close

**Average volume vs. average close**

● Average volume  
● Average close

**Stock: Amazon**

● Average open  
● Average high  
● Average low  
● Average close

**Volume over time**

● Average volume

**Stock: Google**

● Average open  
● Average high  
● Average low  
● Average close

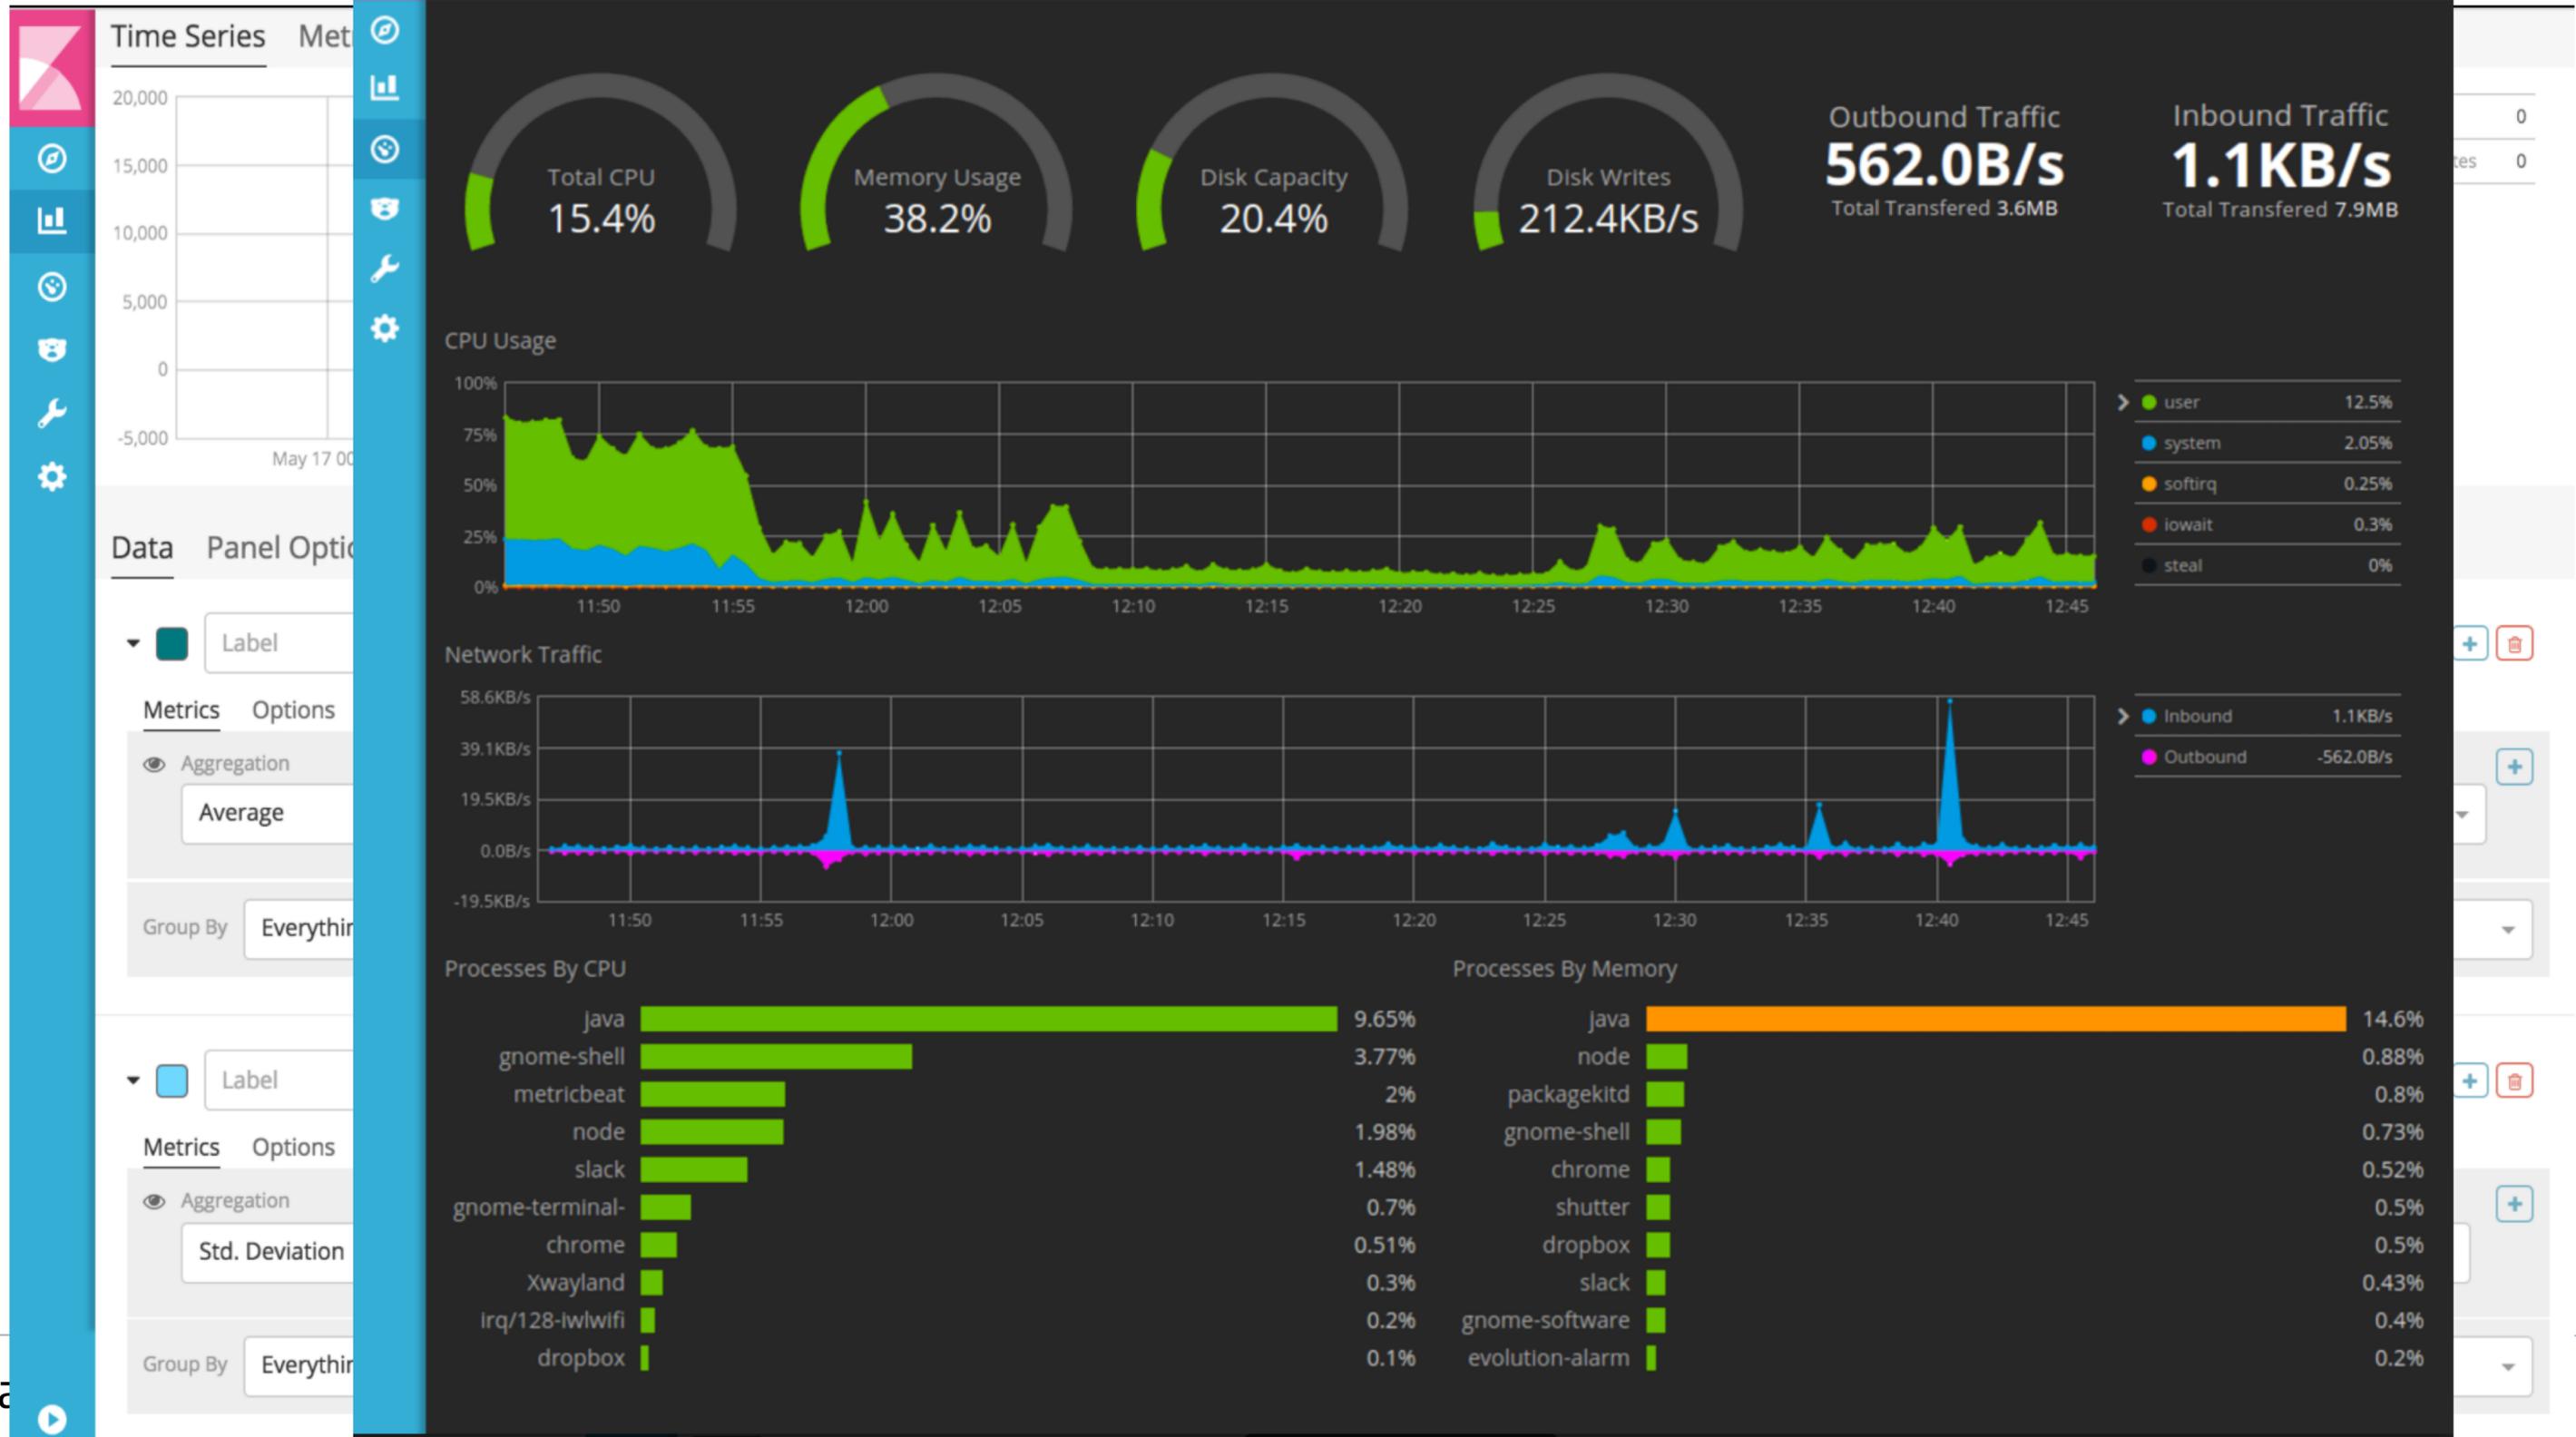
**Stock: Apple**

● Average open  
● Average high  
● Average low  
● Average close

Collapse



# Time Series



# Event Context

Surrounding Documents in logstash-\*

log#AVtdkRwgxIjtxpRZ3Q9K

Load 5 more  newer documents

Time	@message
▶ May 20th 2015, 13:41:28.906	189.244.165.222 - - [2015-05-20T20:41:28.906Z] "GET /uploads/ronald-mcnair.jpg HTTP/1.1" 200 3407 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:6.0a1) Gecko/20110421 Firefox/6.0a1"
▶ May 20th 2015, 13:40:05.434	8.105.89.9 - - [2015-05-20T20:40:05.434Z] "GET /uploads/michael-p-anderson.jpg HTTP/1.1" 200 9179 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:6.0a1) Gecko/20110421 Firefox/6.0a1"
▶ May 20th 2015, 13:36:52.315	49.151.142.71 - - [2015-05-20T20:36:52.315Z] "GET /people/type:astronauts/name:krasimir-stoyanov/profile HTTP/1.1" 200 4663 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:6.0a1) Gecko/20110421 Firefox/6.0a1"
▶ May 20th 2015, 13:36:16.940	167.67.215.75 - - [2015-05-20T20:36:16.940Z] "GET /uploads/john-blaha.jpg HTTP/1.1" 200 8458 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:6.0a1) Gecko/20110421 Firefox/6.0a1"
▶ May 20th 2015, 13:34:26.604	124.254.160.199 - - [2015-05-20T20:34:26.604Z] "GET /uploads/voskhod-1.jpg HTTP/1.1" 200 9979 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:6.0a1) Gecko/20110421 Firefox/6.0a1"
▶ May 20th 2015, 13:32:14.679	161.107.162.58 - - [2015-05-20T20:32:14.679Z] "GET /styles/ad-blocker.css HTTP/1.1" 200 4825 "-" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1; .NET CLR 1.1.4322)"
▶ May 20th 2015, 13:32:06.085	128.155.230.205 - - [2015-05-20T20:32:06.085Z] "GET /uploads/philip-k-chapman.jpg HTTP/1.1" 200 3088 "-" "Mozilla/5.0 (X11; Linux i686) AppleWebKit/534.24 (KHTML, like Gecko) Chrome/11.0.696.50 Safari/534.24"
▶ May 20th 2015, 13:31:24.323	167.137.188.138 - - [2015-05-20T20:31:24.323Z] "GET /uploads/charles-e-brady-jr-.jpg HTTP/1.1" 200 3959 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:6.0a1) Gecko/20110421 Firefox/6.0a1"
▶ May 20th 2015, 13:28:49.156	50.11.235.131 - - [2015-05-20T20:28:49.156Z] "GET /uploads/nicole-marie-passonno-stott.jpg HTTP/1.1" 200 3323 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:6.0a1) Gecko/20110421 Firefox/6.0a1"
▶ May 20th 2015, 13:26:32.000	152.150.148.109 - - [2015-05-20T20:26:32.000Z] "GET /uploads/joseph-m-acaba.png HTTP/1.1" 200 6106 "-" "Mozilla/5.0 (X11; Linux x86_64; rv:6.0a1) Gecko/20110421 Firefox/6.0a1"
▶ May 20th 2015, 13:23:17.483	163.15.240.221 - - [2015-05-20T20:23:17.483Z] "GET /uploads/sergei-ryazanski.jpg HTTP/1.1" 200 3251 "-" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1; .NET CLR 1.1.4322)"



# Grok Debugger

Dev Tools

Console Search Profiler Grok Debugger

**Input**

```
183.60.215.50 - - [11/Sep/2014:22:00:00 +0000] "GET /scripts/netcat-webserver HTTP/1.1" 200 182 "-" "Mozilla/5.0 (compatible; EasouSpider; +http://www.easou.com/search/spider.html)"
```

**Pattern**

```
%{COMBINEDAPACHELOG}
```

▶ Custom Patterns

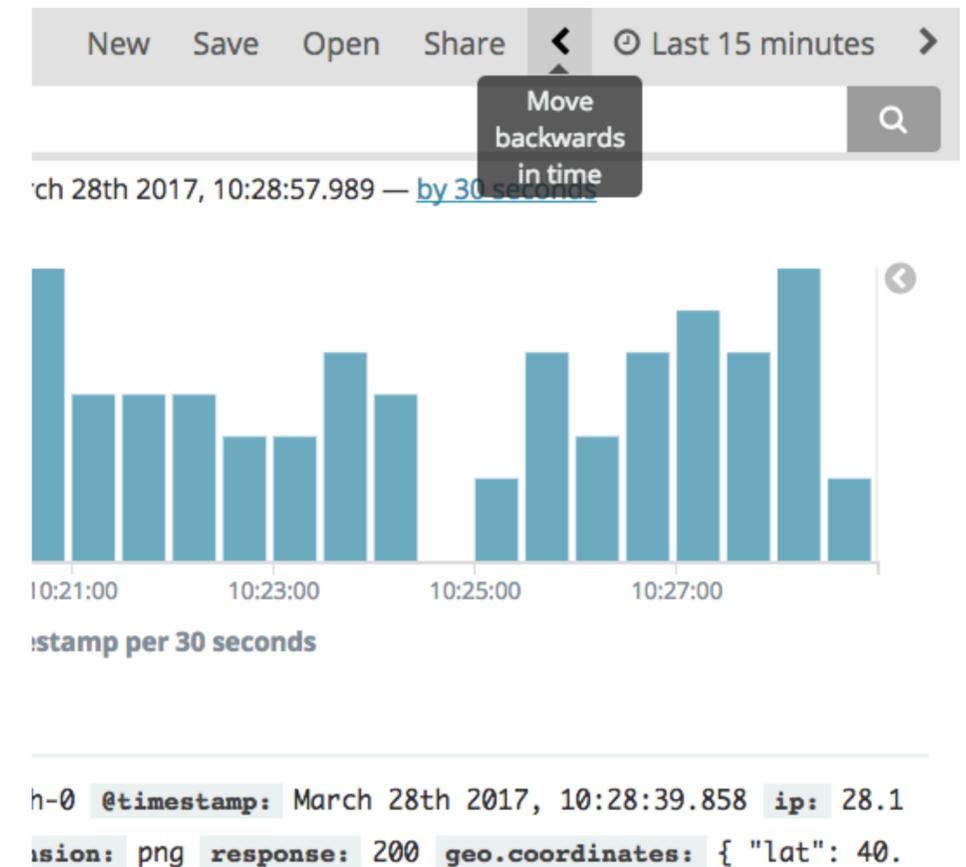
Simulate

**Output**

```
1 {
2   "request": "/scripts/netcat-webserver",
3   "agent": "\"Mozilla/5.0 (compatible; EasouSpider; +http://www.easou.com/search/spider.html)\\"",
4   "auth": "-",
5   "ident": "-",
6   "verb": "GET",
7   "referrer": "\"-\"",
8   "response": "200",
9   "bytes": "182",
10  "clientip": "183.60.215.50",
11  "httpversion": "1.1",
12  "timestamp": "11/Sep/2014:22:00:00 +0000"
13 }
```

# User Experience Improvement

- Easier time scroll with the new timepicker
- Simplified “add to dashboard” flow
- Expand a visualization in a dashboard
- Targeted document highlighting
- Easy filter on document table values
- Geo-Centroid

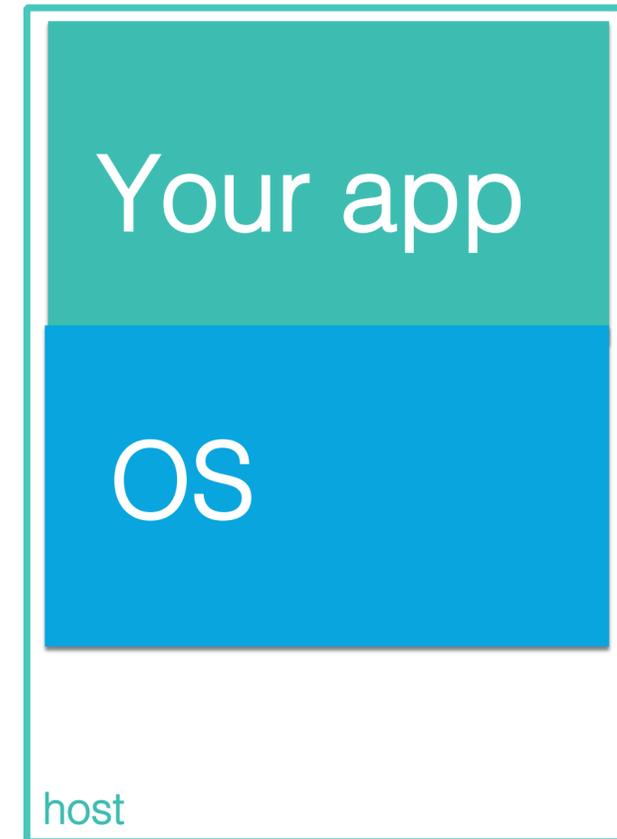
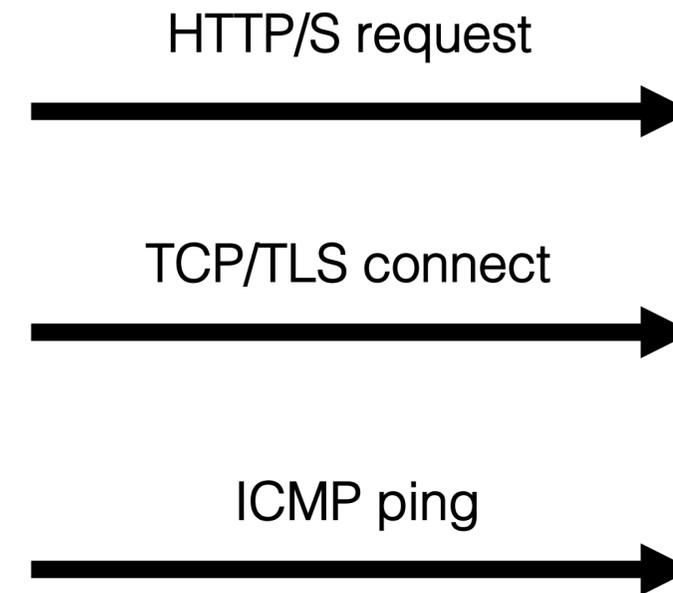




# Beats since 5.0

# New Beat: Heartbeat (beta in 5.2)

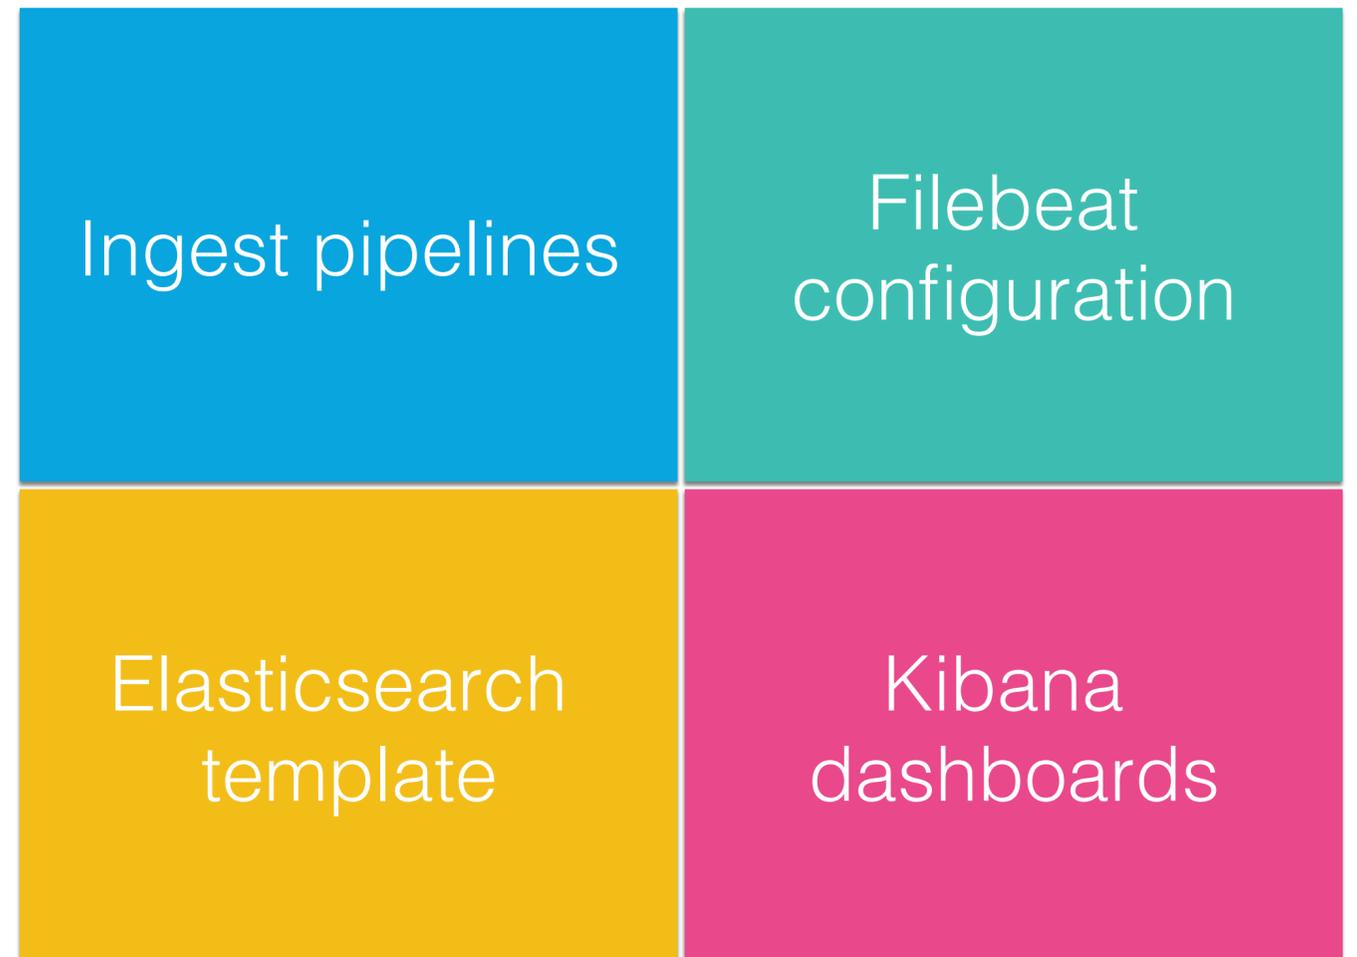
- Ping all the things
- Gather round trip metrics
- Many to many
- Ping IPs behind load balancers



# Filebeat Modules (5.3)

They are like Metricbeat modules...only different

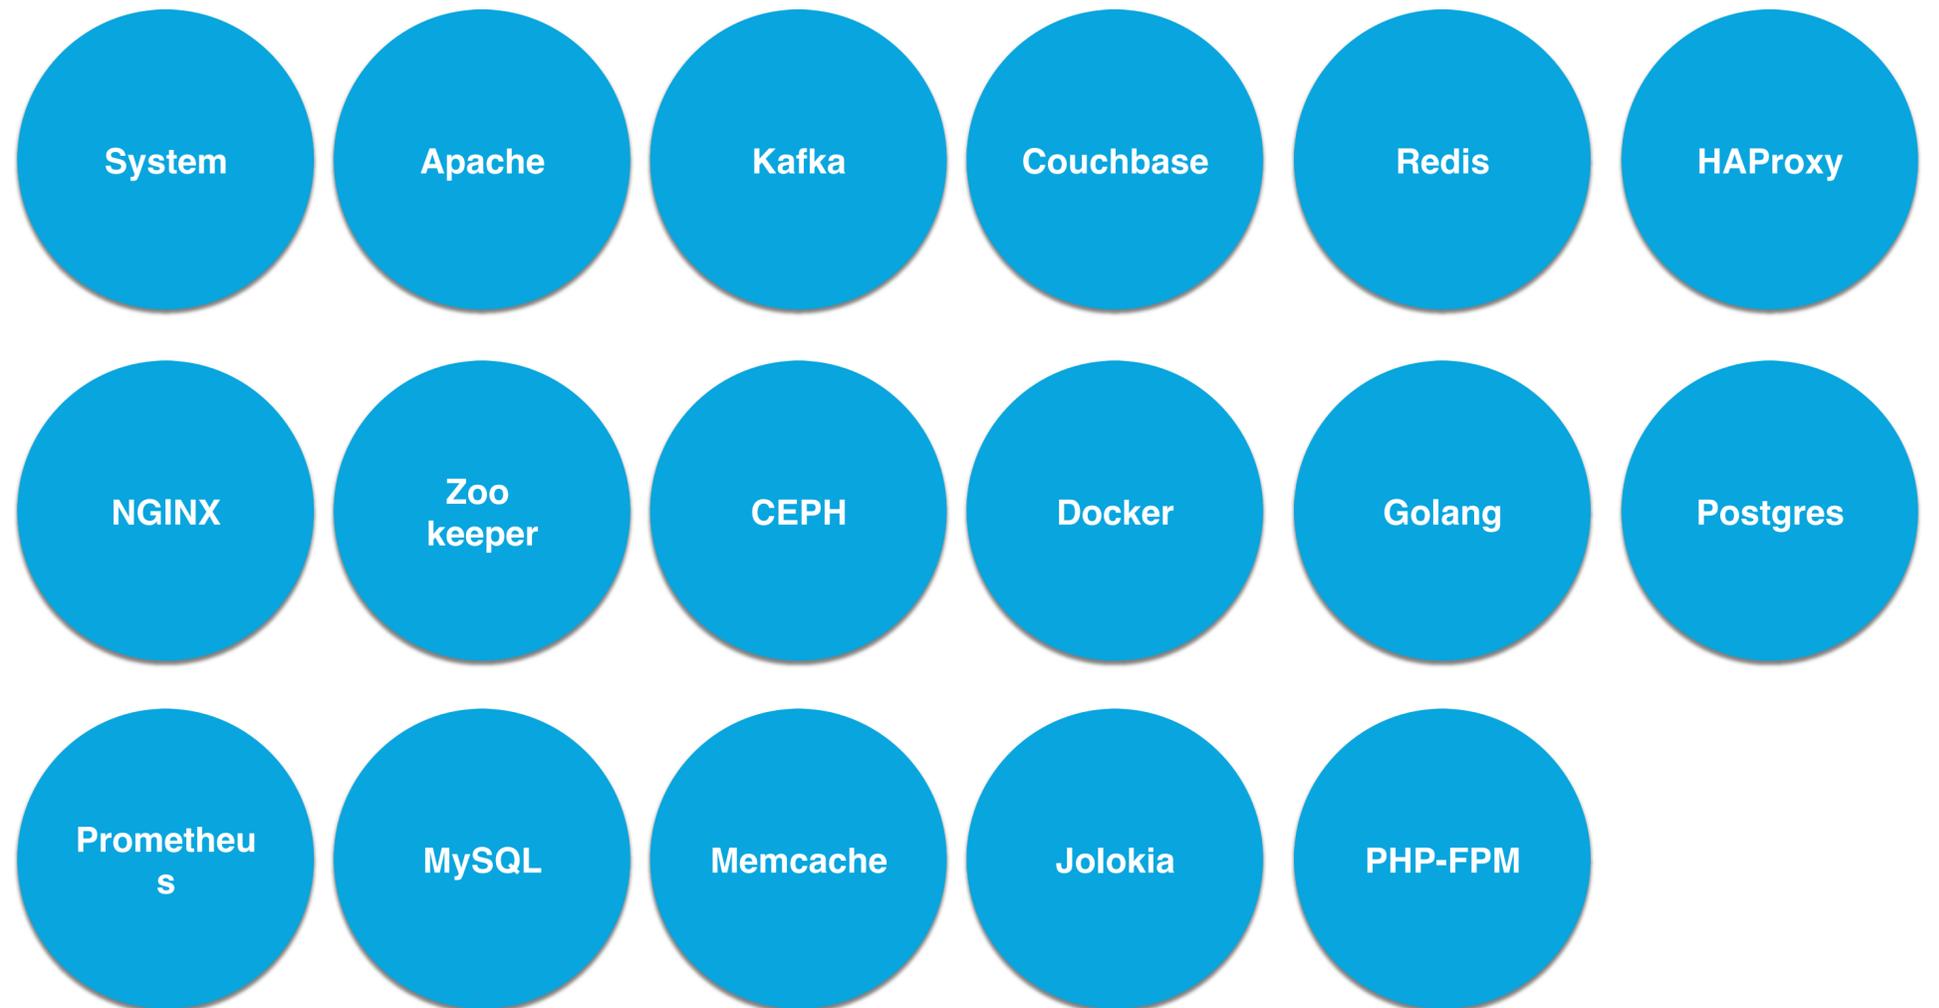
- Because simple things should be simple
- Prepackaged configs for common log formats
- Get to a dashboard in minutes
- First release includes Apache, Nginx, MySQL, system modules. More to come.



# New Modules in Metricbeat

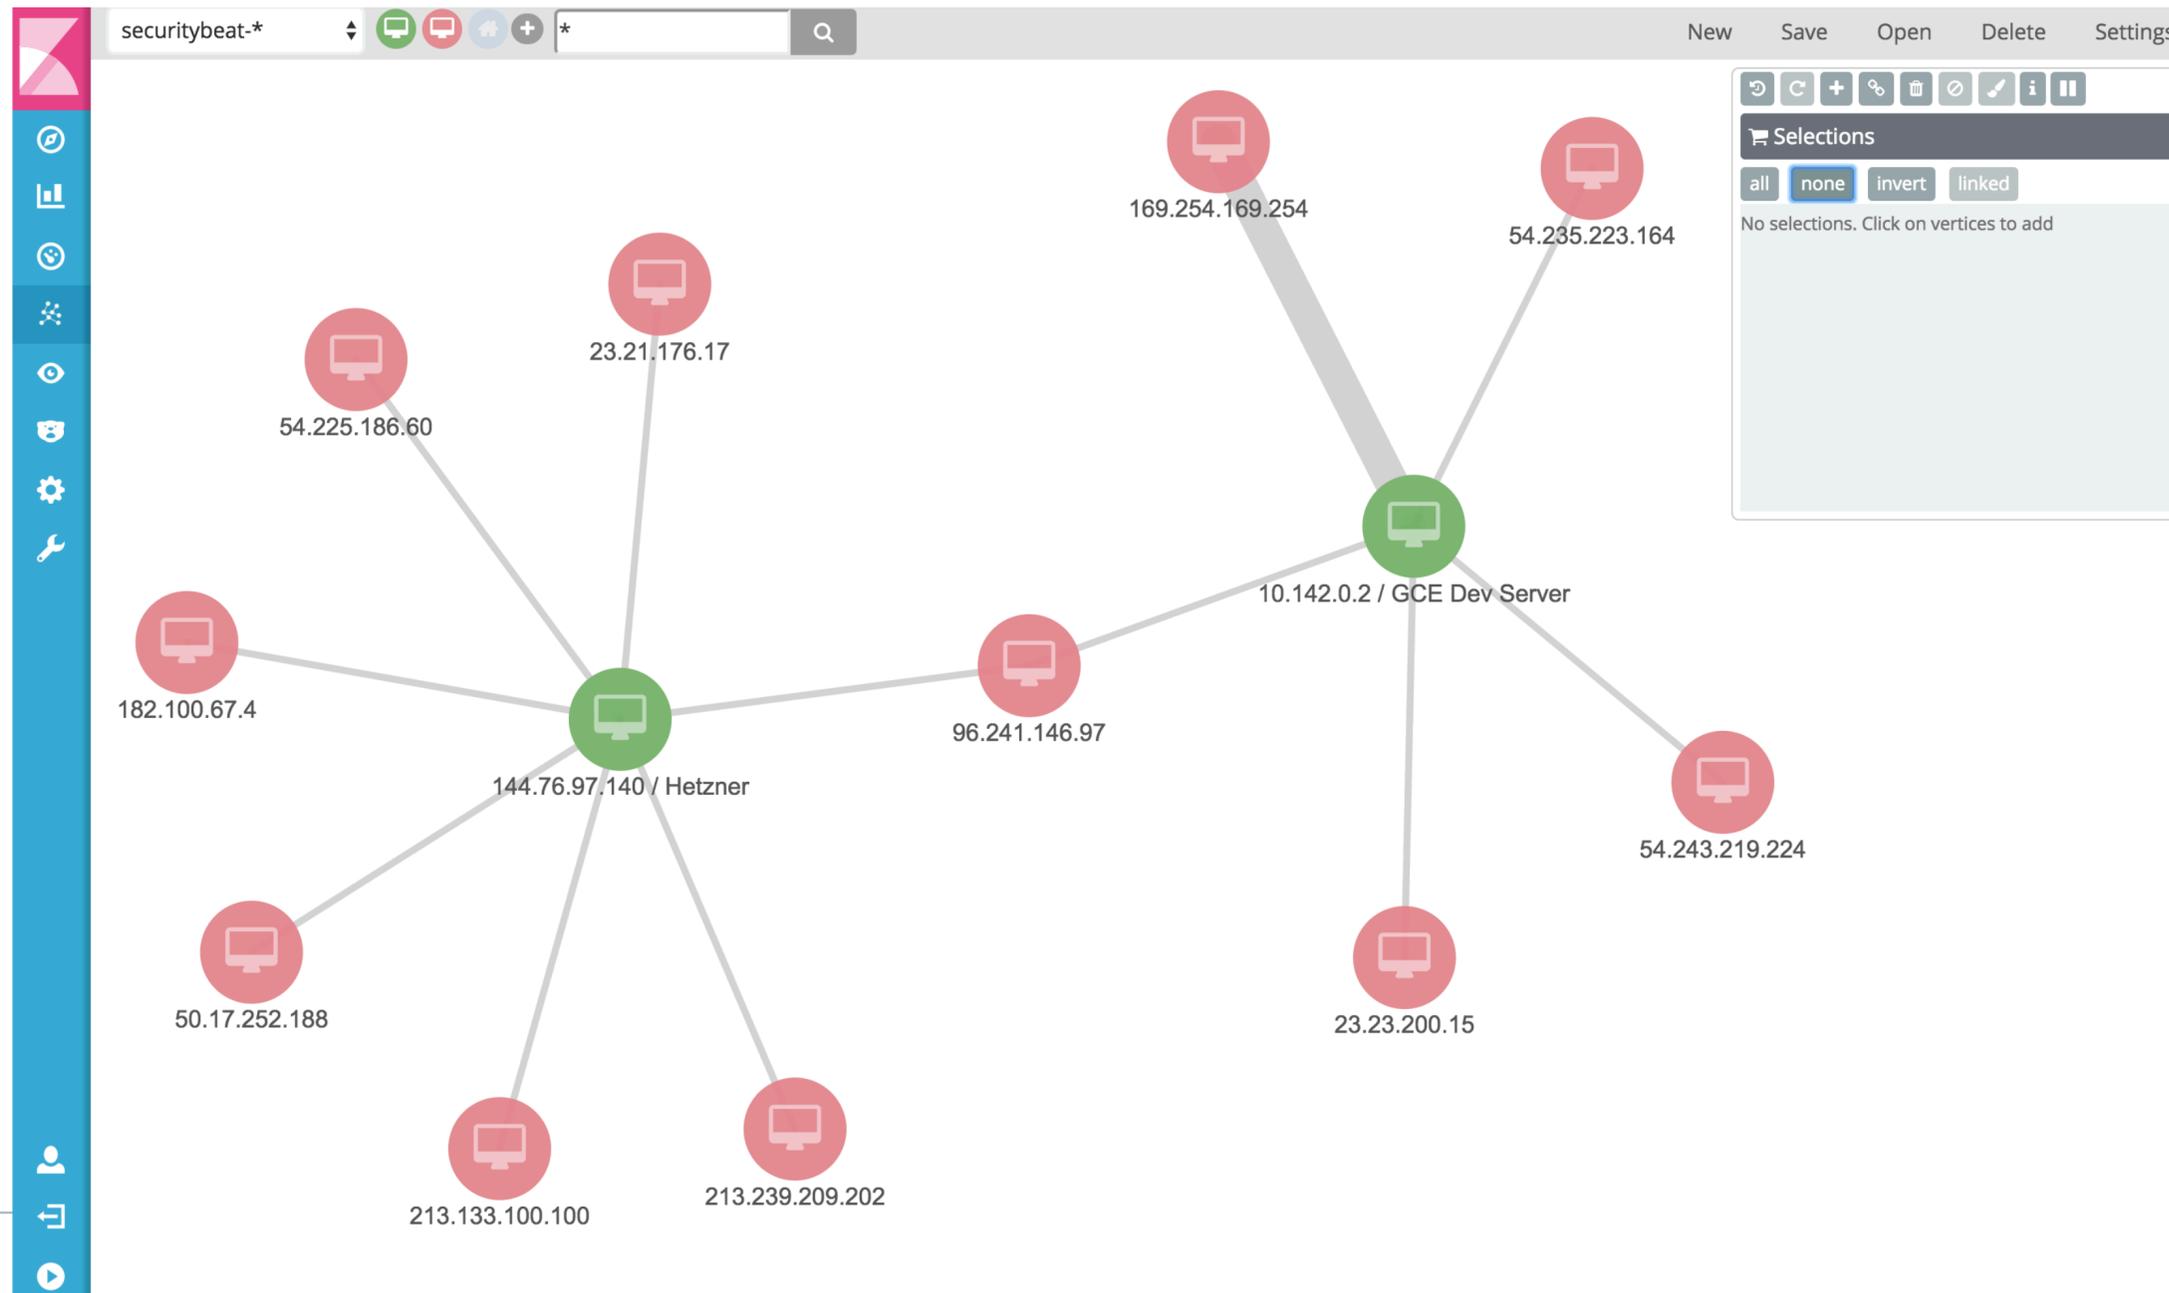
17 modules and growing

- Docker (5.1)
- Kafka (5.1)
- Prometheus (5.2)



# Track Network Connection

## Metricbeat system modules adds TCP Socket monitoring (5.2)



# Monitor applications running in cloud

add\_cloud\_metadata Beat processor



Google Cloud Platform





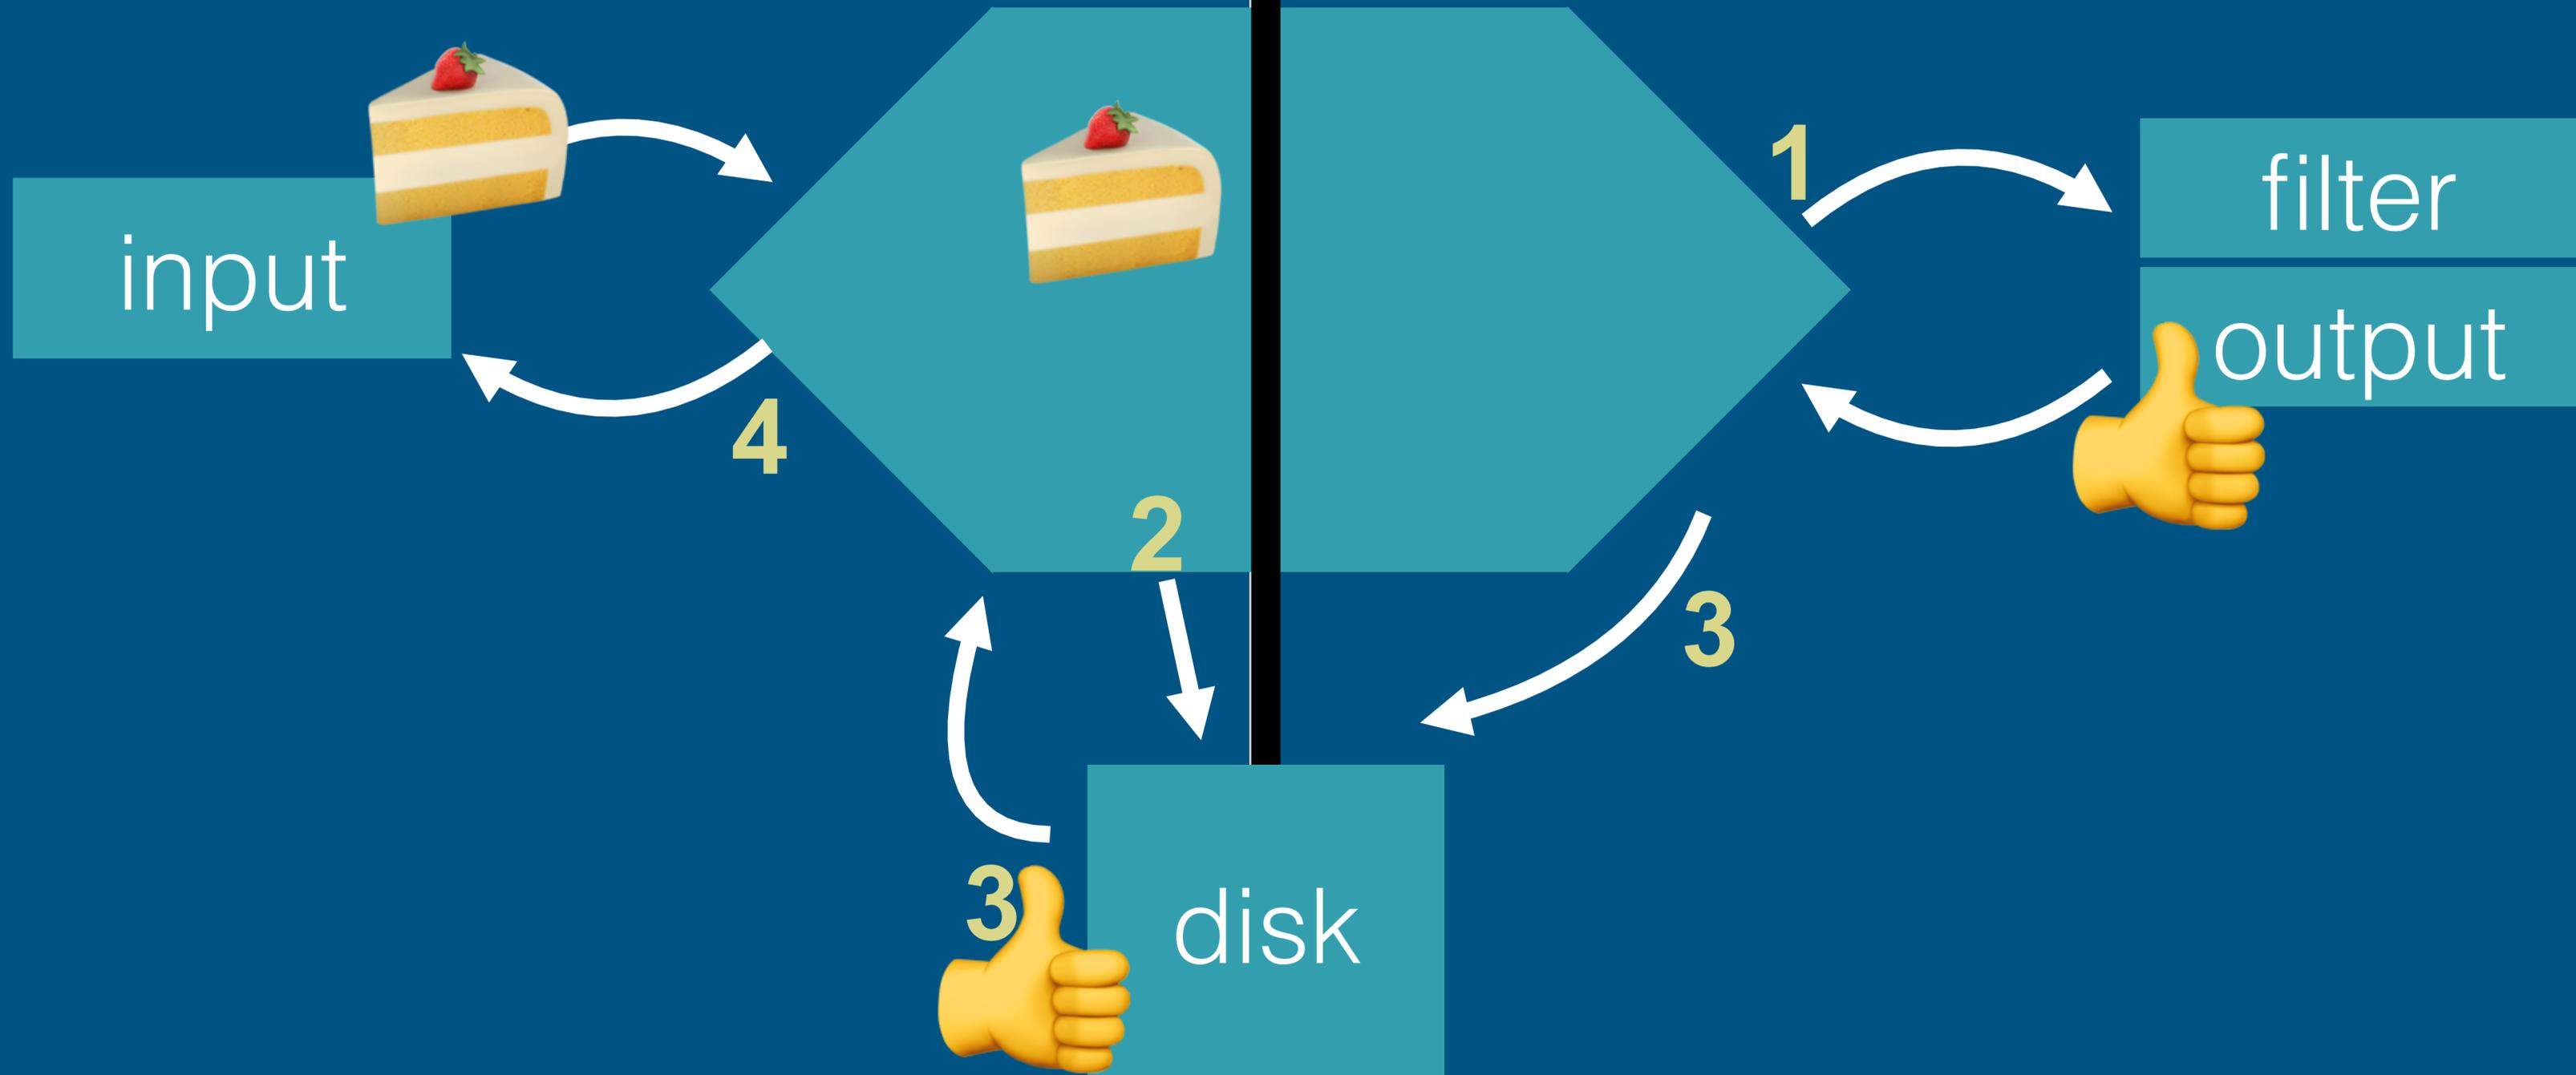
# Logstash since 5.0

# Centralised Management

- Elasticsearch as a remote config store
- Manage configurations via UI
- Group multiple Logstash under roles
- Simple alternative to puppet, chef

# Persistent Queues (beta in 5.2)

- ✦ Survives (temporary) machine failures
- ✦ Adaptive, native buffering to handle ingestion spikes
- ✦ Limited impact on performance
- ✦ View queue stats in monitoring UI



# Offline Plugin Management (5.2)

## Air-gapped Networks and Offline Environments

### Prepare and Pack Plugins on Staging Box

```
$ bin/logstash-plugin prepare-offline-pack logstash-filter-* logstash-input-beats
```

### Move Offline Pack to Offline Boxes

- Default pack location: `/LOGSTASH_HOME/logstash-offline-plugins-5.2.0.zip`
- Change pack location using `--output /path/to/pack` parameter

### Install or Update Plugins

```
$ bin/logstash-plugin install file:///path/to/logstash-offline-plugins-{logstash_version}.zip
```

# Database Lookup Enrichment

## JDBC\_streaming filter (5.3)

- Enrich Logstash events with DB data (streaming joins)
- Executes JDBC lookup queries per event (add one or more fields)

```
jdbc_streaming {  
  jdbc_connection_string => "jdbc:mysql://localhost:3306/mydatabase"  
  statement => "select * from PRODUCTS.FRUITES WHERE SKU = :sku"  
  parameters => { "sku" => "sku_code"  
  }  
}
```



# Pipelines: Visual Builder

Management / Logstash / Pipelines

**kibana**

Pipelines

- Discover
- Visualize
- Dashboard
- Timelion
- Machine Learning
- Graph
- Dev Tools
- Monitoring
- Management

elastic  
Logout  
Collapse

stdin (testStdin)

beats (filebeat)

grok (parseApache)

date (apacheDate)

condition (tag == \_grokparsefailure)

true false



# X-Pack since 5.0

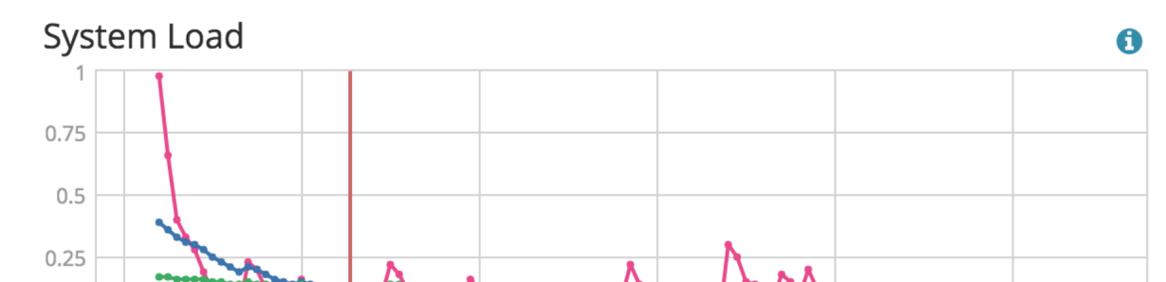
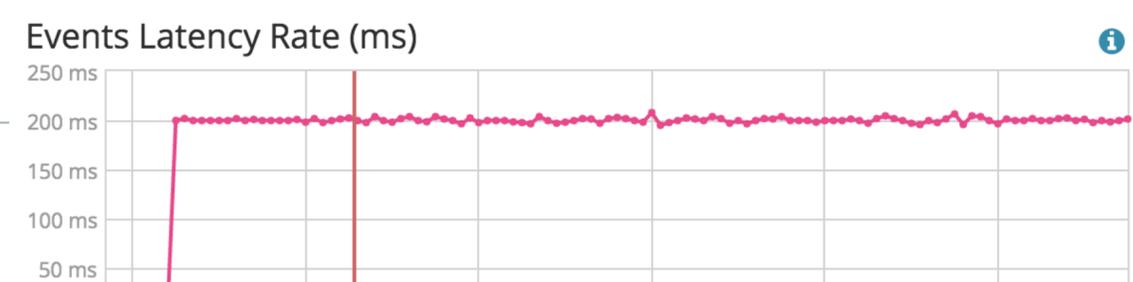
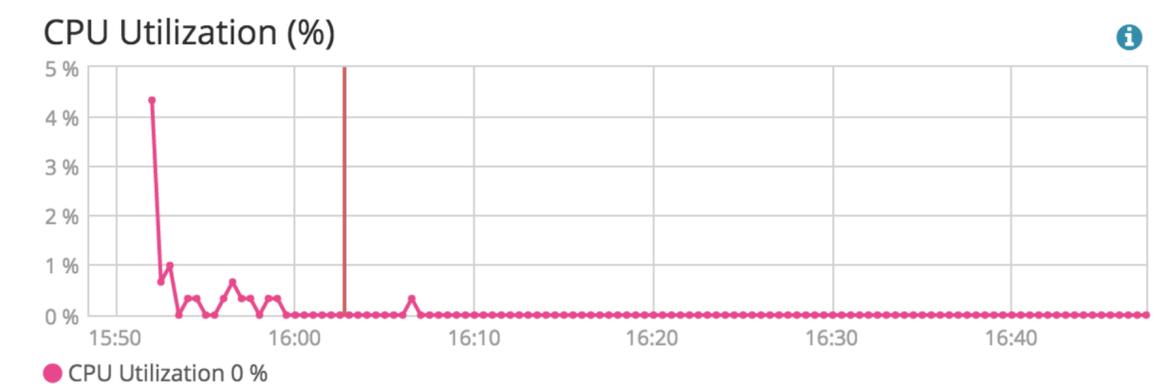
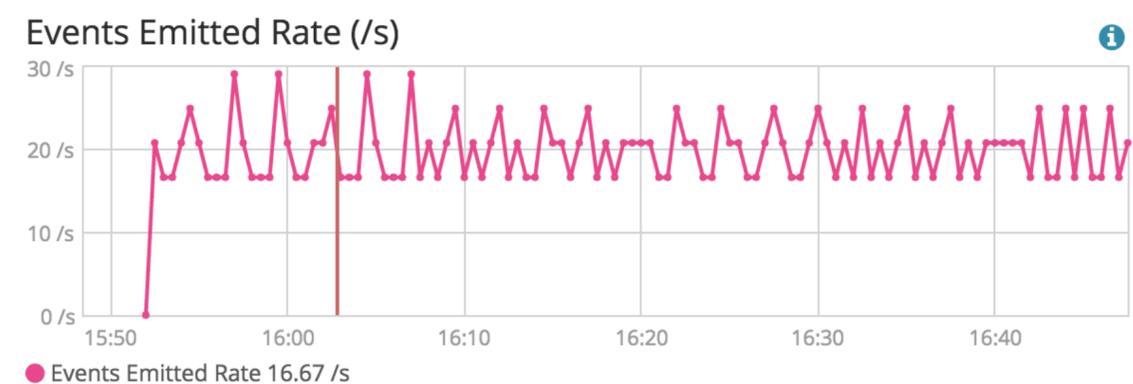
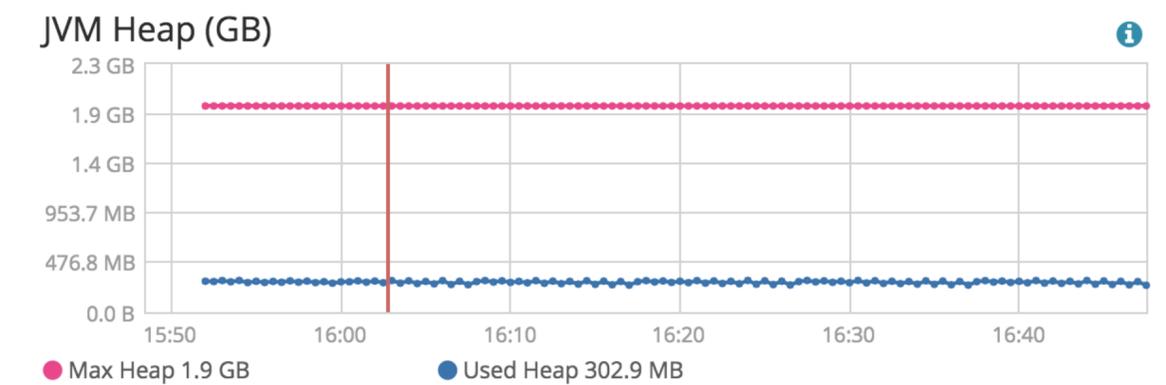
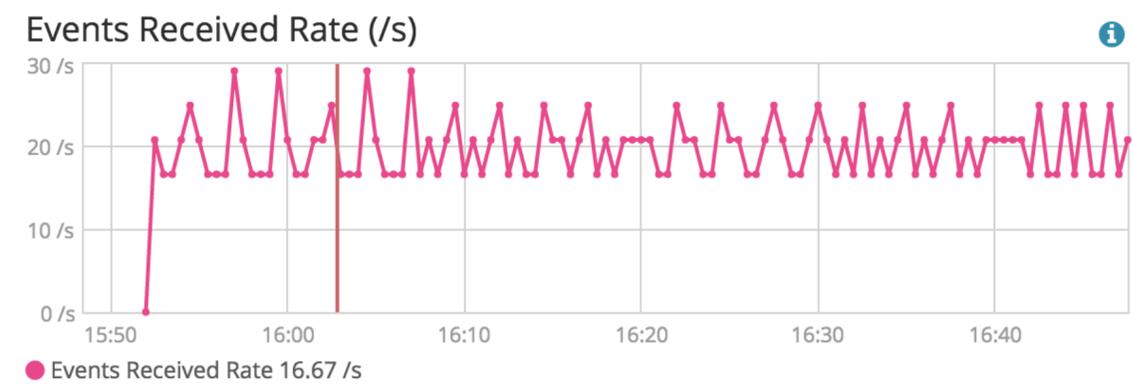
# Monitoring Enhancements

## Logstash Monitoring UI added to X-Pack

Clusters / monitoring-cluster / Logstash / Nodes 10 seconds ◀ Last 1 hour ▶

logstash-node1

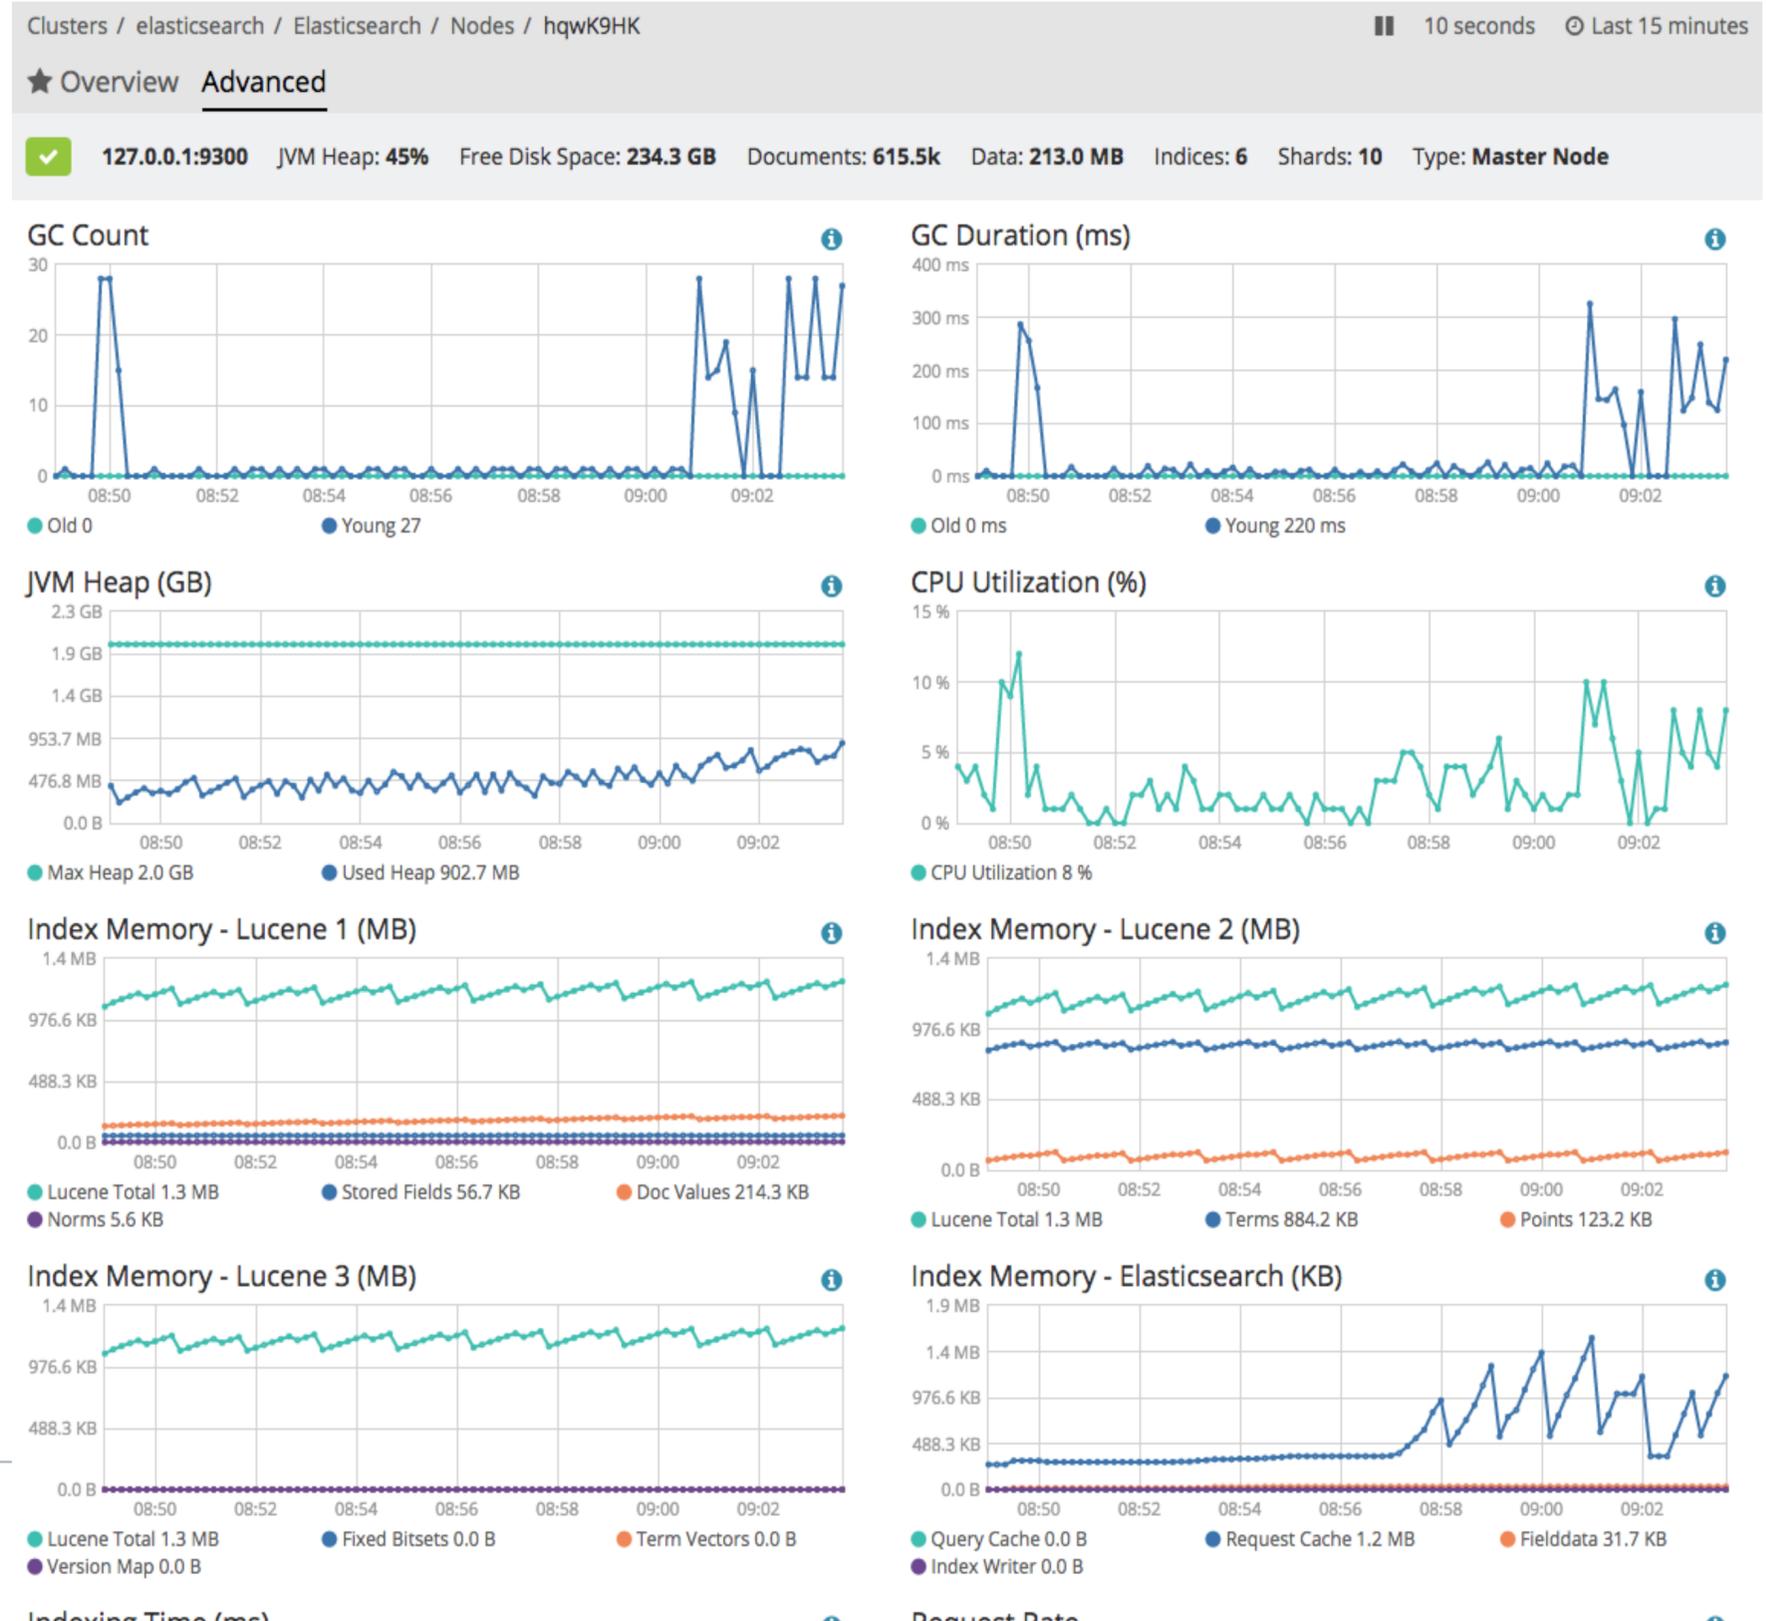
**127.0.0.1:9600** Events Received: **67.4k** Events Emitted: **66.9k** Config Reloads: **0** Pipeline Workers: **4** Batch Size: **125** Version: **5.2.0**



# Monitoring Enhancement

## More charts! Better Charts!

- Improved Charts with Multiple Series (5.0)
- More charts! Advanced Node and Index views (5.1)
- Capture cgroup (Container) metrics for Elasticsearch (5.2)



# Security: New since 5.0

- Utility to simplify TLS Certificate Generation Process
- Consistent TLS configuration across the stack
- TLS required for node-to-node transport
- Goodbye default passwords
- Goodbye clear-text passwords in config files

# Graph: New since 5.0

- Explore across multiple indices
- Simplified field configuration
- Saveable/shareable workspaces
- Deep linking in to Graph
- Deep linking out of Graph

# Machine Learning

Already released in 5.4



# New job from index pattern it\_ops-kpi

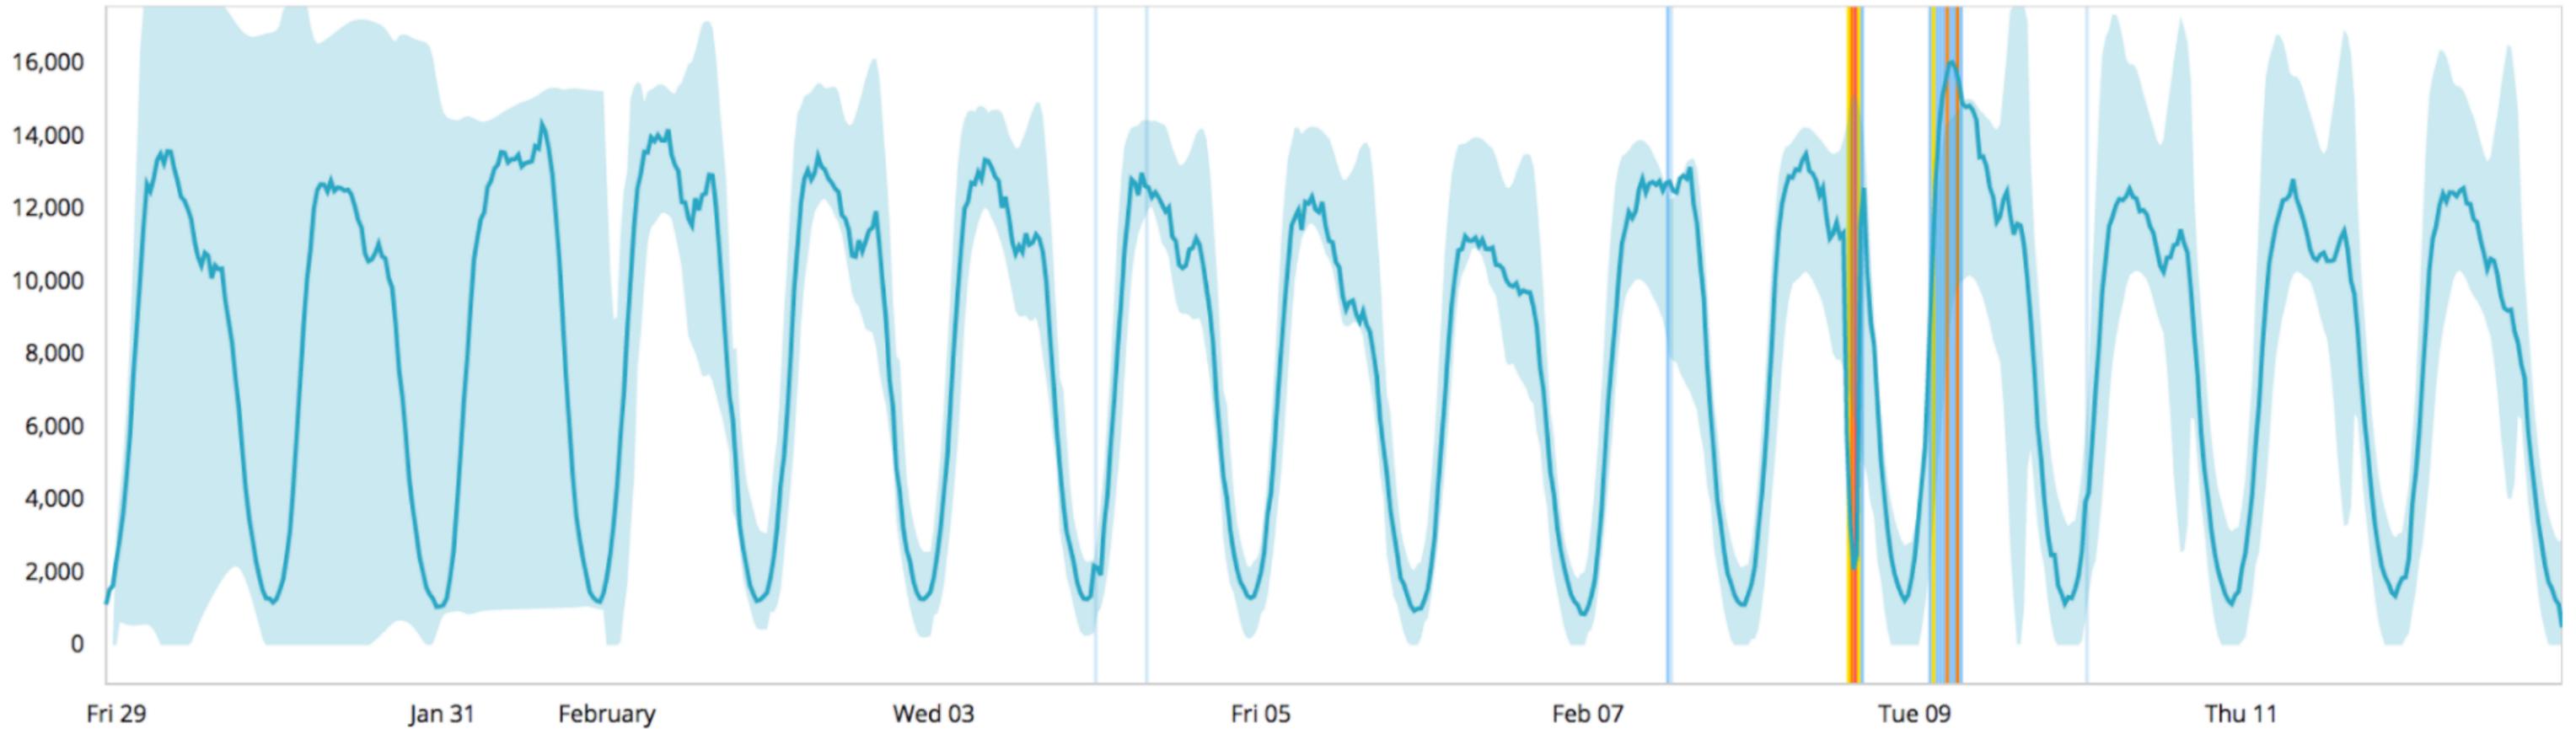
Aggregation ⓘ

Count

Field ⓘ

Bucket span ⓘ

5m



Job it\_ops-kpi created

Reset

View Results

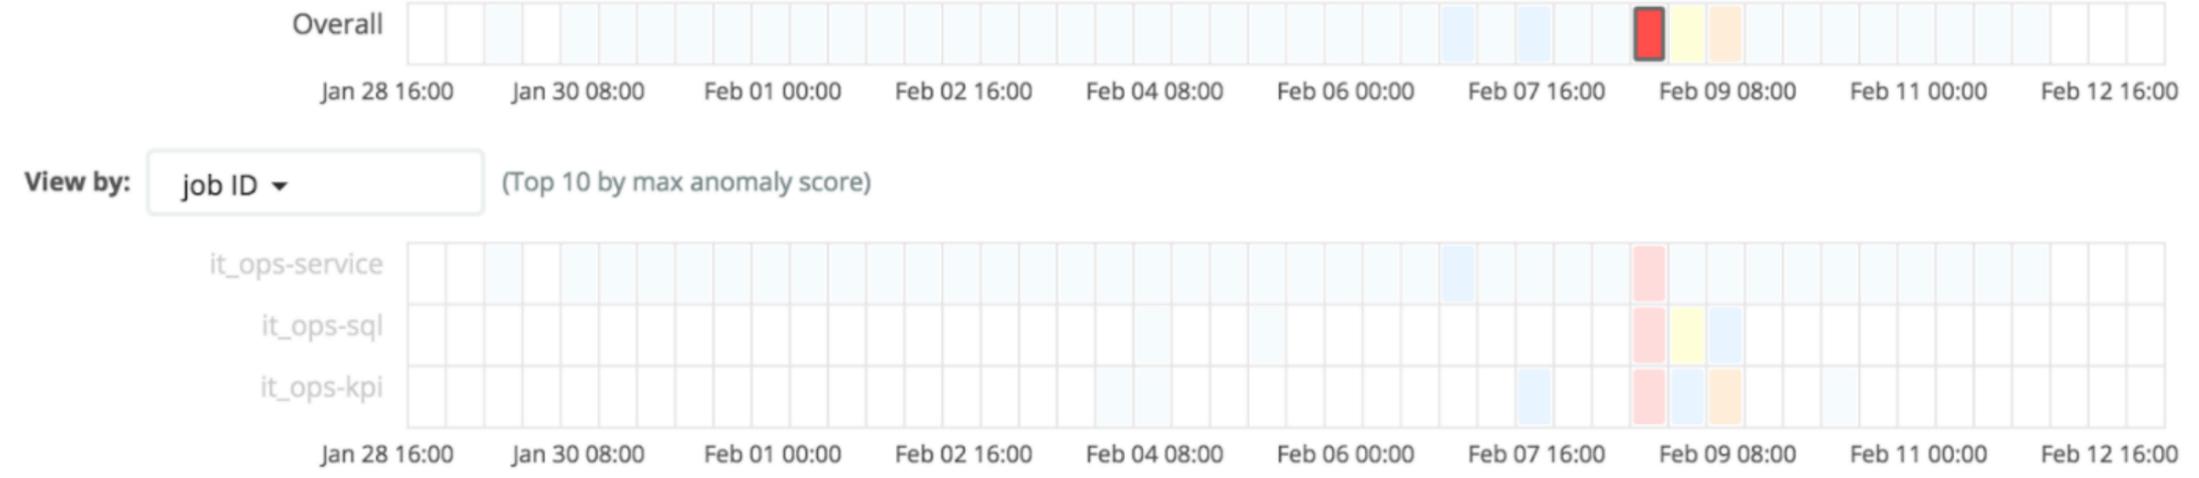


Job it\_ops-kpi and 2 others

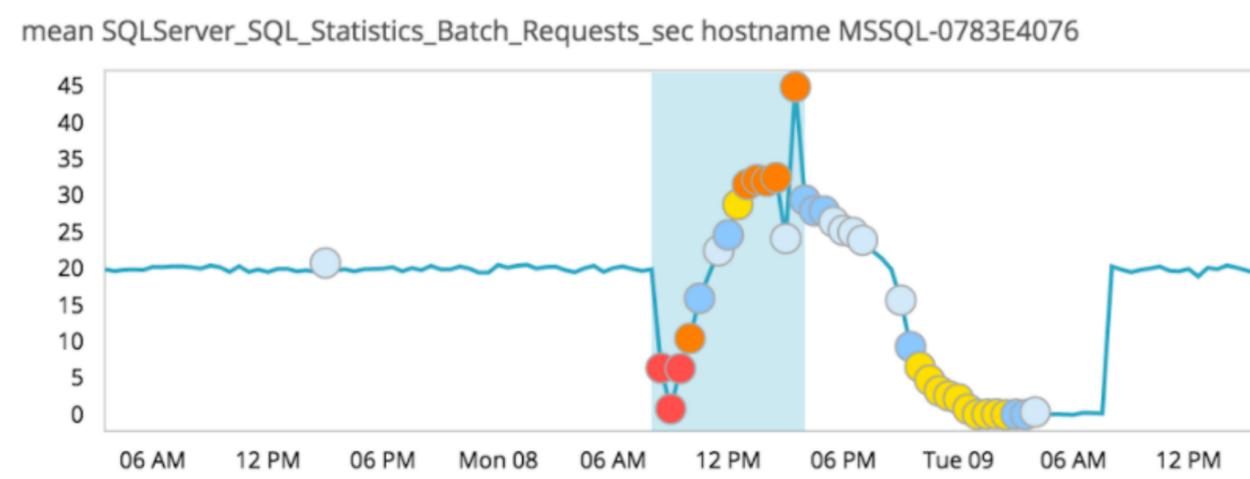
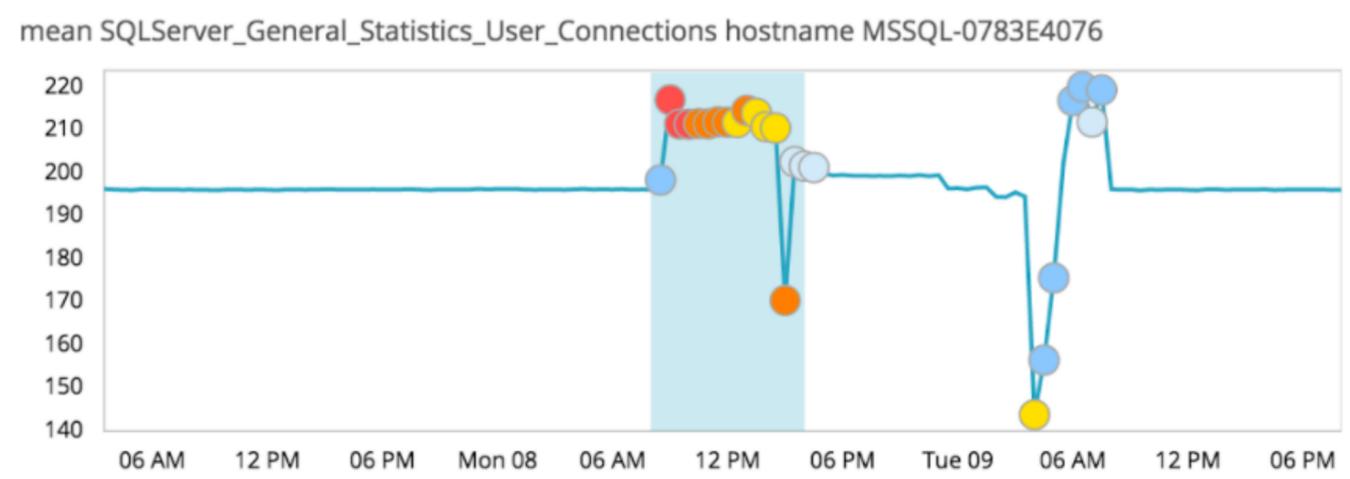
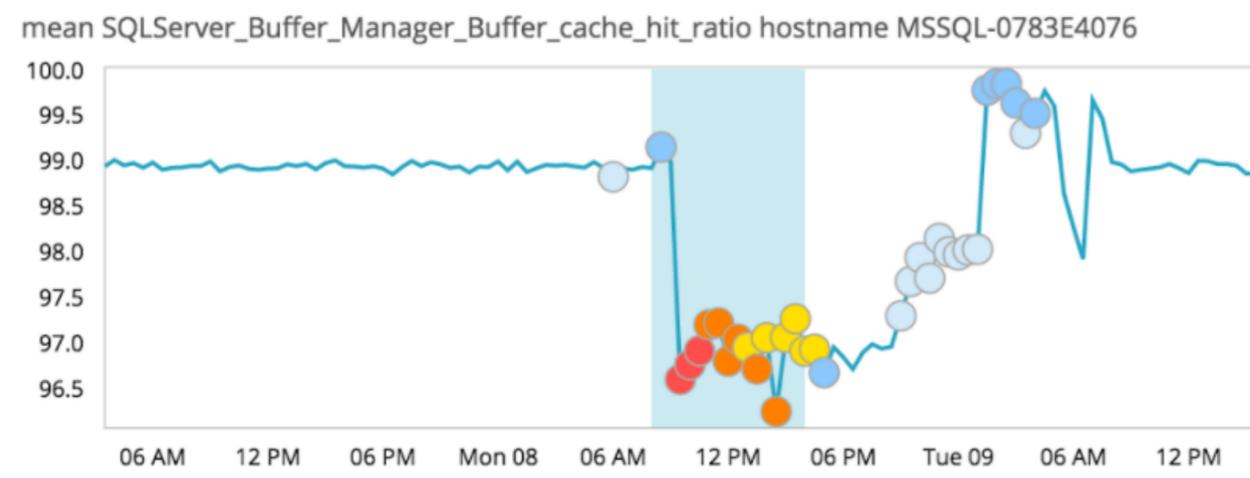
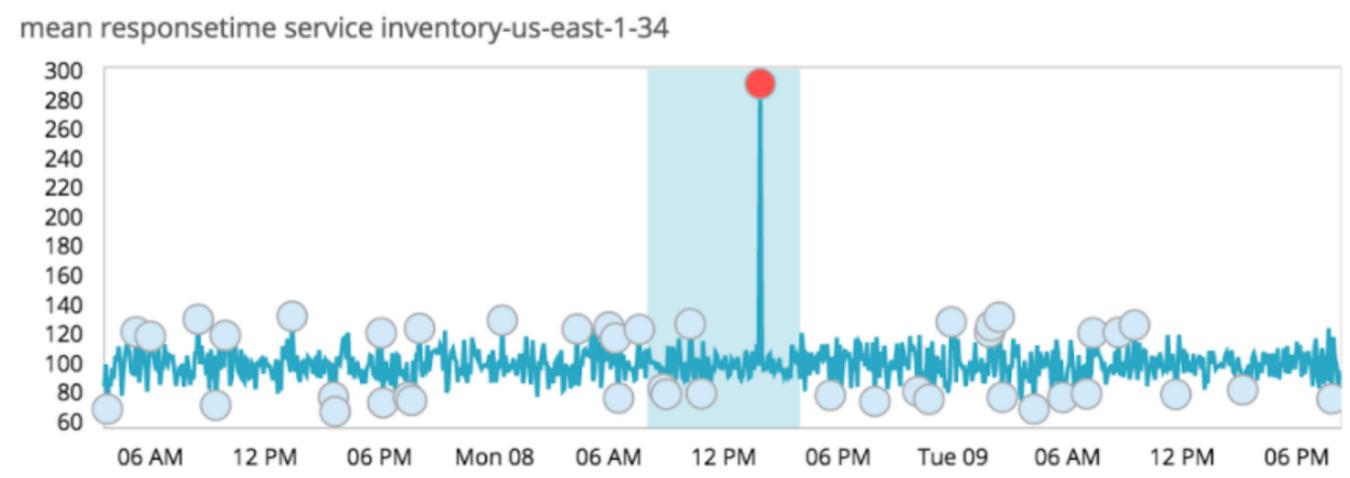
### Top Influencers

service	Count	Score
inventory-us-east-1-34	94	97
auth-us-west-1-1e	7	51
test-srv-02	5	29
elasticsearch-22	3	16
elasticsearch-77	2	11
payment-srv-21	2	6
payment-srv-11	1	5
backup-srv-13	1	9
test-srv-01	1	7
inventory-us-west-1-4e	1	7
hostname	Count	Score
MSSQL-0783E4076	94	1237

### Anomaly timeline



### Anomalies



# Coming soon



# Elastic Cloud Enterprise

1.0.0-beta2 released



# Elasticsearch SQL

Coming soon!



# Kibana Canvas

Coming soon!

Expression >

Axis Style >

Text along the x and y axis

B I U [font icons]

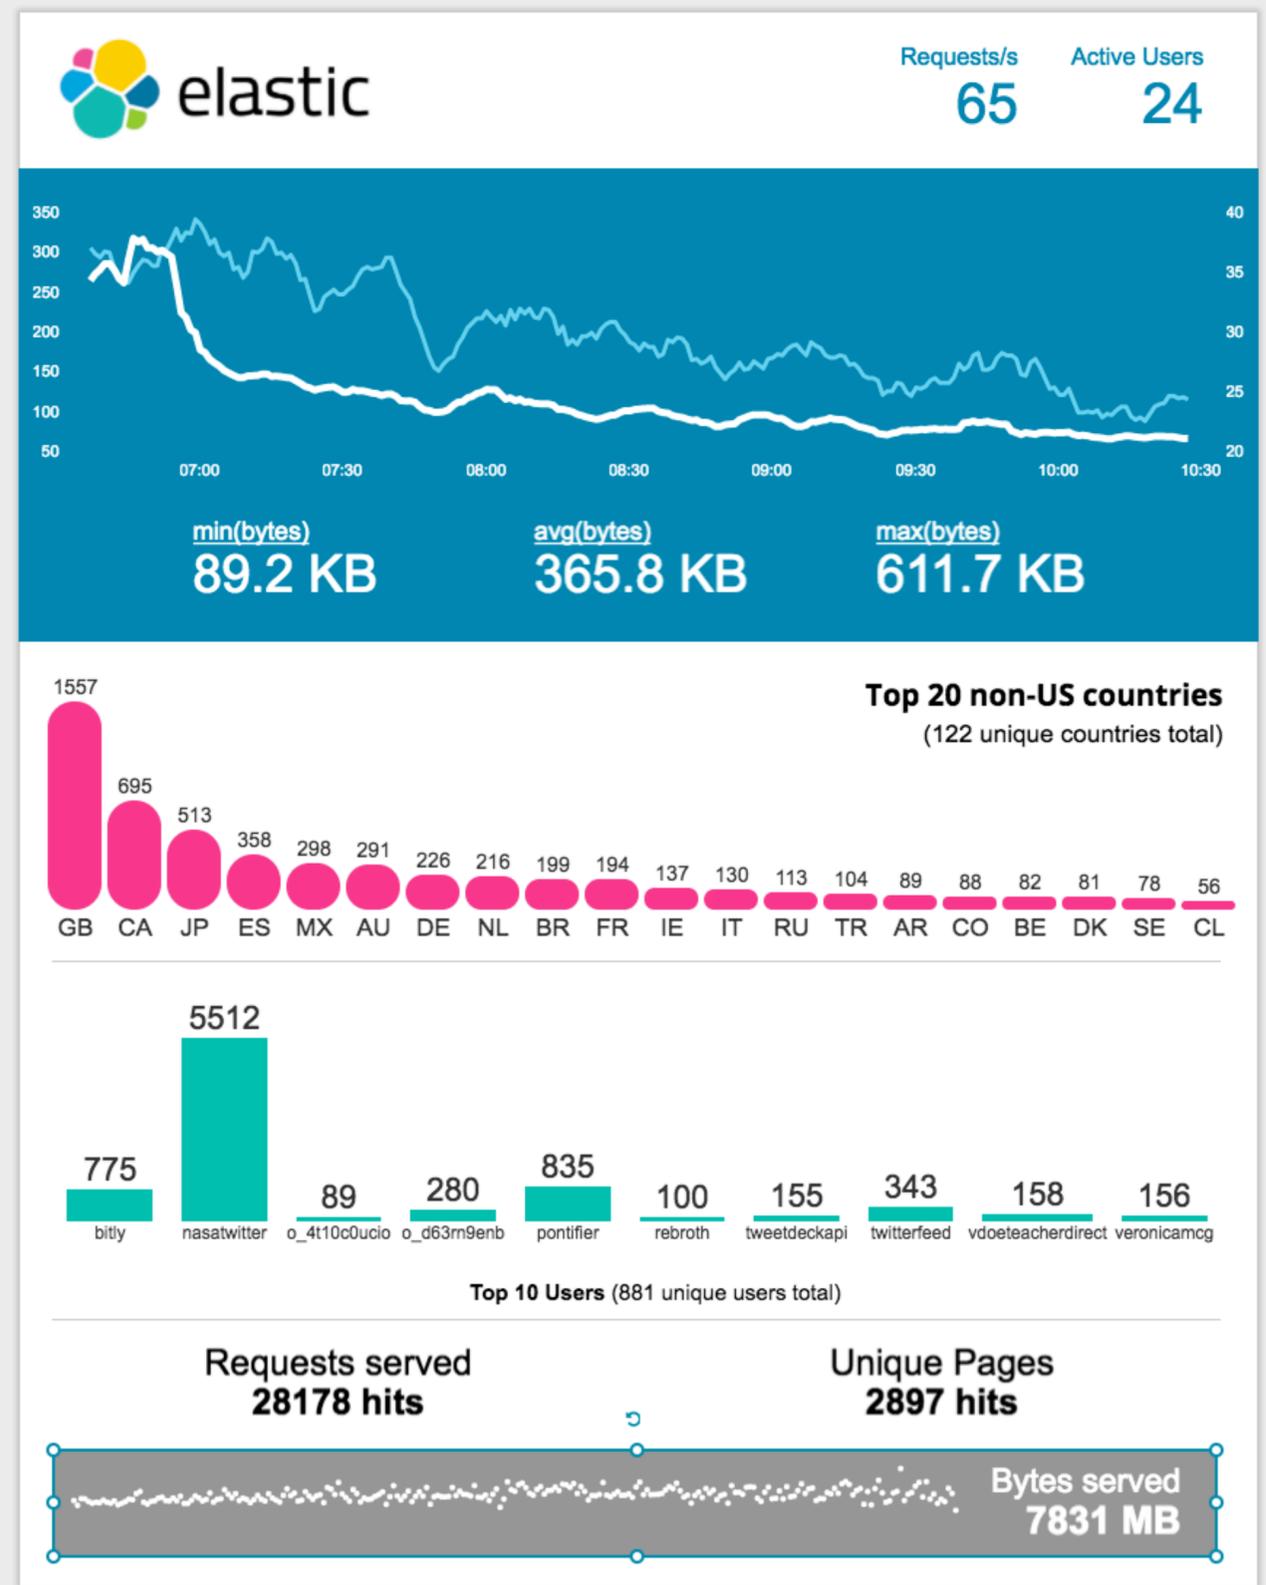
Font Face Size Color

Grid Color >

Panel Style >

Custom Style Sheet >

Arrangement >



- New visualization application on top of Elasticsearch data
- Use Case:
  - live infographics
  - presentations with live data feeds
  - highly customized reports
- Currently, in the prototyping phase
- Release date: TBD

Page Background

Image Lightbox Vertical\_b...

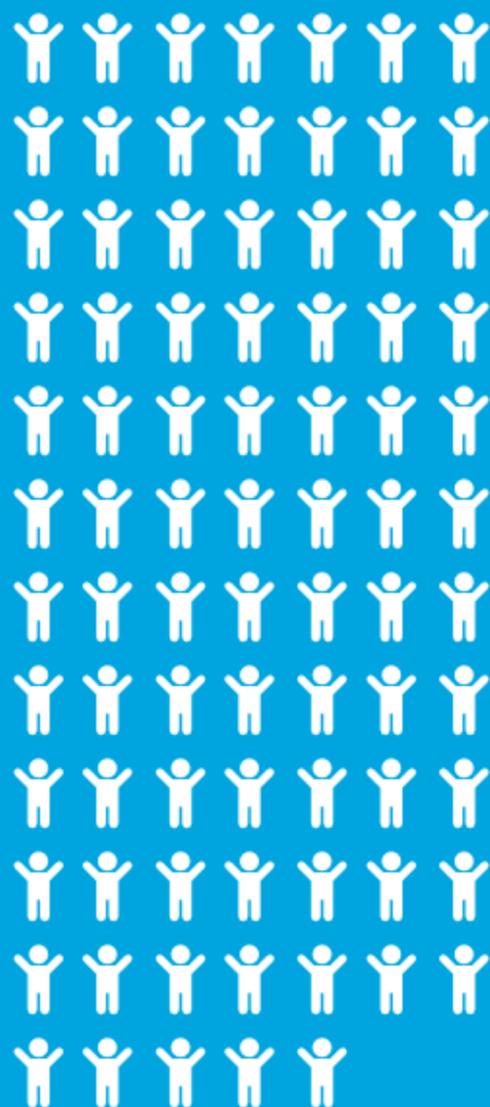
Horizontal... Variable S... String Input

Number S... Circles Grow

Number Repeat Markdown

Timechart Debug

82% of those surveyed have participated in volunteer activities



12% 18-35

While the Deloitte study targeted younger employees, aged 21 - 35, other surveys have found that volunteers are spread fairly evenly across all age ranges with 12% under 25 years of age and 34% over 55. A full 51% were in the prime working age range of 25 to 54

51% 25-54

34% 55+

# VOLUNTEERISM

from Deloitte's Volunteer Impact study on employee engagement and volunteering activities



82%

Volunteer because it...



77%

Improves the community



77%

Makes me happy



63%

Lowers stress levels



48%

Helps meet people



21%

Helps manage an illness



17%

Provides professional development

We can easily quantify the value we give, in units of hours and resources. The part we're less likely to calculate is the benefit we receive from volunteering. We donate our time and skills for many reasons, but community involvement benefits us both in our work life and our personal life.

# 6.0.0 alpha1