Elastic Stack 5.x and beyond …

Medcl

Elastic Stack & X-Pack

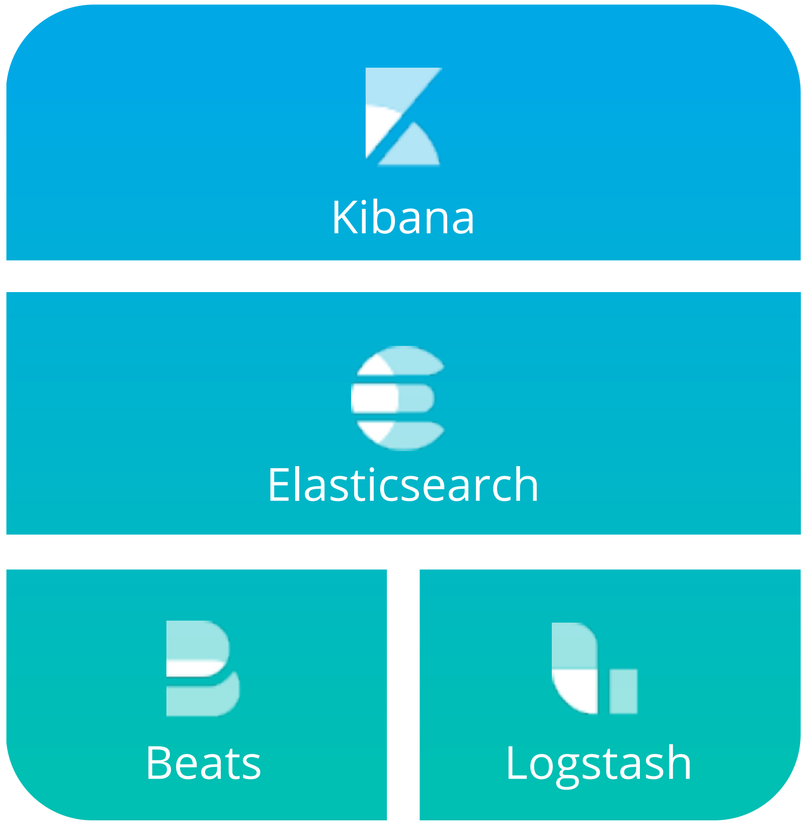

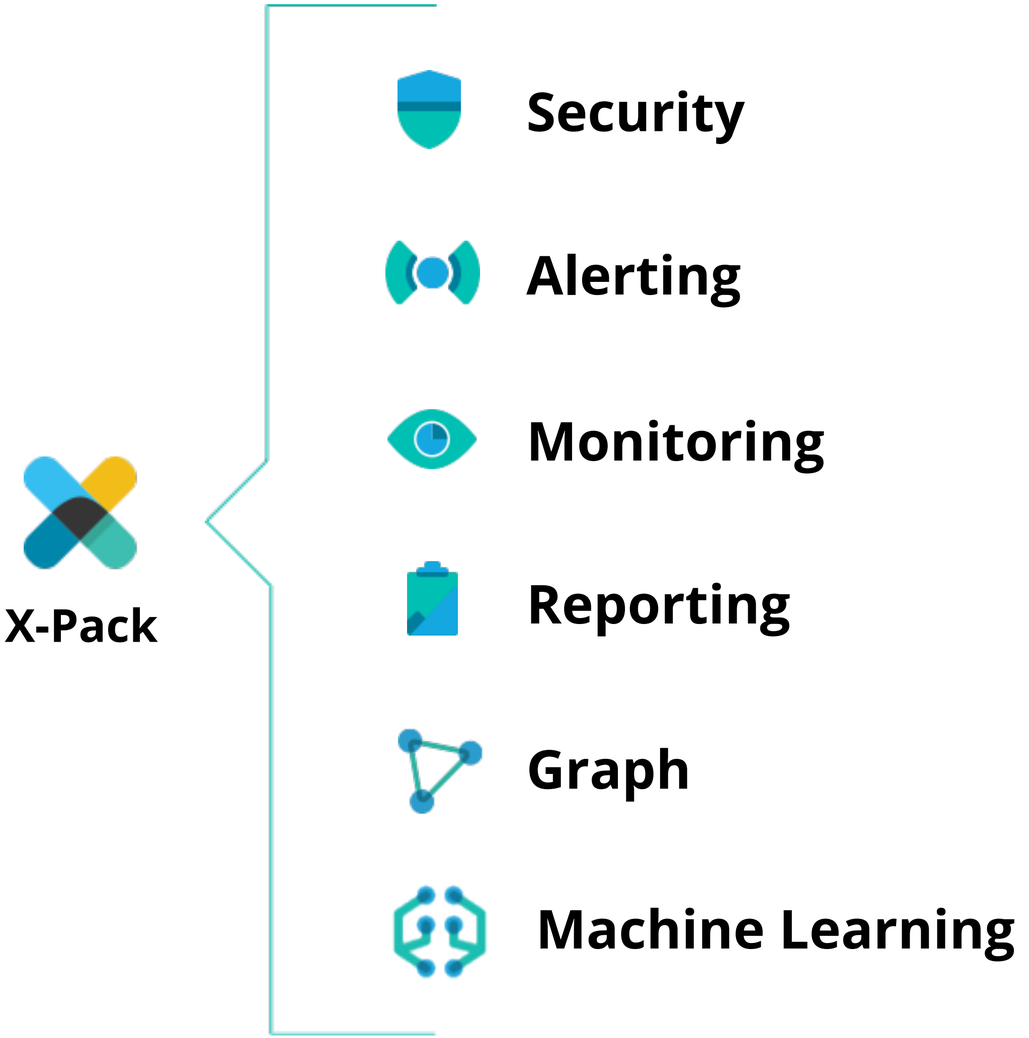



Elasticsearch since 5.0

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

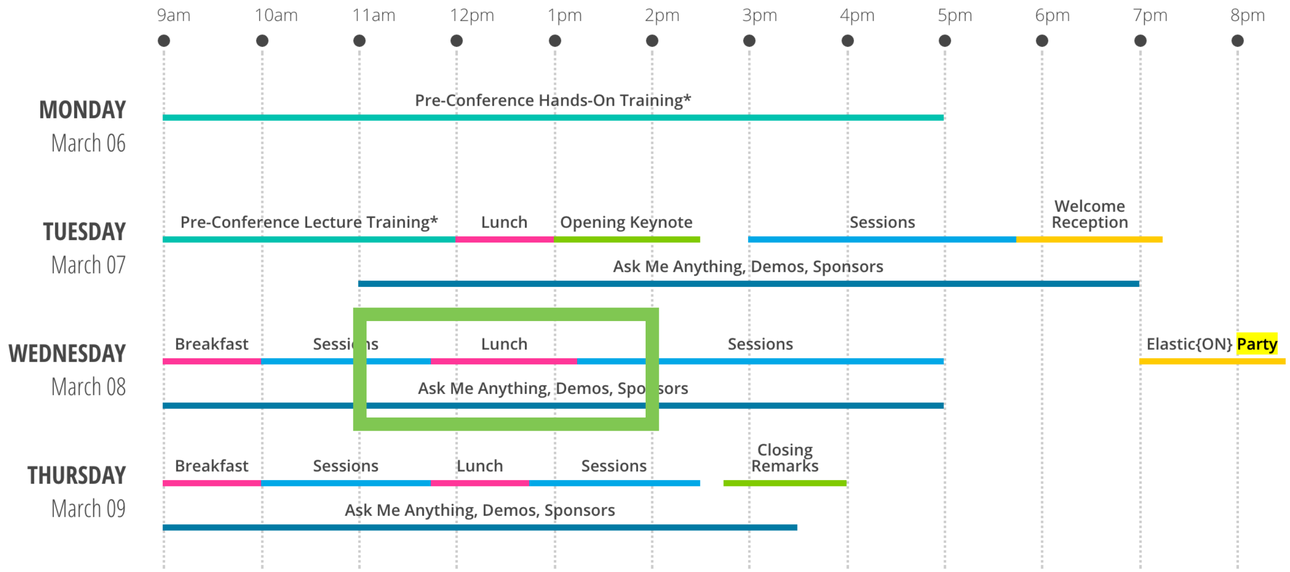

Keyword Normalizer (5.2)

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`

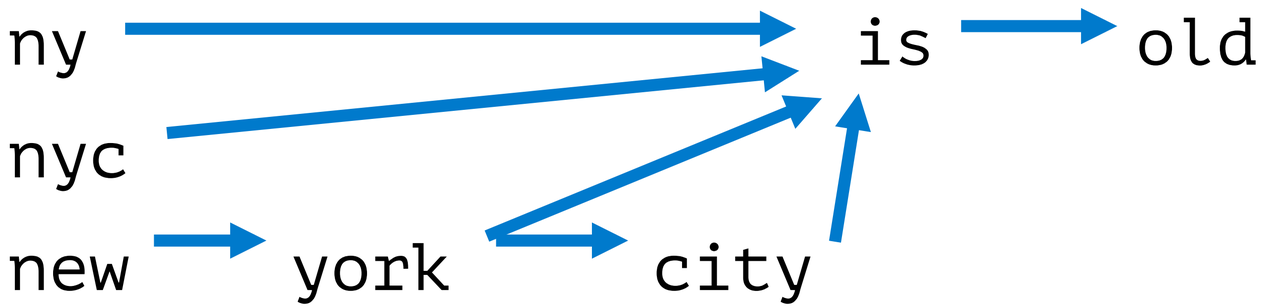

Cluster Allocation Explain API (5.2)

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", ❶
        "inner_hits": {
            "name": "last_tweets", ❷
            "size": 5, ❸
            "sort": [{ "date": "asc" }] ❹
        },
        "max_concurrent_group_searches": 4 ❺
    },
    "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
- Adjacency matrix



Kibana since 5.0

New visualizations

Tag Cloud (5.1)

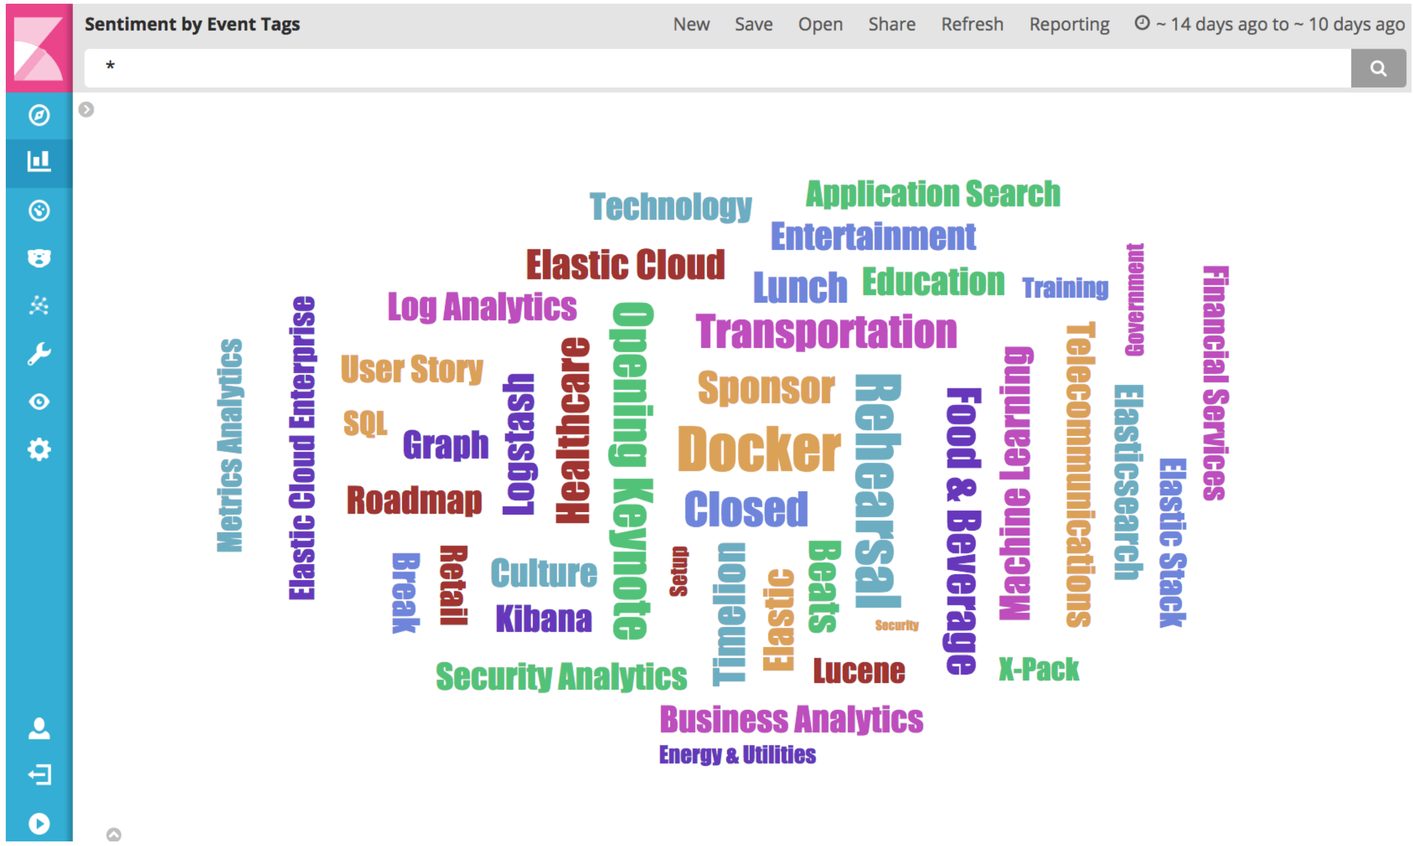

New visualizations

Heatmap (5.2)

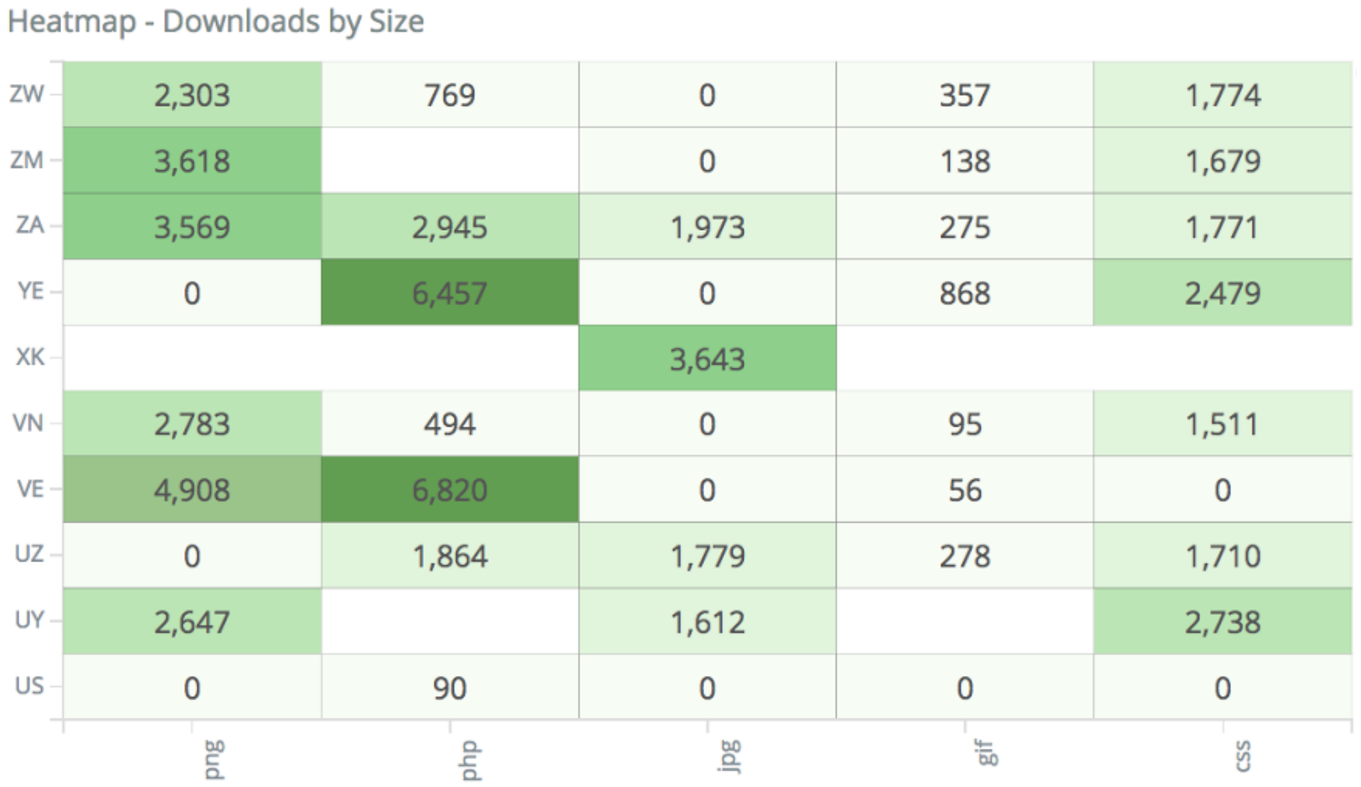

New Analytics

Top Hits Aggregation (5.3) - Visualize the ‘latest’ metric

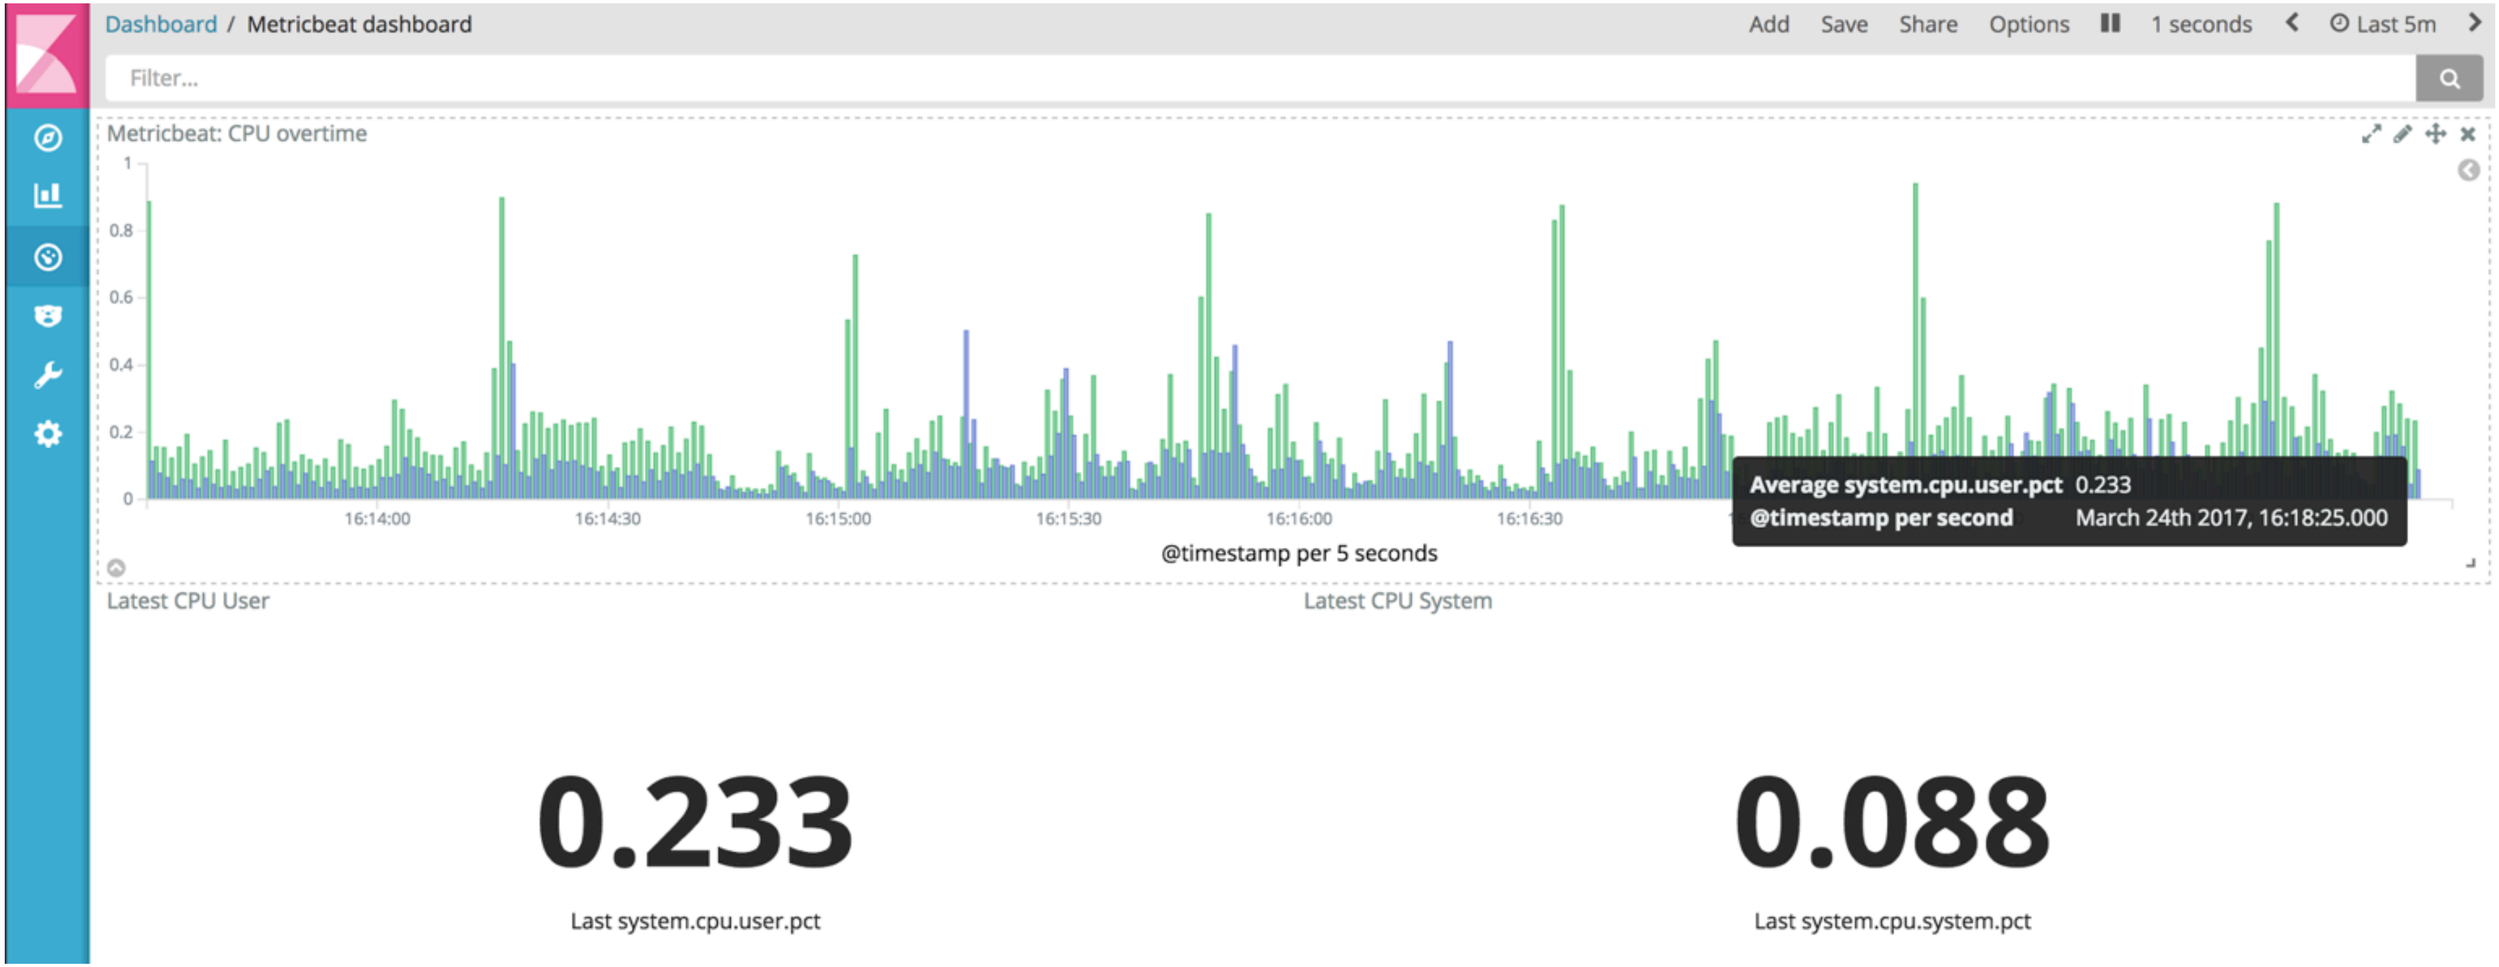

elastic

Profile your Search Queries

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

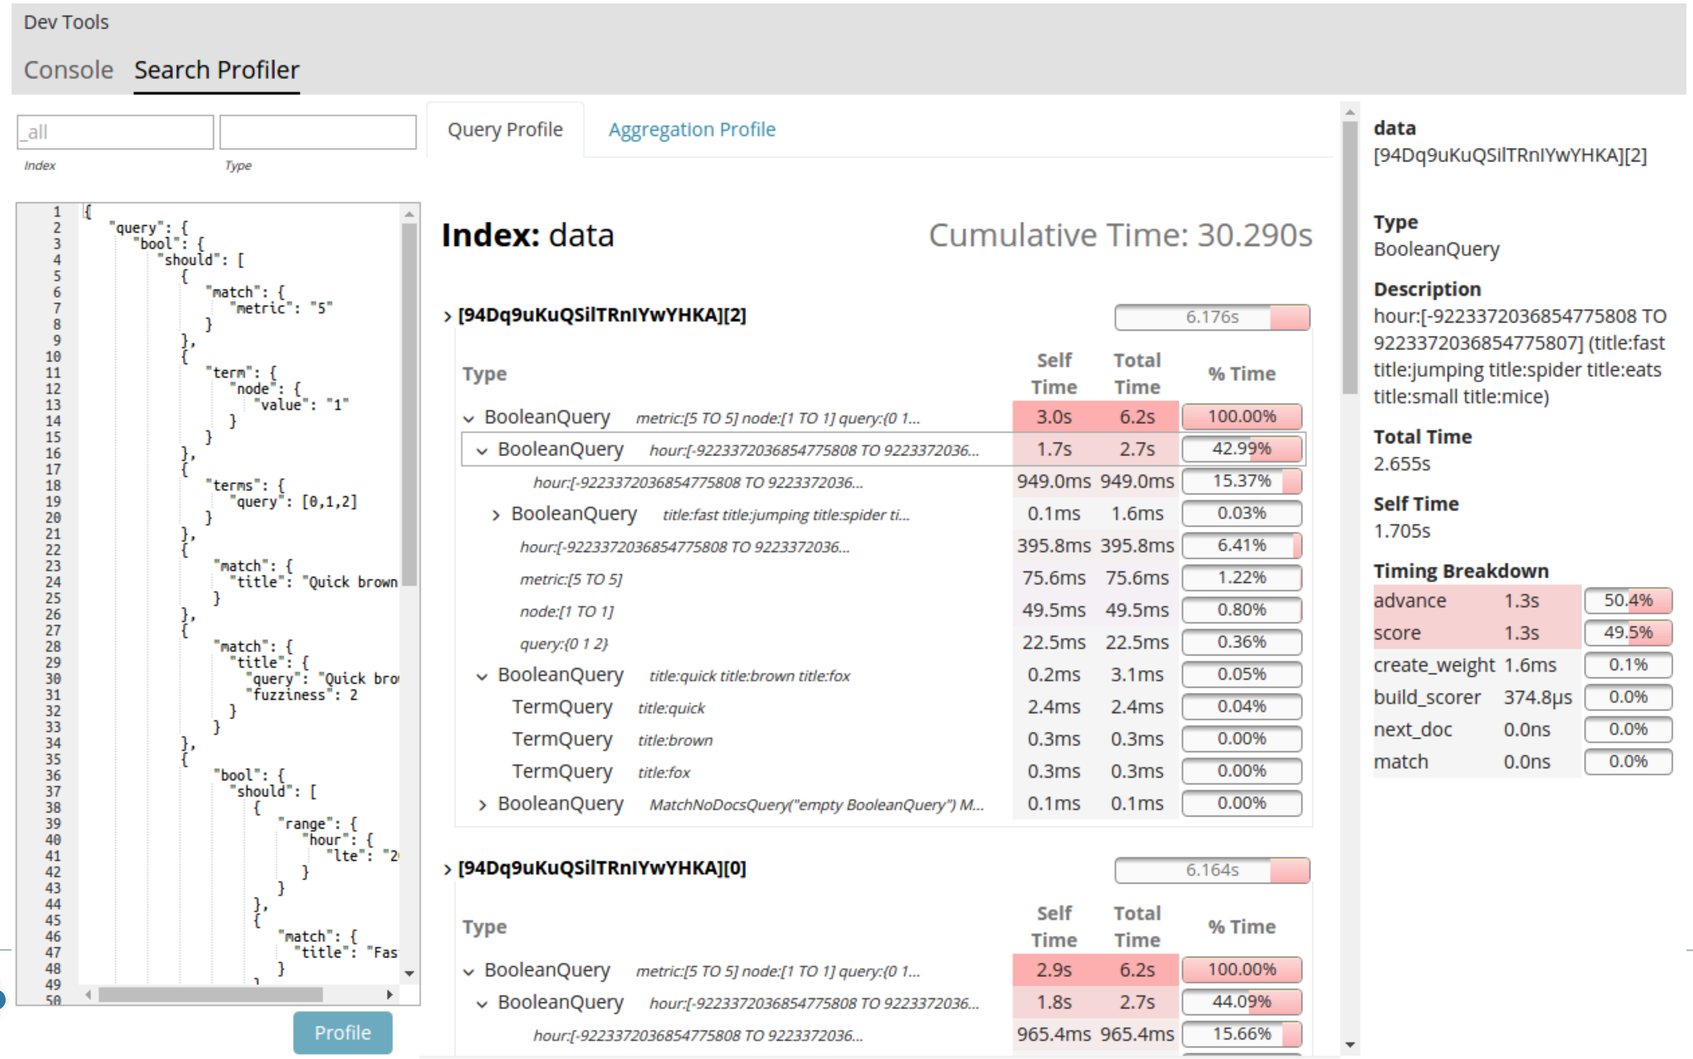

* requires X-Pack (Basic)

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

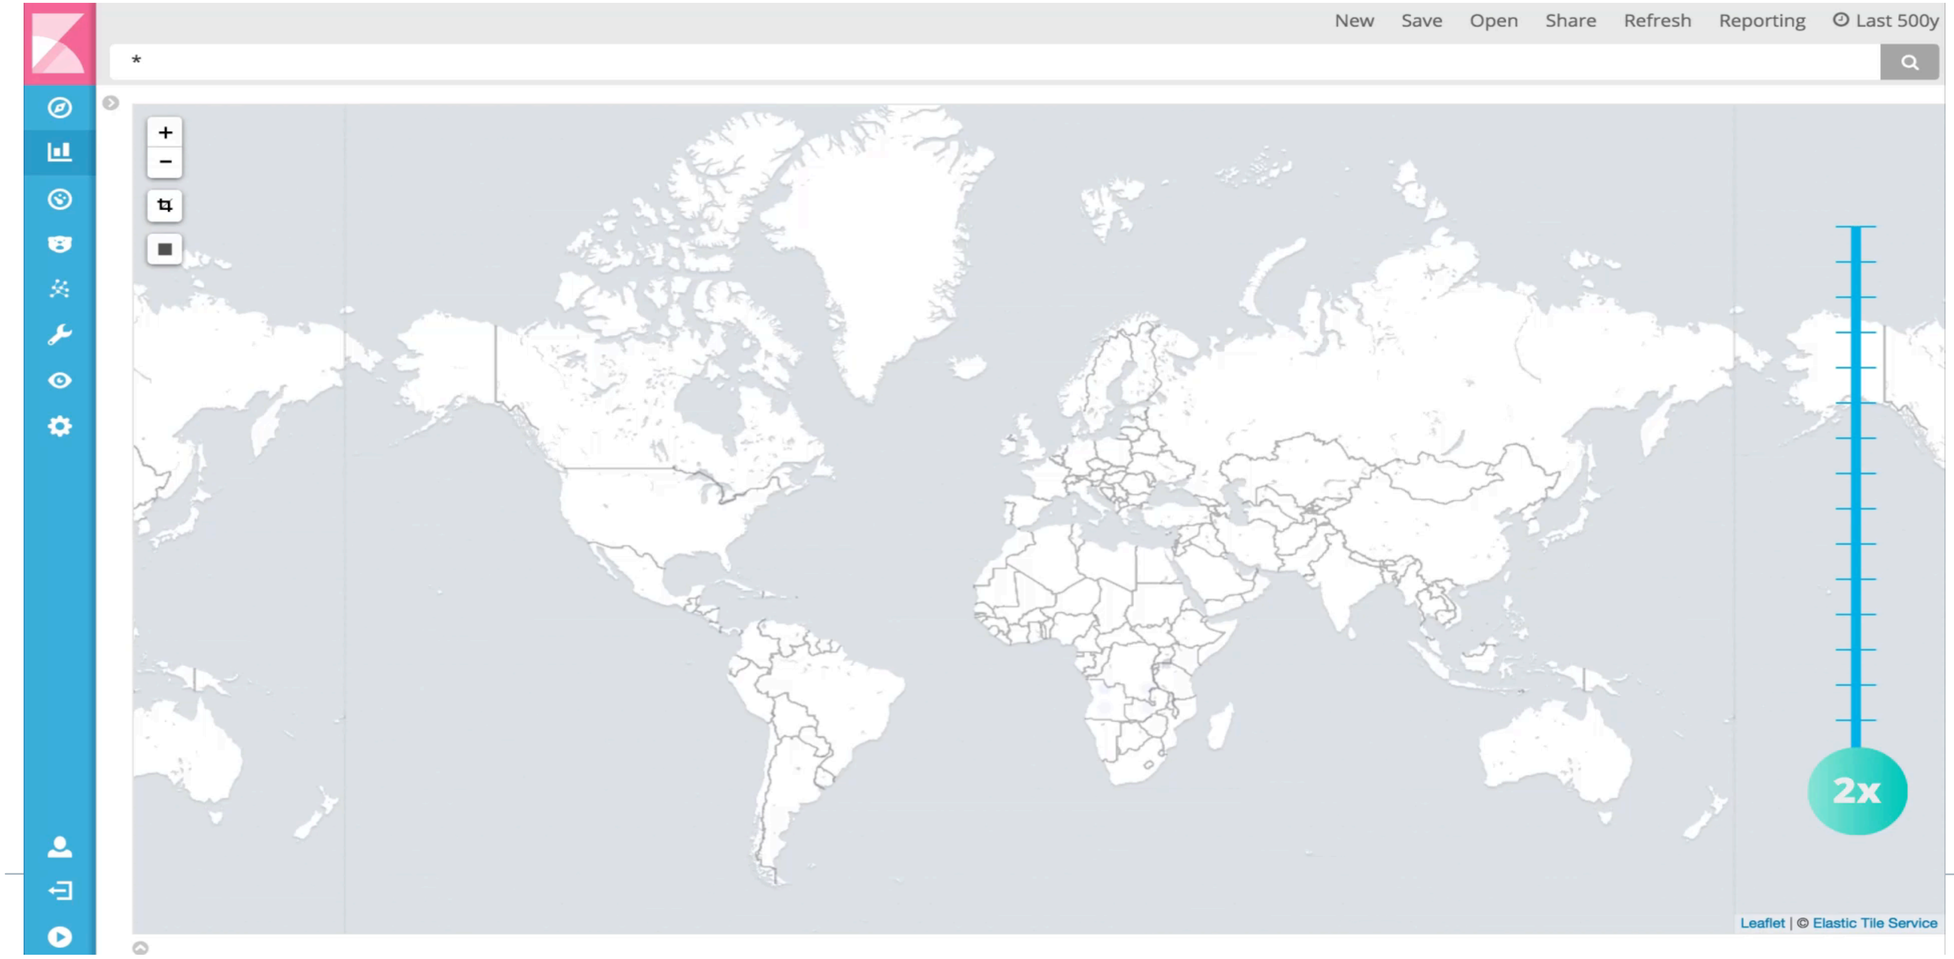

X-Pack (Basic) required for zoom > 10

Added support for Tribe

Tribe Support

As simple as:

`elasticsearch.tribe.url`

No jumping through hoops.

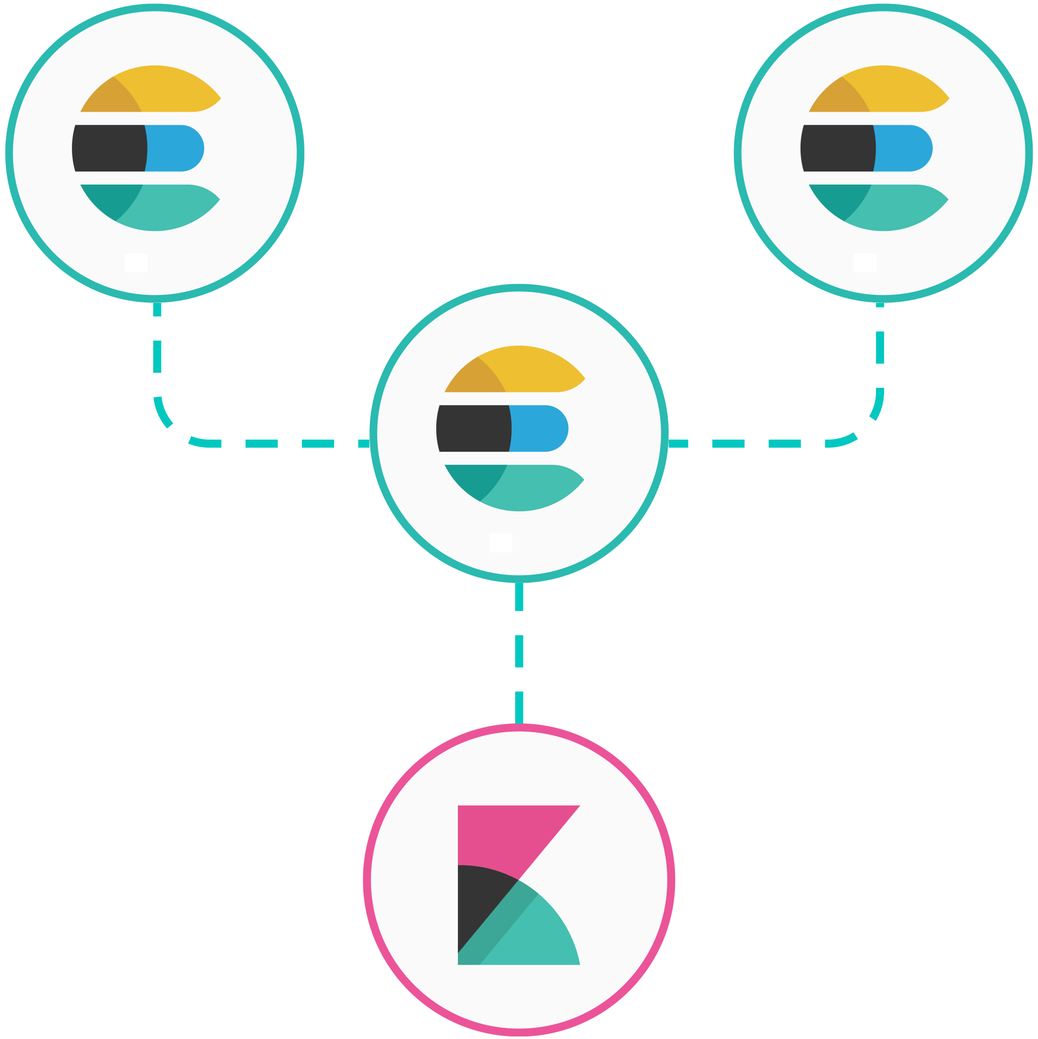

View/Edit mode for Dashboards

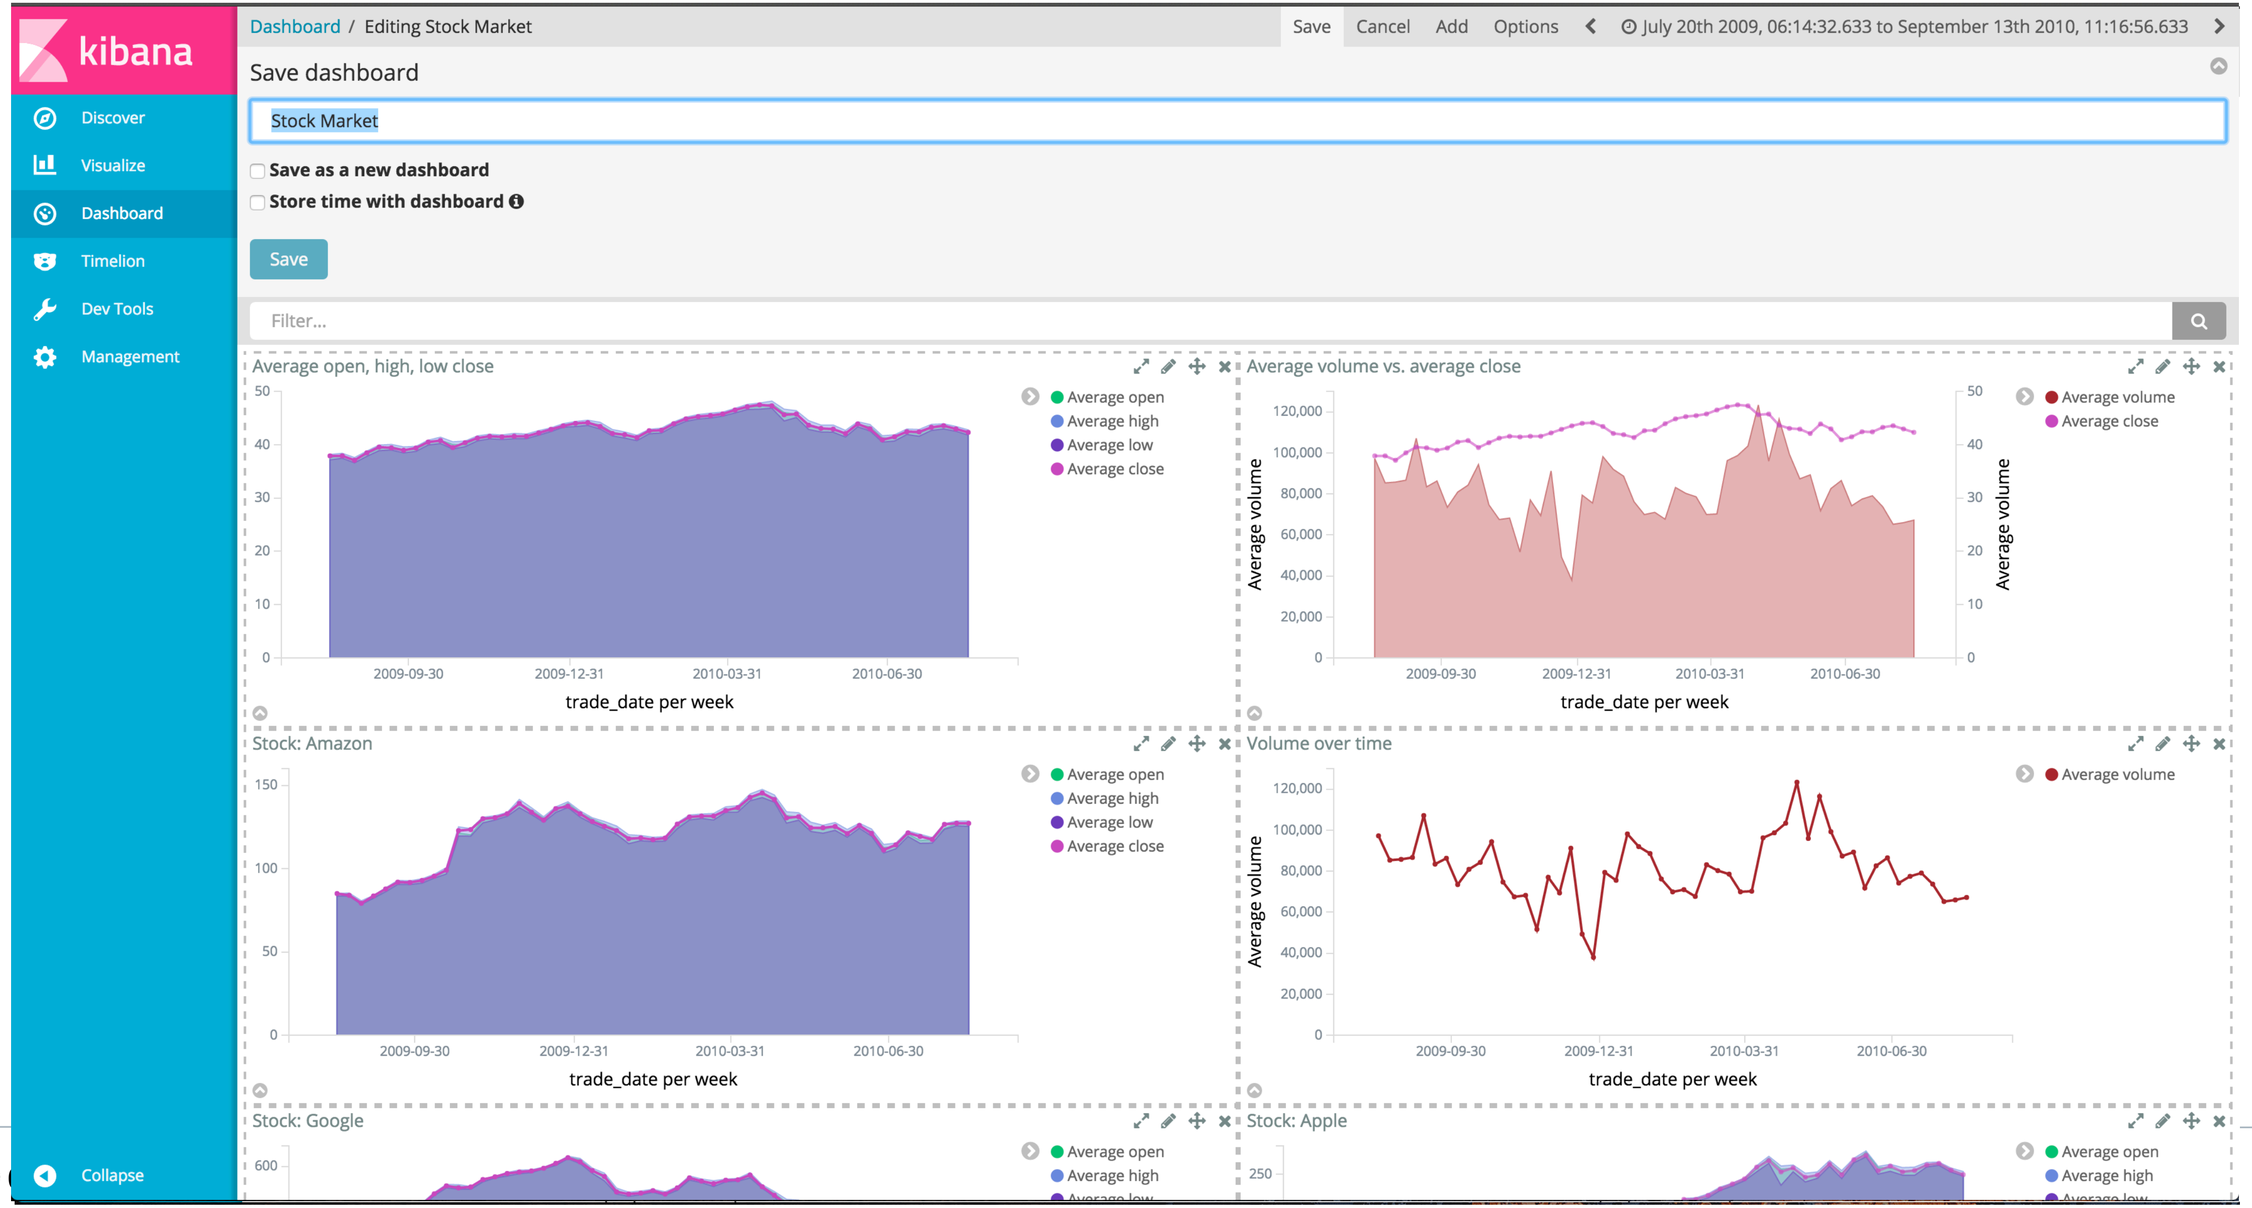

Time Se

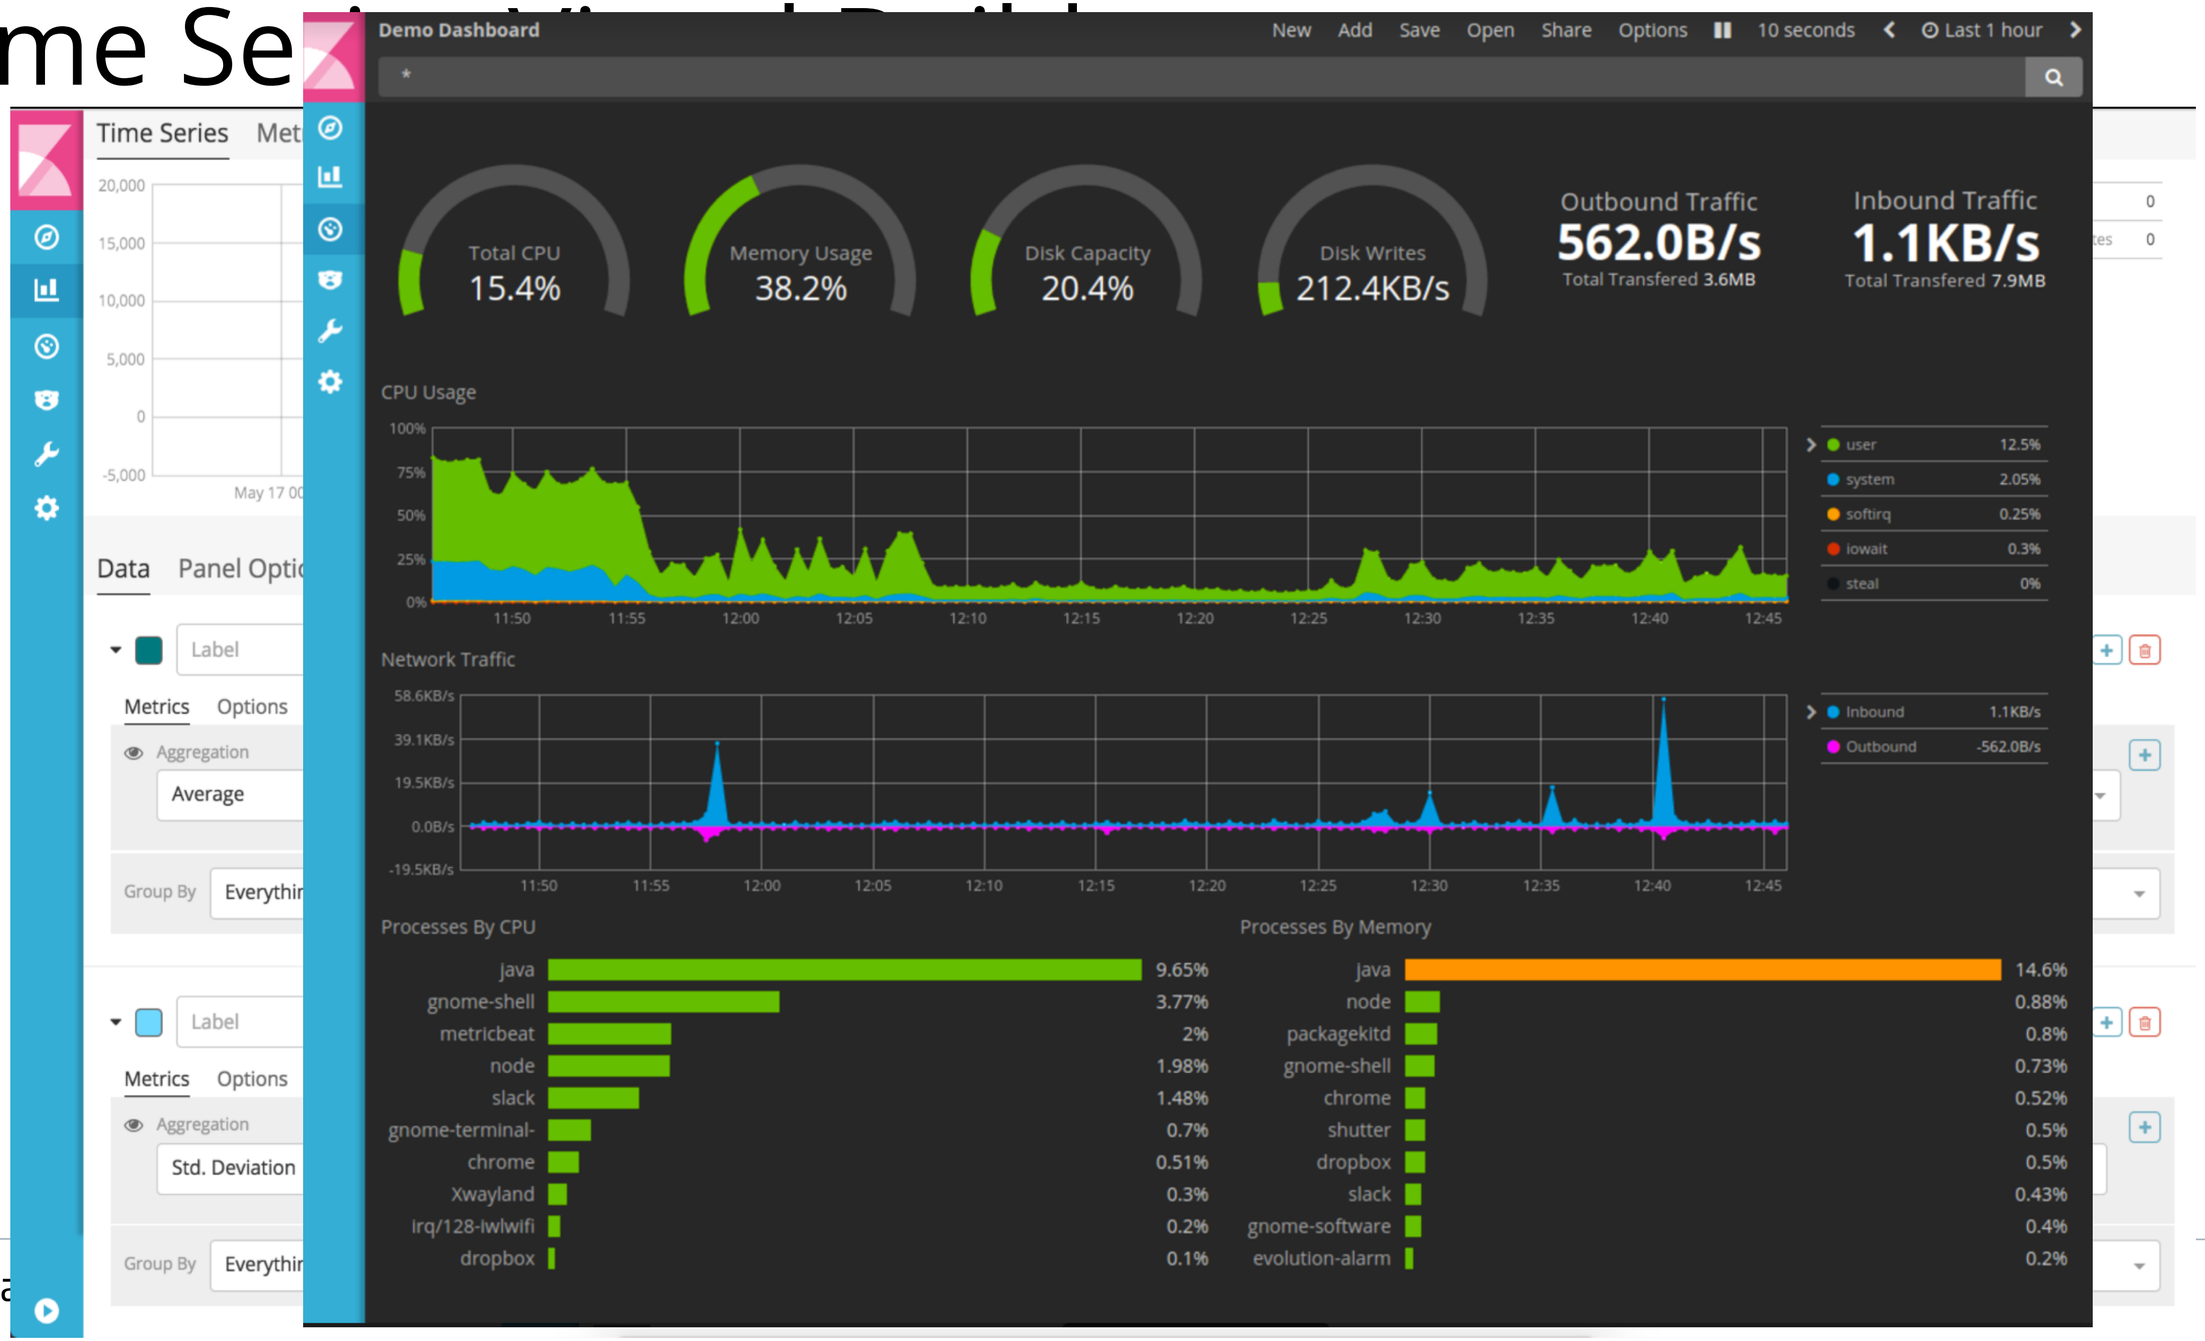
Time Series
Met
20,000
15,000
10,000
5,000
0
-5,000
May 17 00
Data
Panel Optic
Label
Metrics
Options
Aggregation
Average
Group By
Everythir
Label
Metrics
Options
Aggregation
Std. Deviation
Group By
Everythir
Demo Dashboard
New
Add
Save
Open
Share
Options
10 seconds
Last 1 hour
*
Total CPU
15.4%
Memory Usage
38.2%
Disk Capacity
20.4%
Disk Writes
212.4KB/s
Outbound Traffic
562.0B/s
Total Transfered 3.6MB
Inbound Traffic
1.1KB/s
Total Transfered 7.9MB
CPU Usage
100%
75%
50%
25%
0%
11:50
11:55
12:00
12:05
12:10
12:15
12:20
12:25
12:30
12:35
12:40
12:45
user 12.5%
system 2.05%
softirq 0.25%
iowait 0.3%
steal 0%
Network Traffic
58.6KB/s
39.1KB/s
19.5KB/s
0.0B/s
-19.5KB/s
Inbound 1.1KB/s
Outbound -562.0B/s
Processes By CPU
java 9.65%
gnome-shell 3.77%
metricbeat 2%
node 1.98%
slack 1.48%
gnome-terminal- 0.7%
chrome 0.51%
Xwayland 0.3%
irq/128-iwlwifi 0.2%
dropbox 0.1%
Processes By Memory
java 14.6%
node 0.88%
packagekitd 0.8%
gnome-shell 0.73%
chrome 0.52%
shutter 0.5%
dropbox 0.5%
slack 0.43%
gnome-software 0.4%
evolution-alarm 0.2%

Event Context

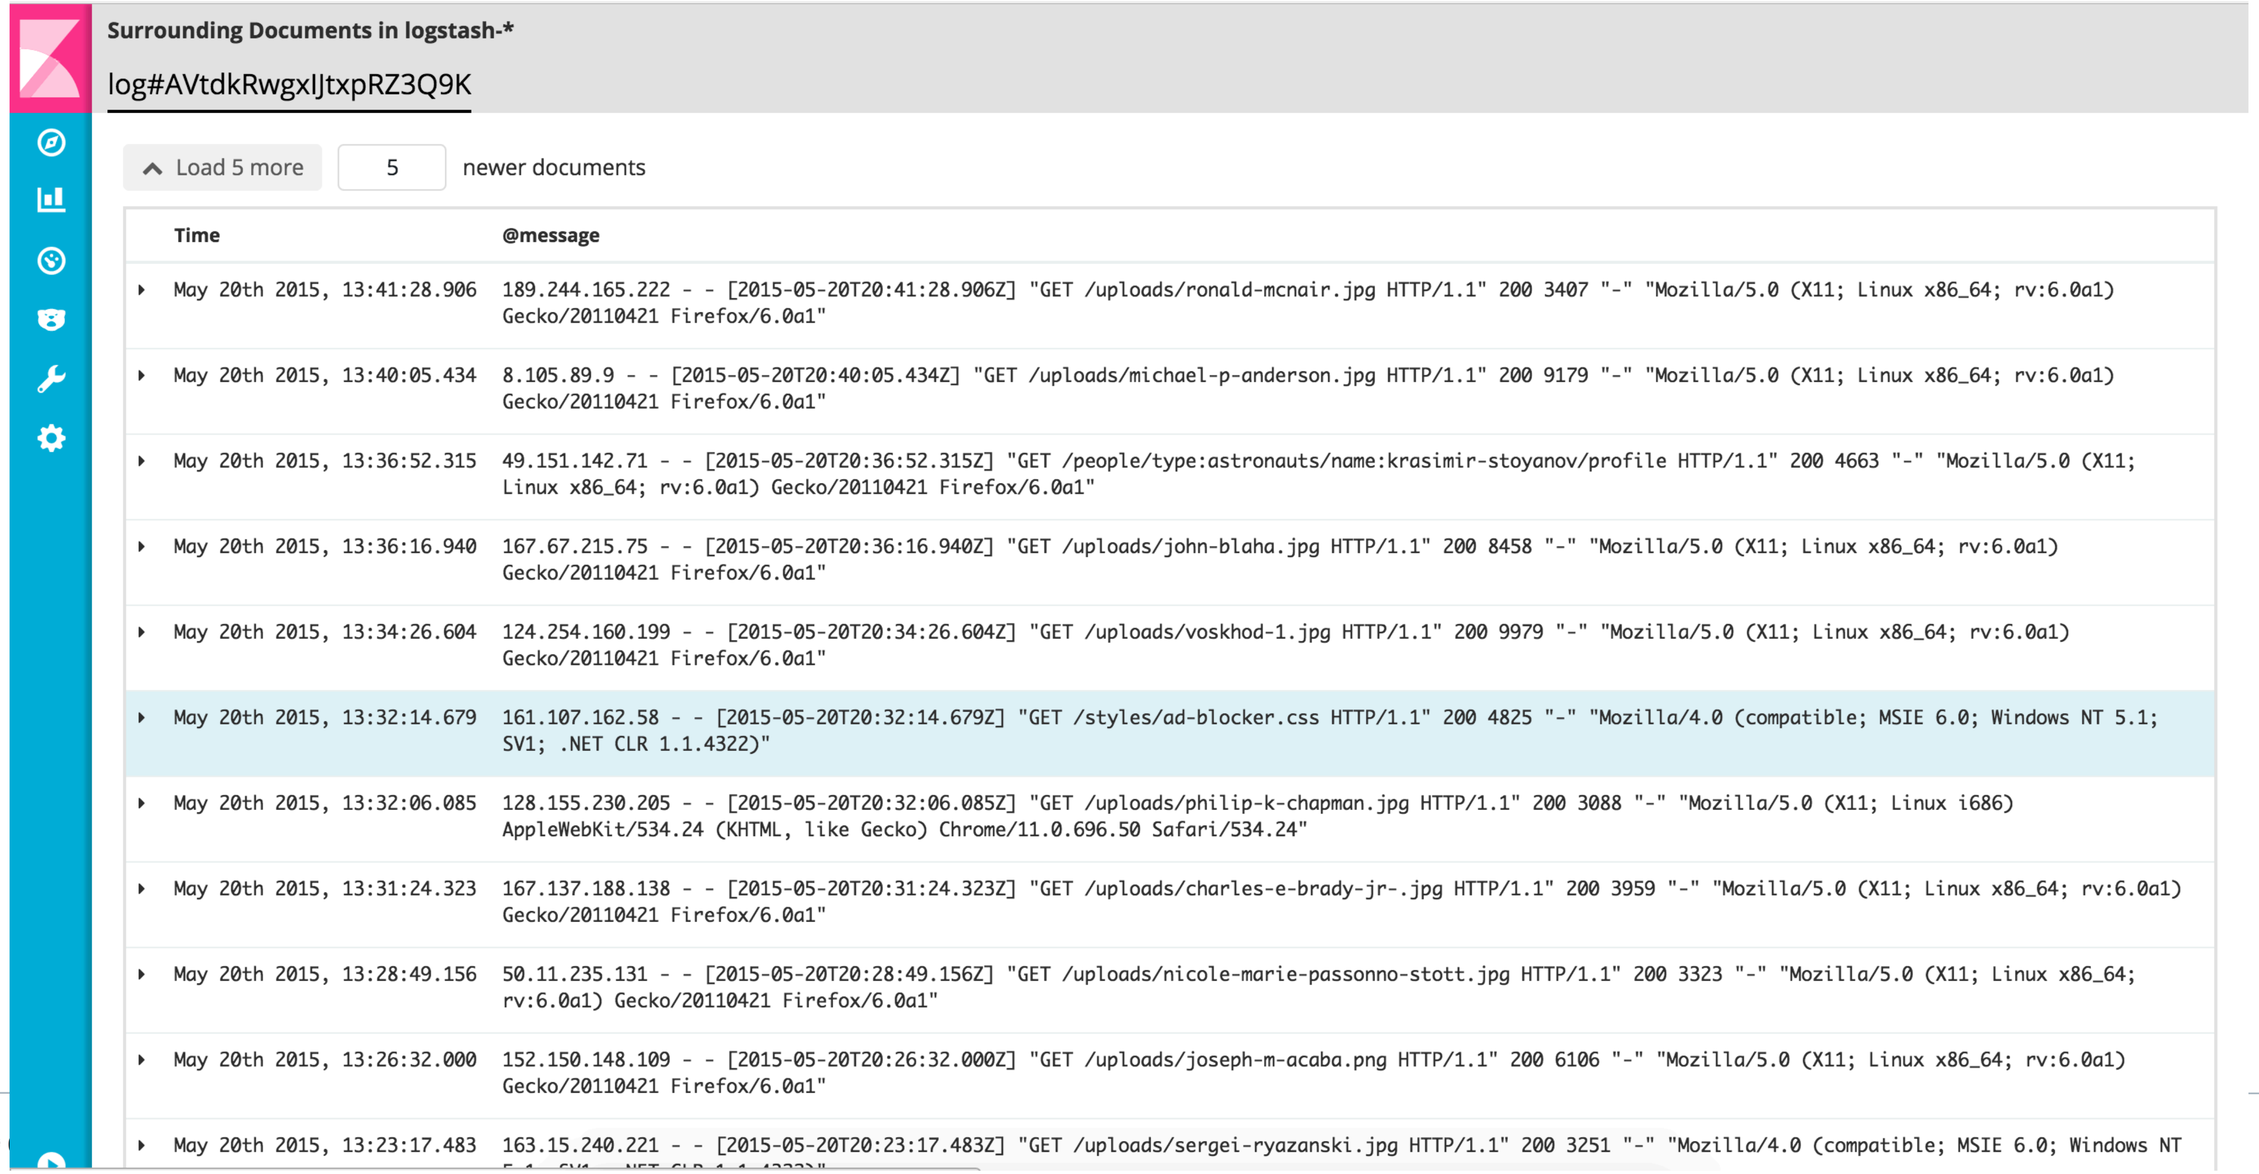

Grok Debugger

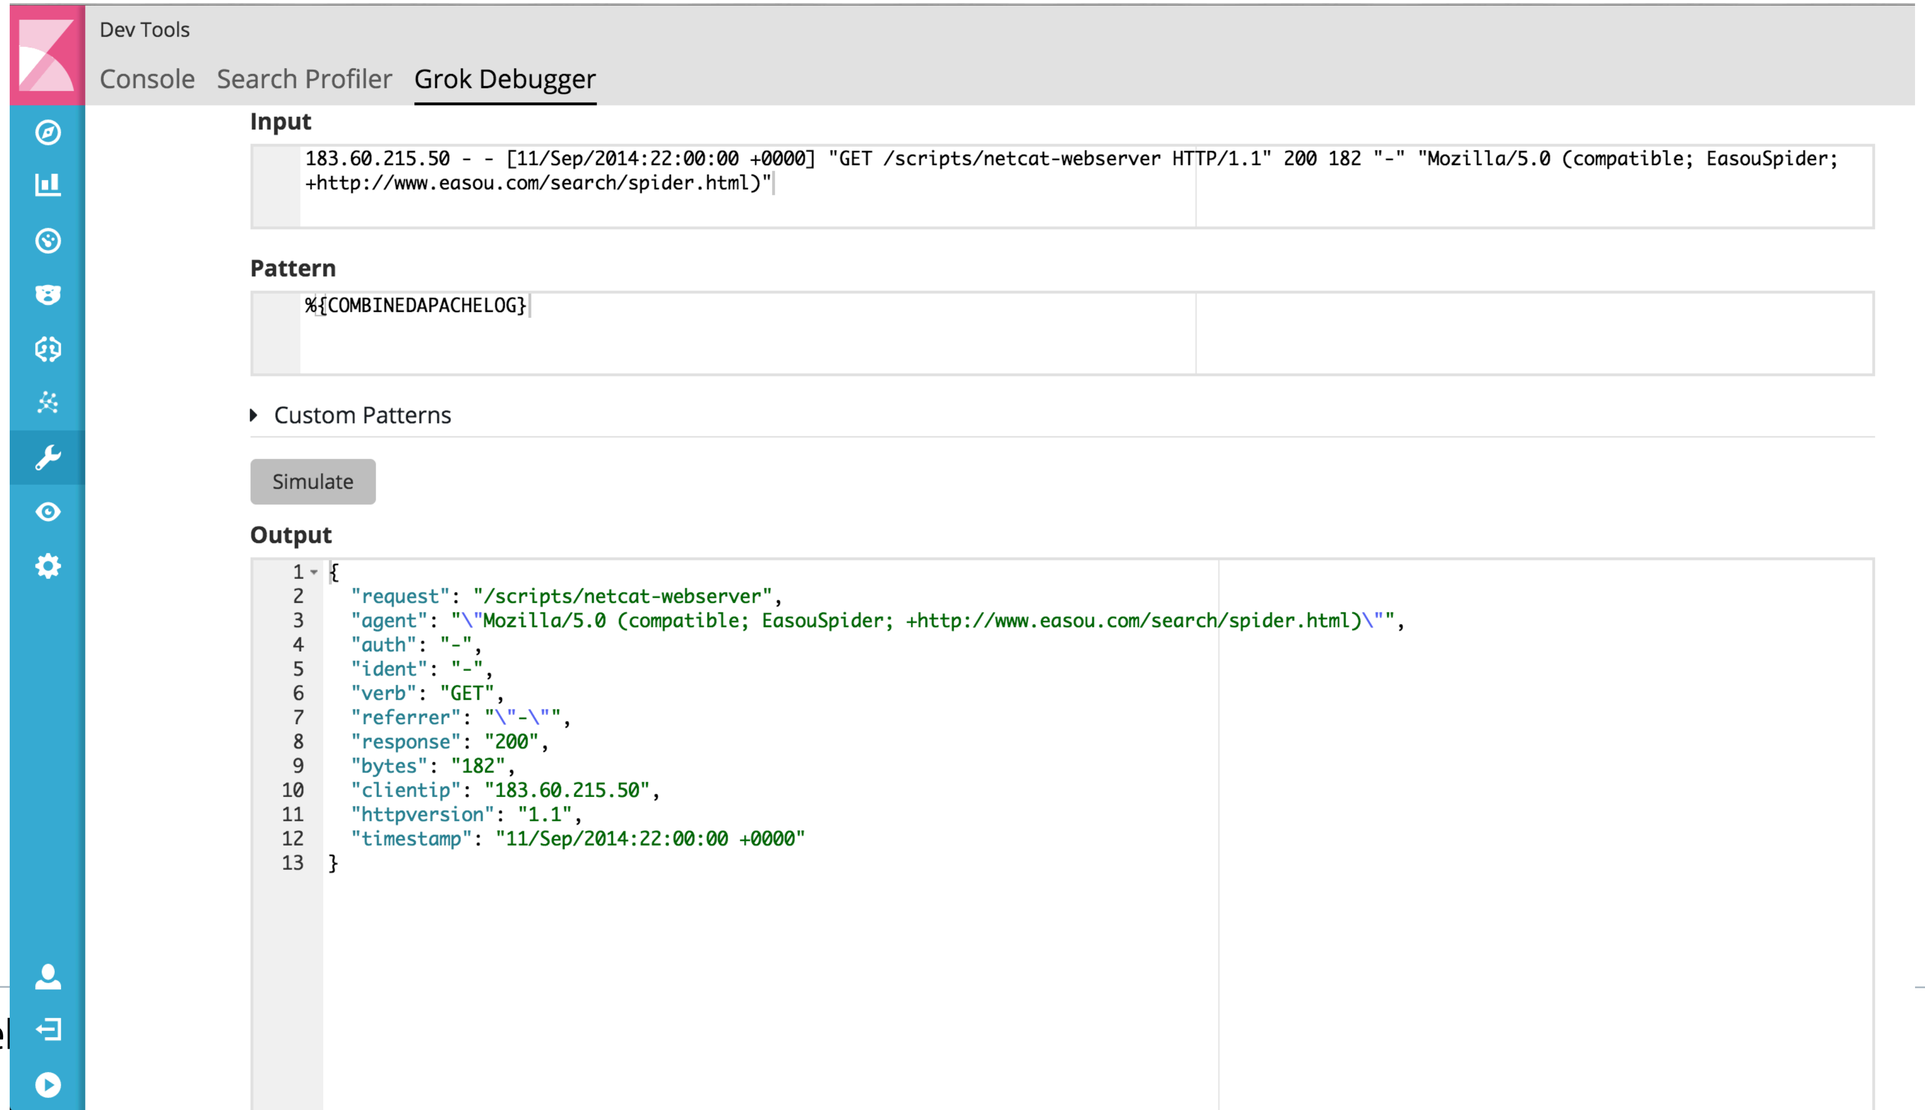

User Experience Improvement

- 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

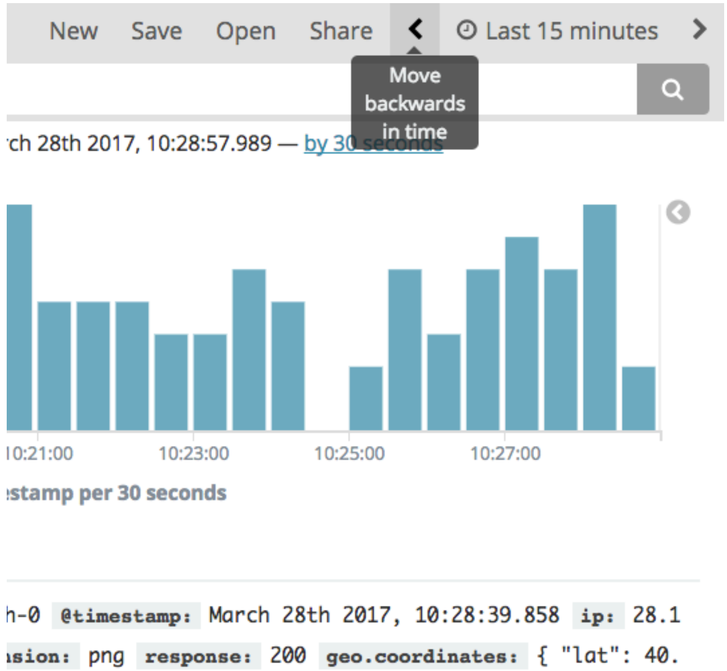



Beats since 5.0

New Beat: Heartbeat (beta in 5.2)

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



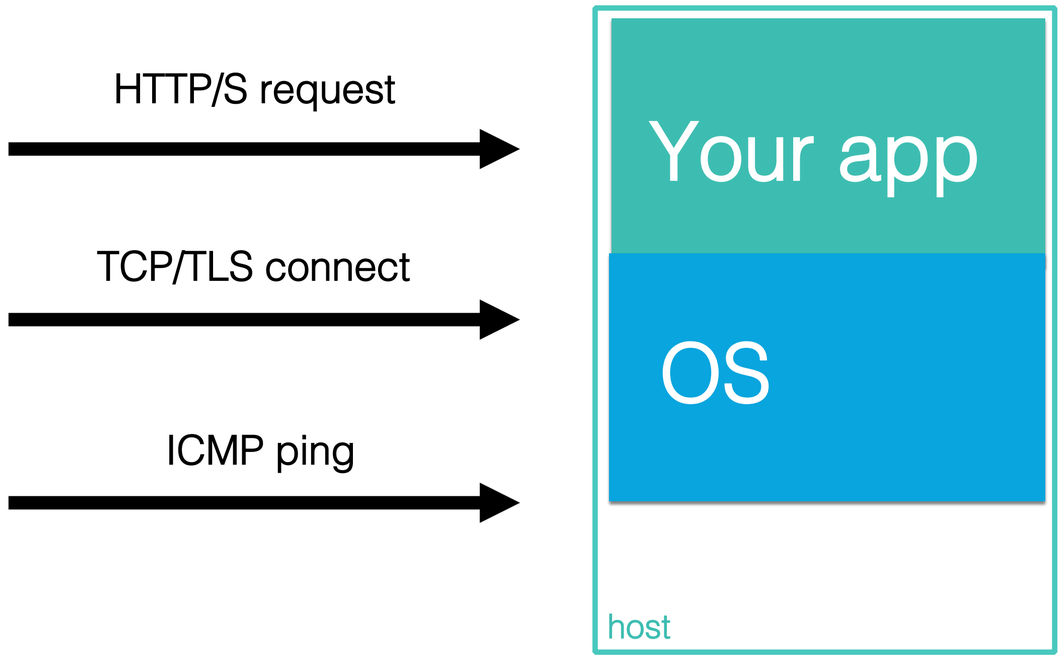

Filebeat Modules (5.3)

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.

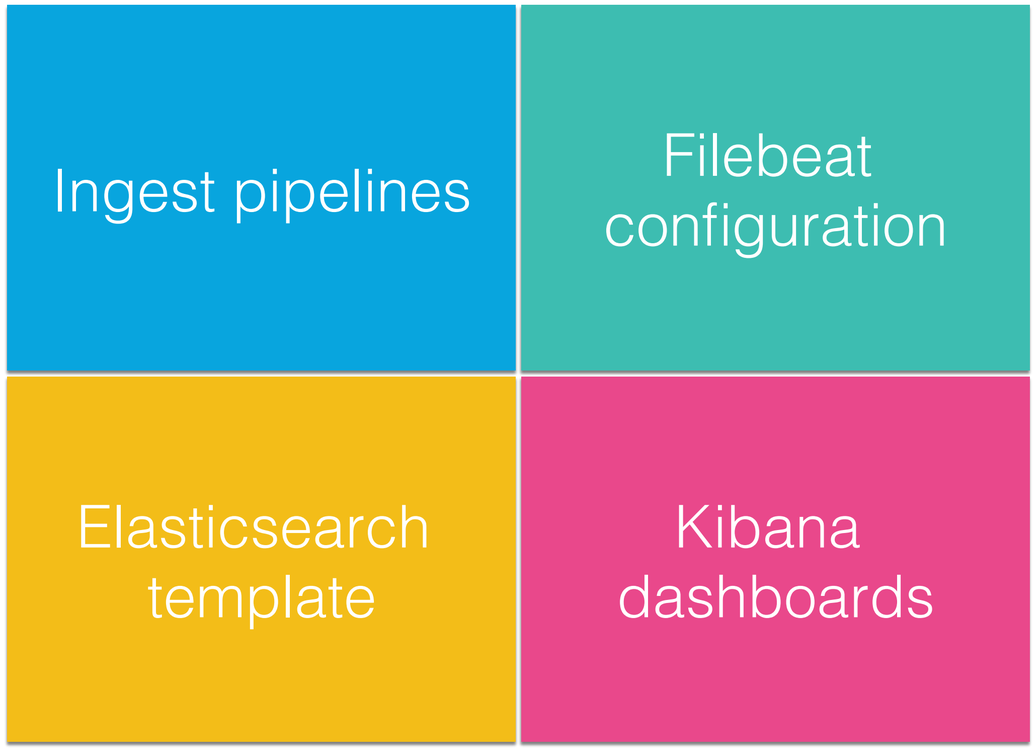

New Modules in Metricbeat

17 modules and growing

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

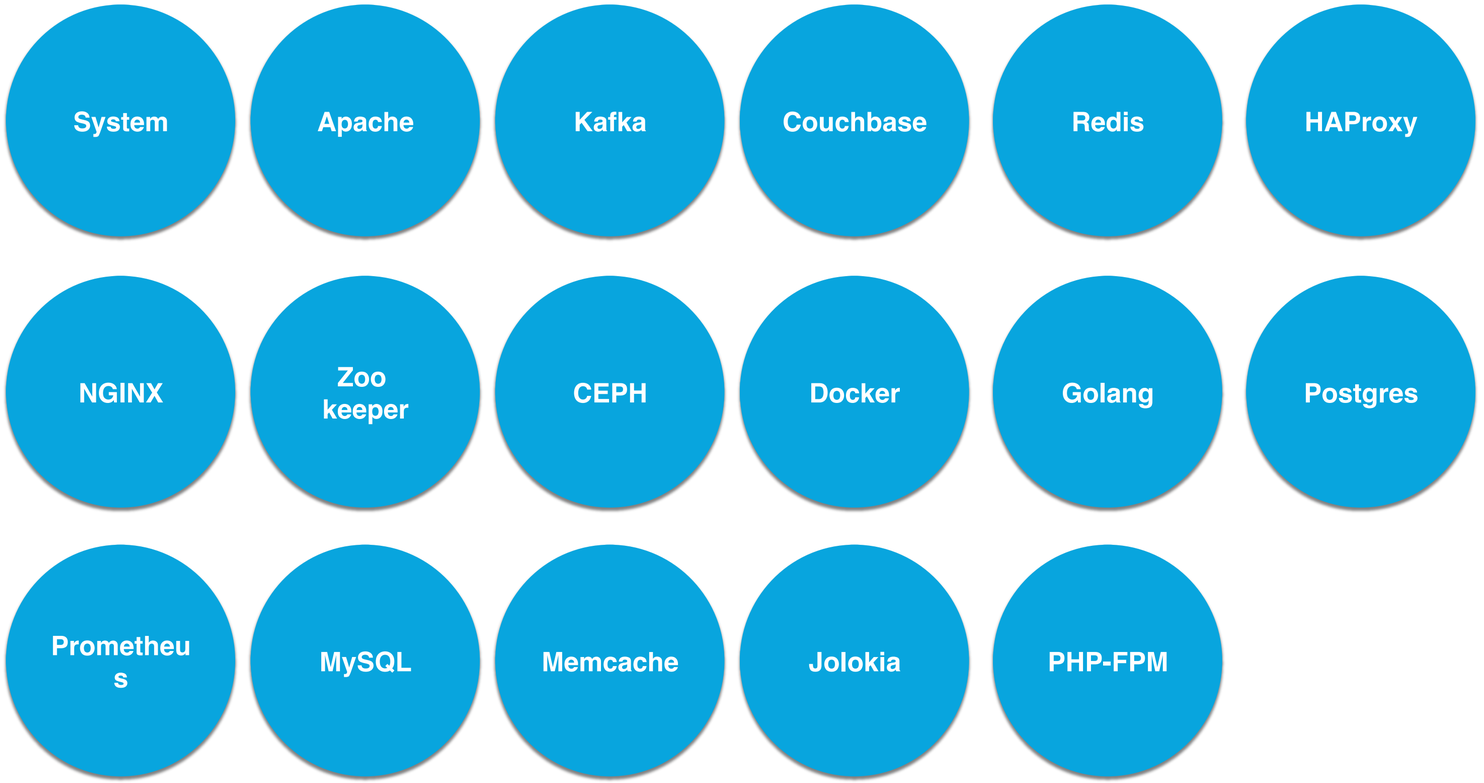

Track Network Connection

Metricbeat system modules adds TCP Socket monitoring (5.2)

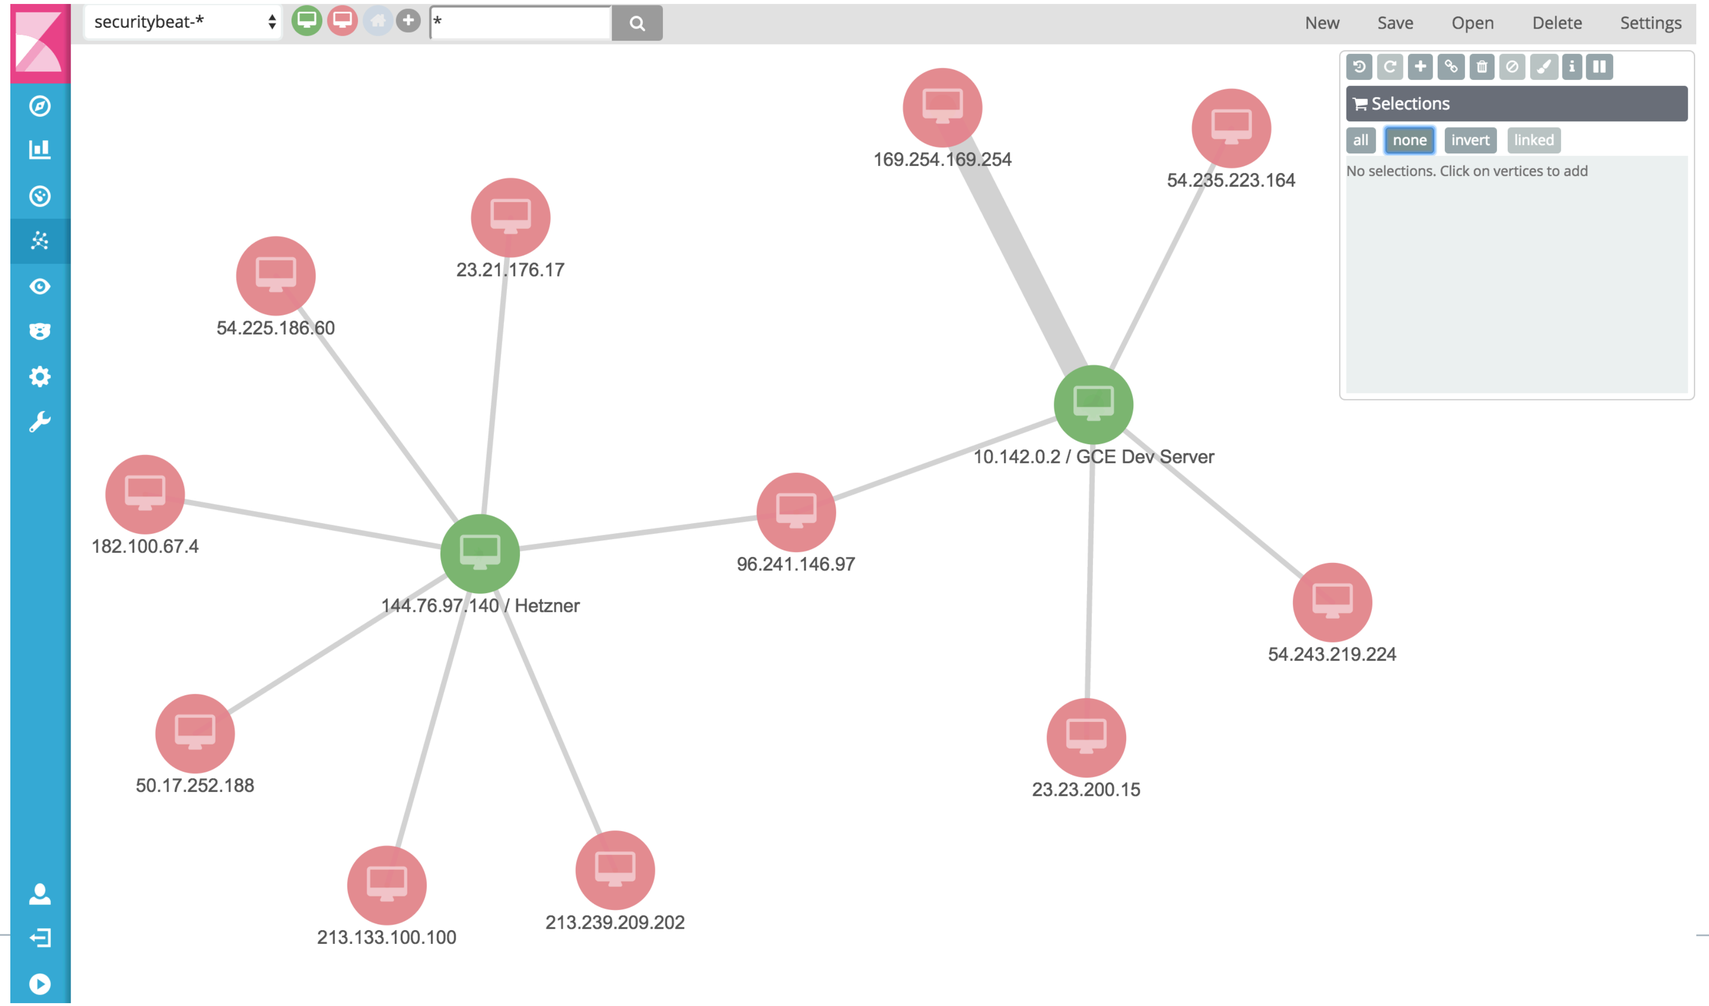

elastic

Monitor applications running in cloud

add_cloud_metadata Beat processor









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

elastic

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

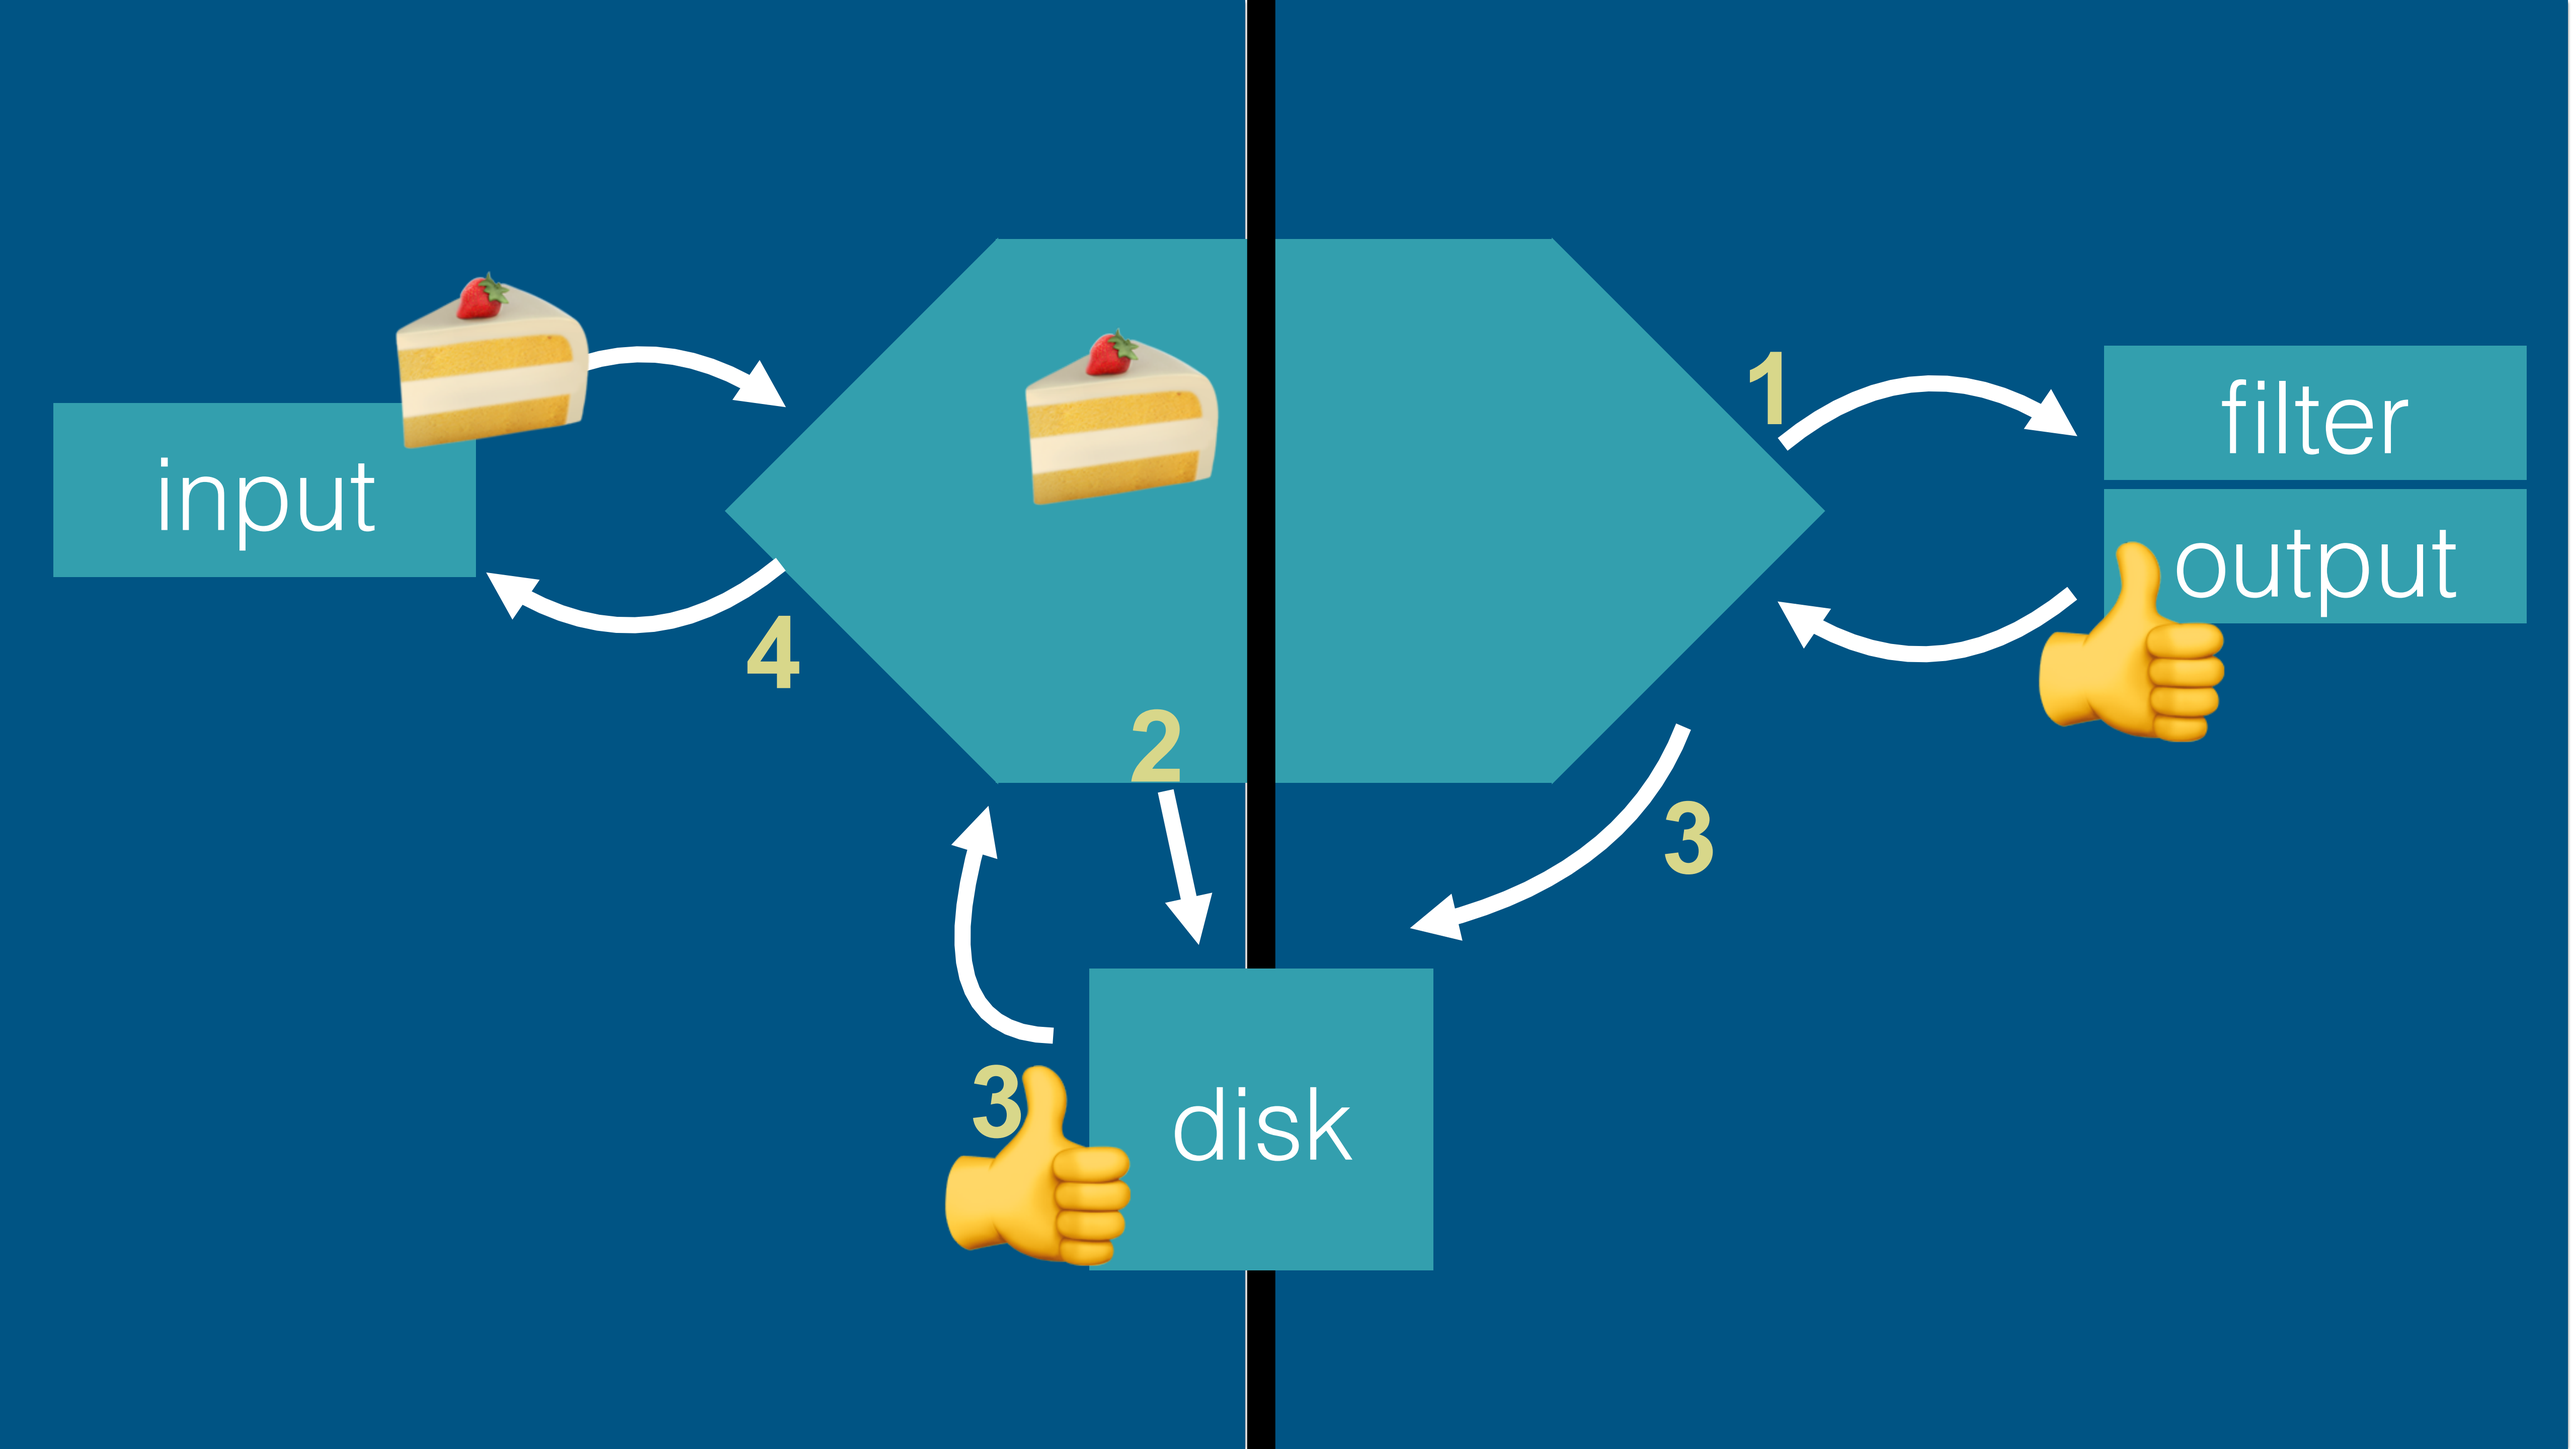
input
4
2
3
1
filter
output
3
disk

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.FRUITS WHERE SKU = :sku"
  parameters => { "sku" => "sku_code"}
}
```

Pipeline Throughput

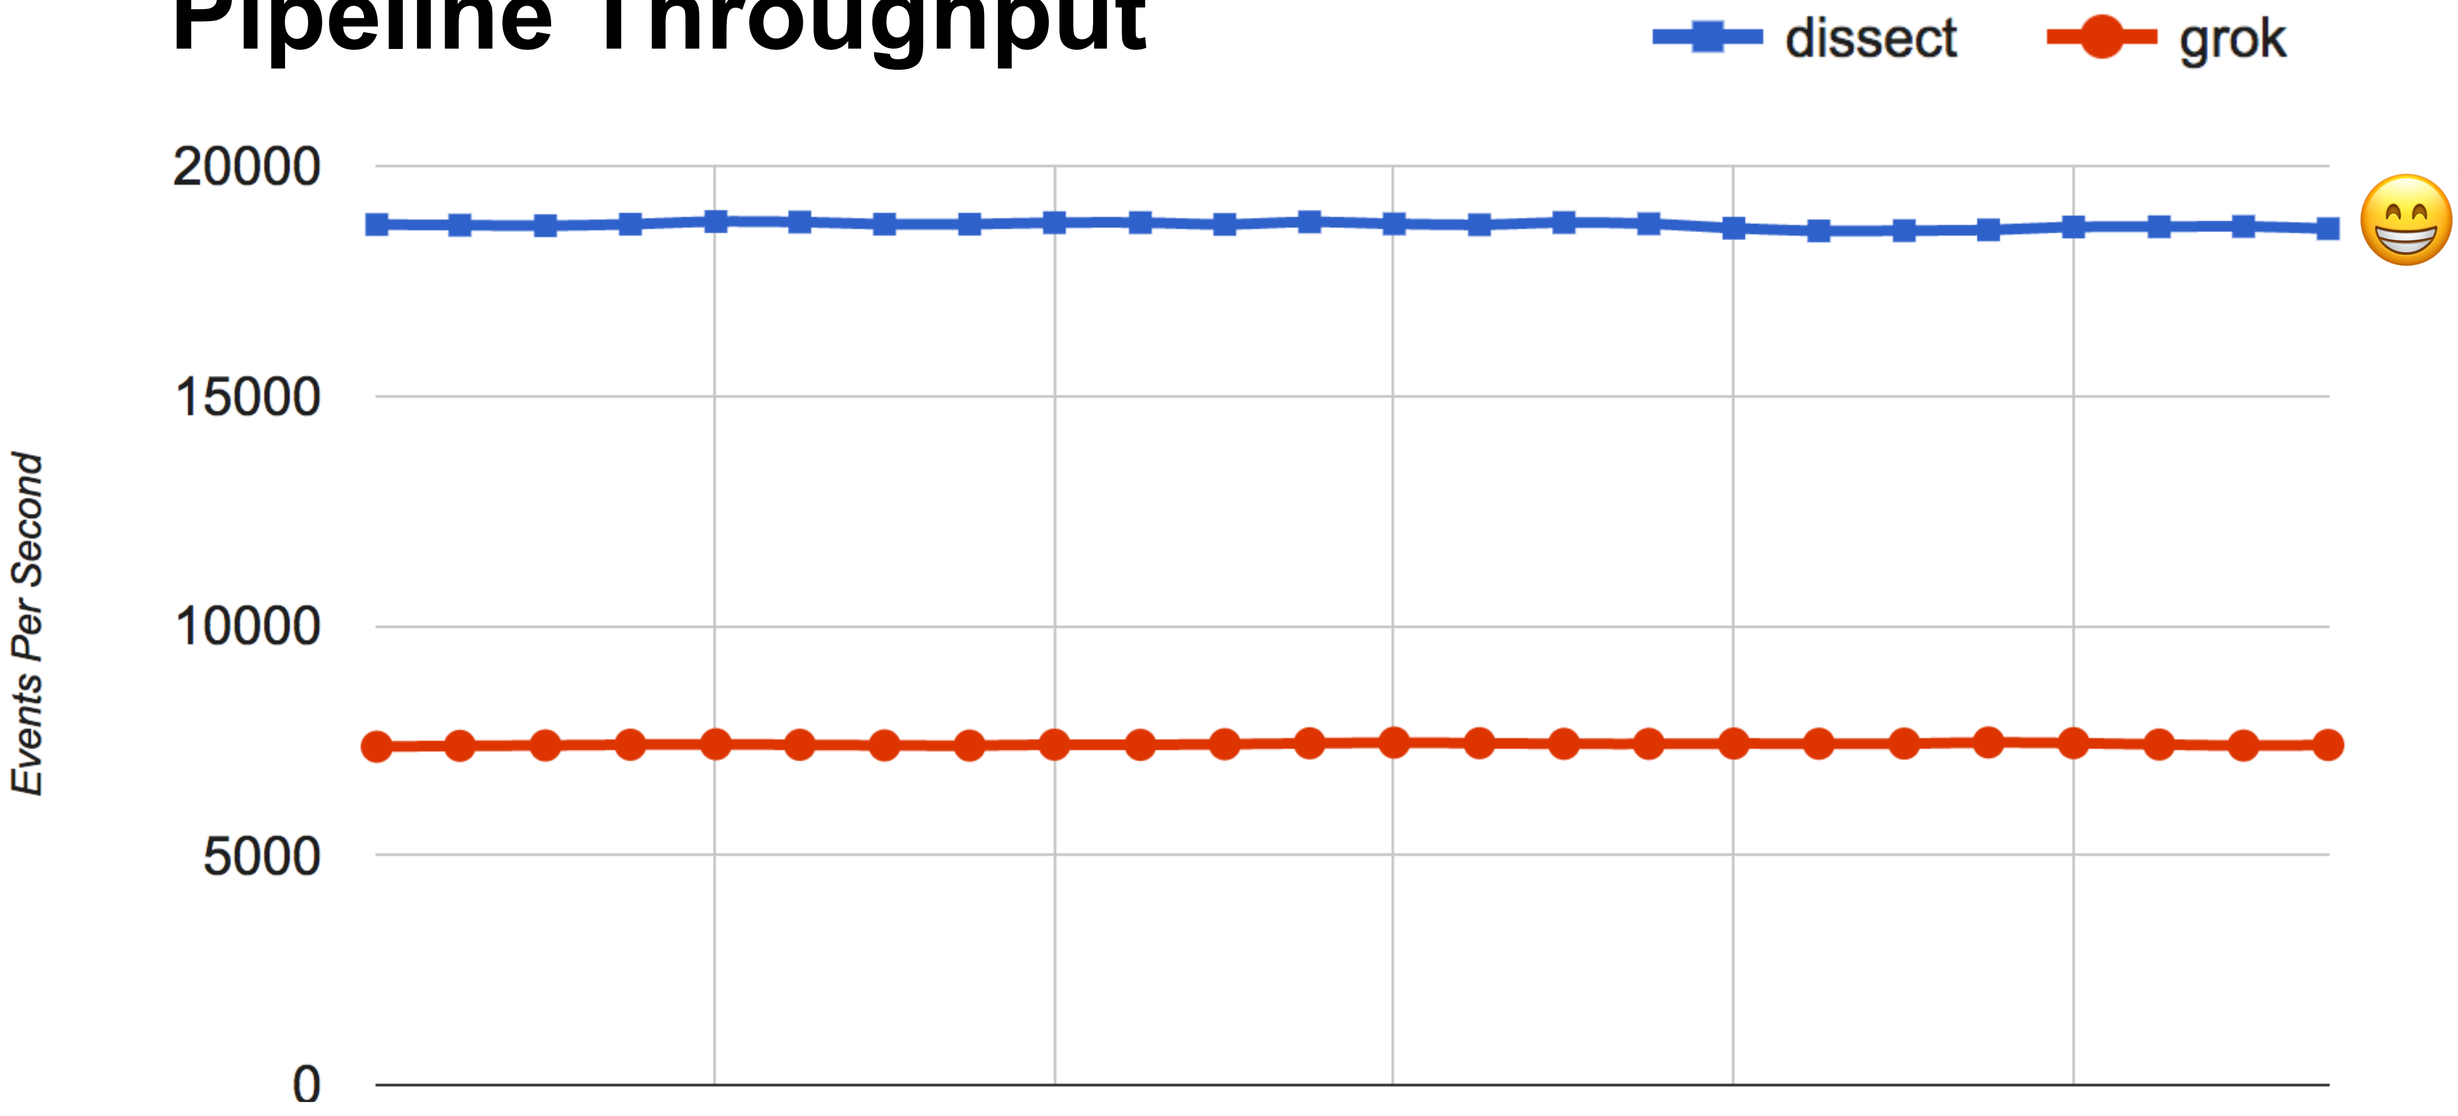

Pipelines: Visual Builder

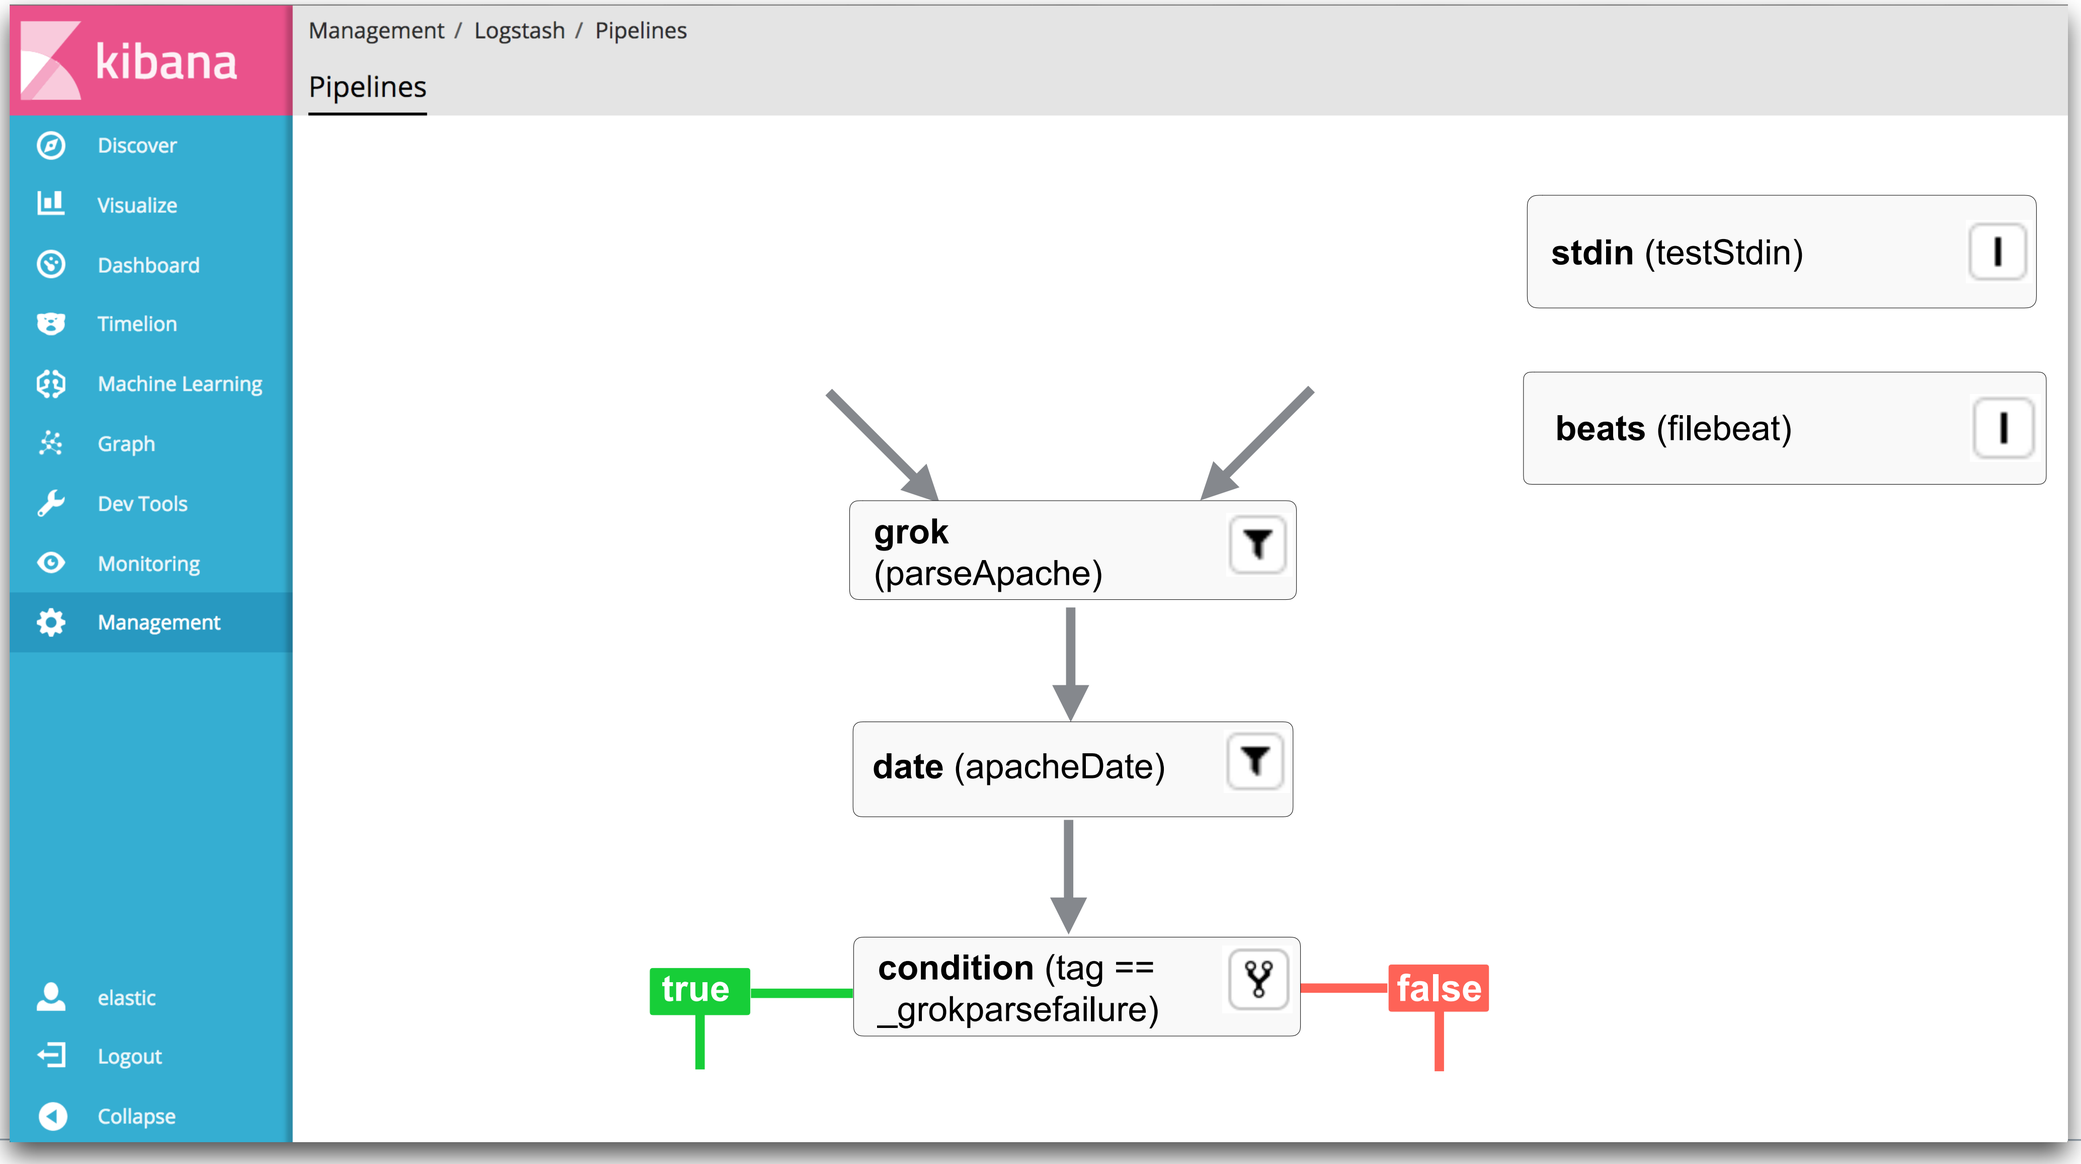



X-Pack since 5.0

Monitoring Enhancements

Logstash Monitoring UI added to X-Pack



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)

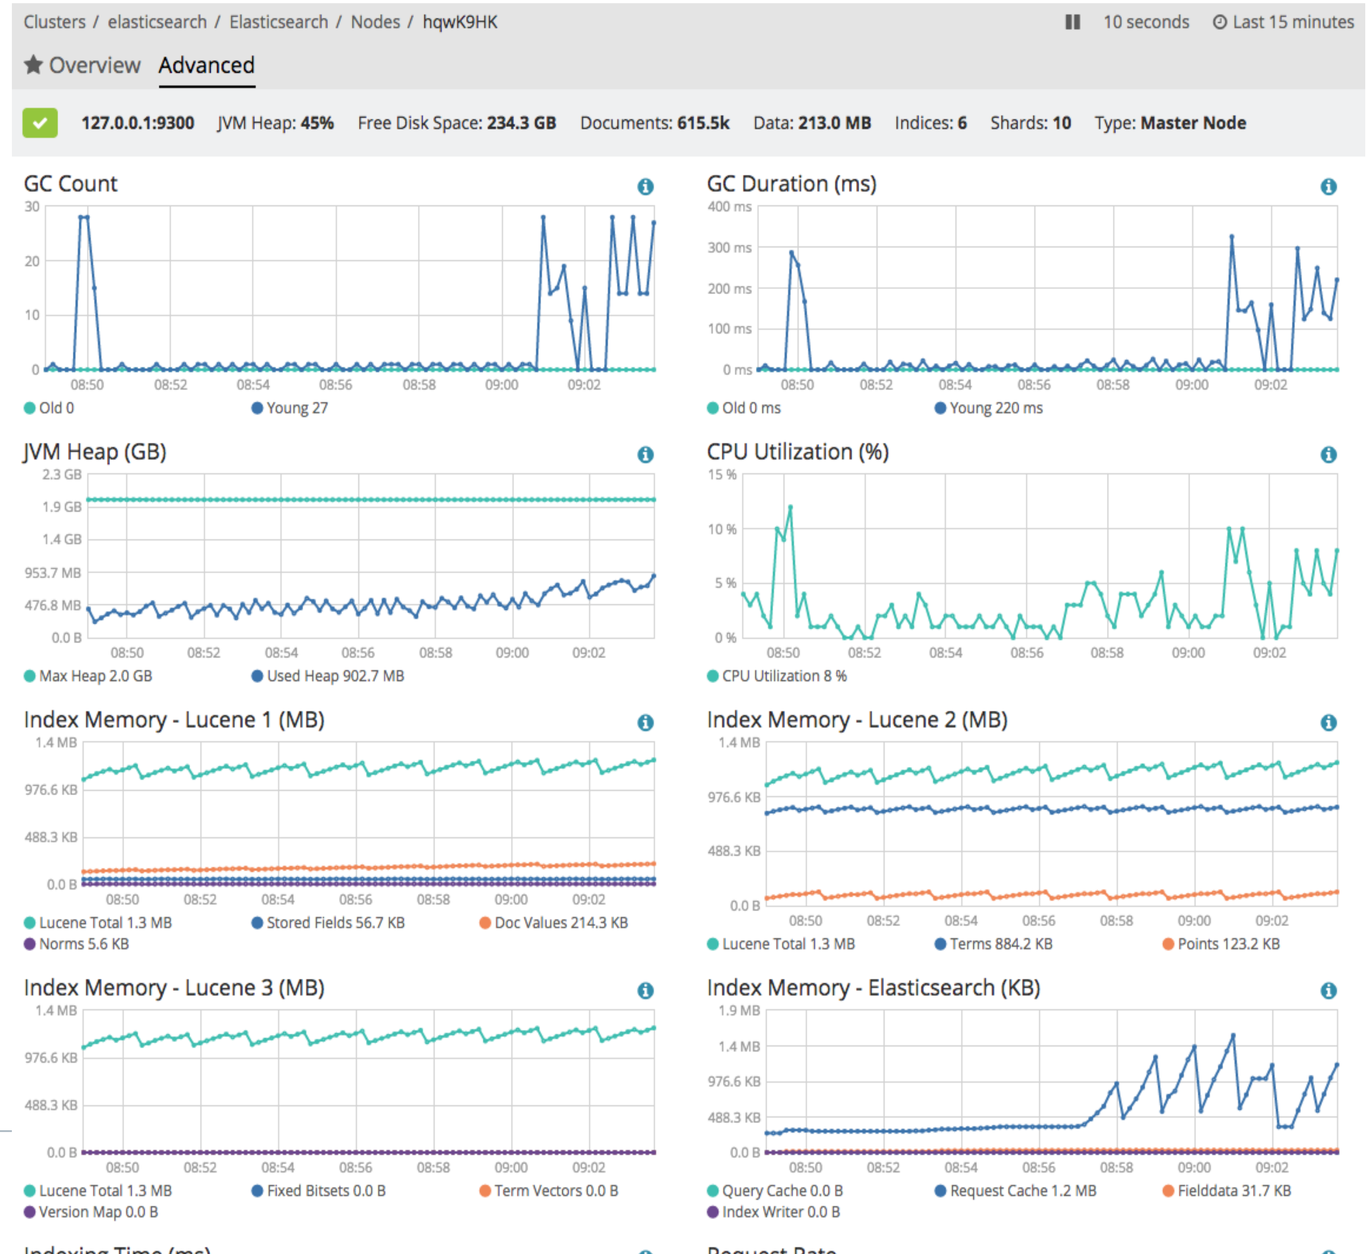

elastic

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

elastic

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

elastic

Machine Learning

Already released in 5.4



Chart interval: 30m Use full it_ops-kpi data

New job from index pattern it_ops-kpi

Aggregation **Field** **Bucket span**

Count 5m

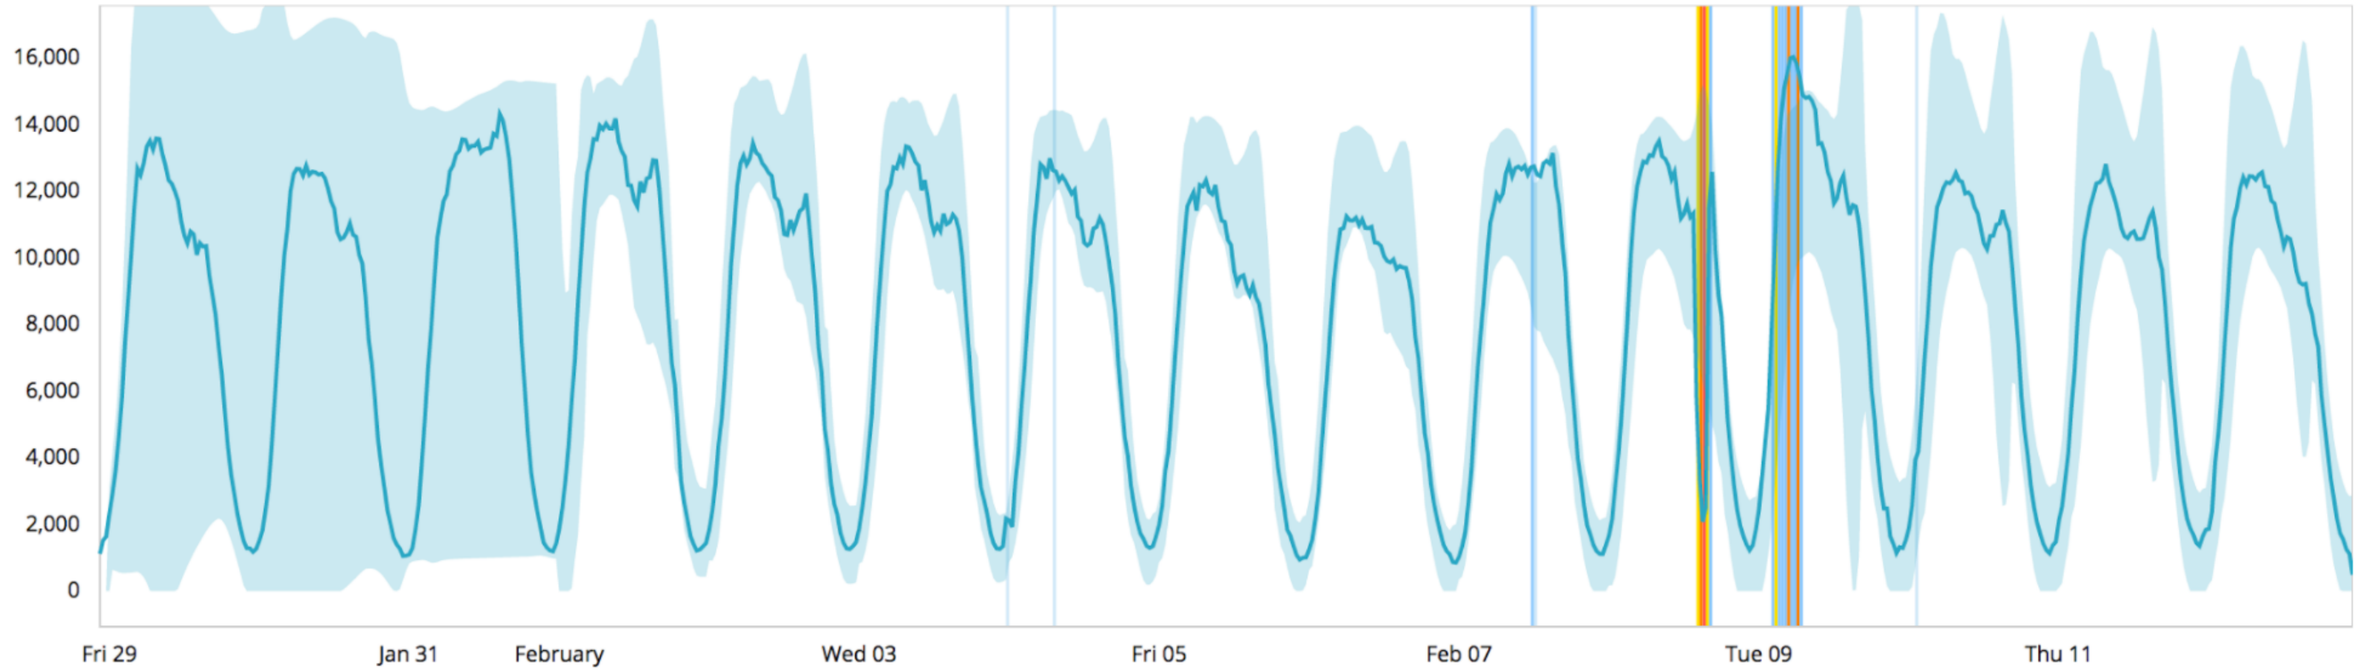

Job it_ops-kpi created

Reset View Results

Job: it_ops-kpi and 2 others

Top Influencers

service

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

MSSQL-0783E4076	94	1237

Anomaly timeline

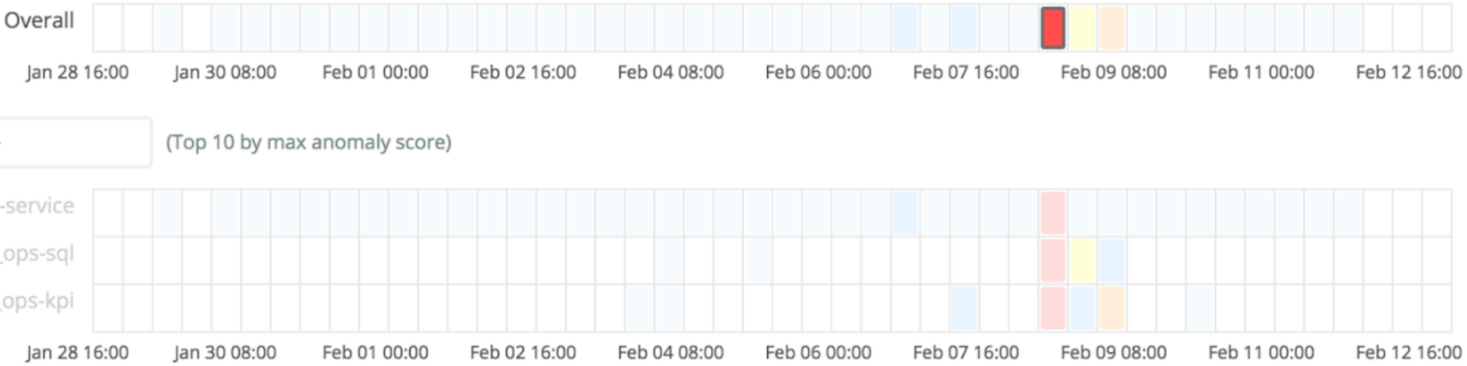

Anomalies

mean responsetime service inventory-us-east-1-34

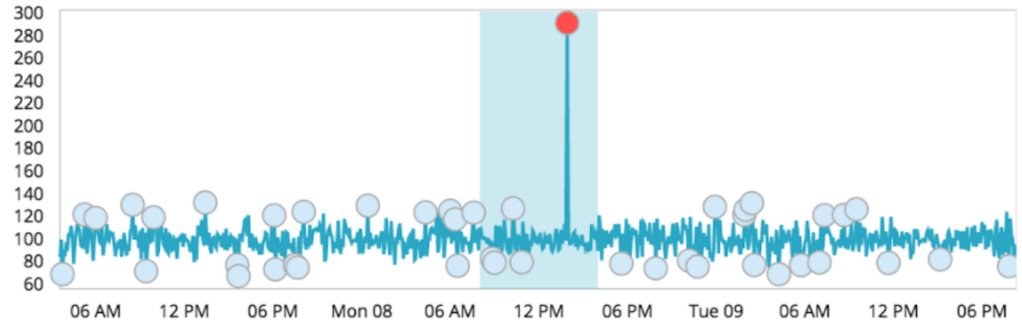

mean SQLServer_Buffer_Manager_Buffer_cache_hit_ratio hostname MSSQL-0783E4076

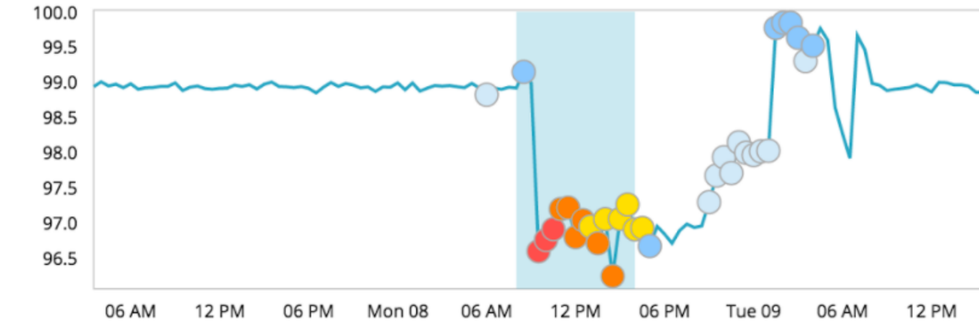

mean SQLServer_General_Statistics_User_Connections hostname MSSQL-0783E4076

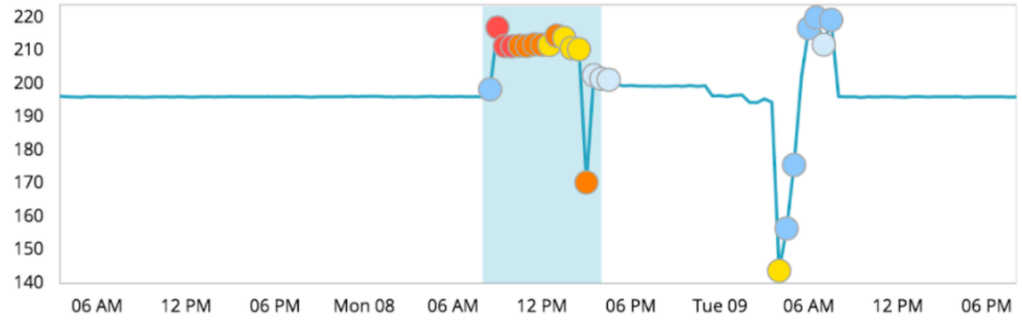

mean SQLServer_SQL_Statistics_Batch_Requests_sec hostname MSSQL-0783E4076

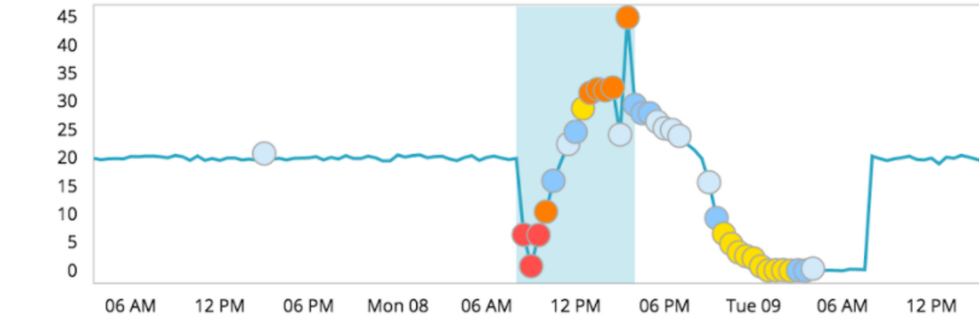

Coming soon

Elastic Cloud Enterprise

1.0.0-beta2 released

Elasticsearch SQL

Coming soon!



Kibana Canvas

Coming soon!

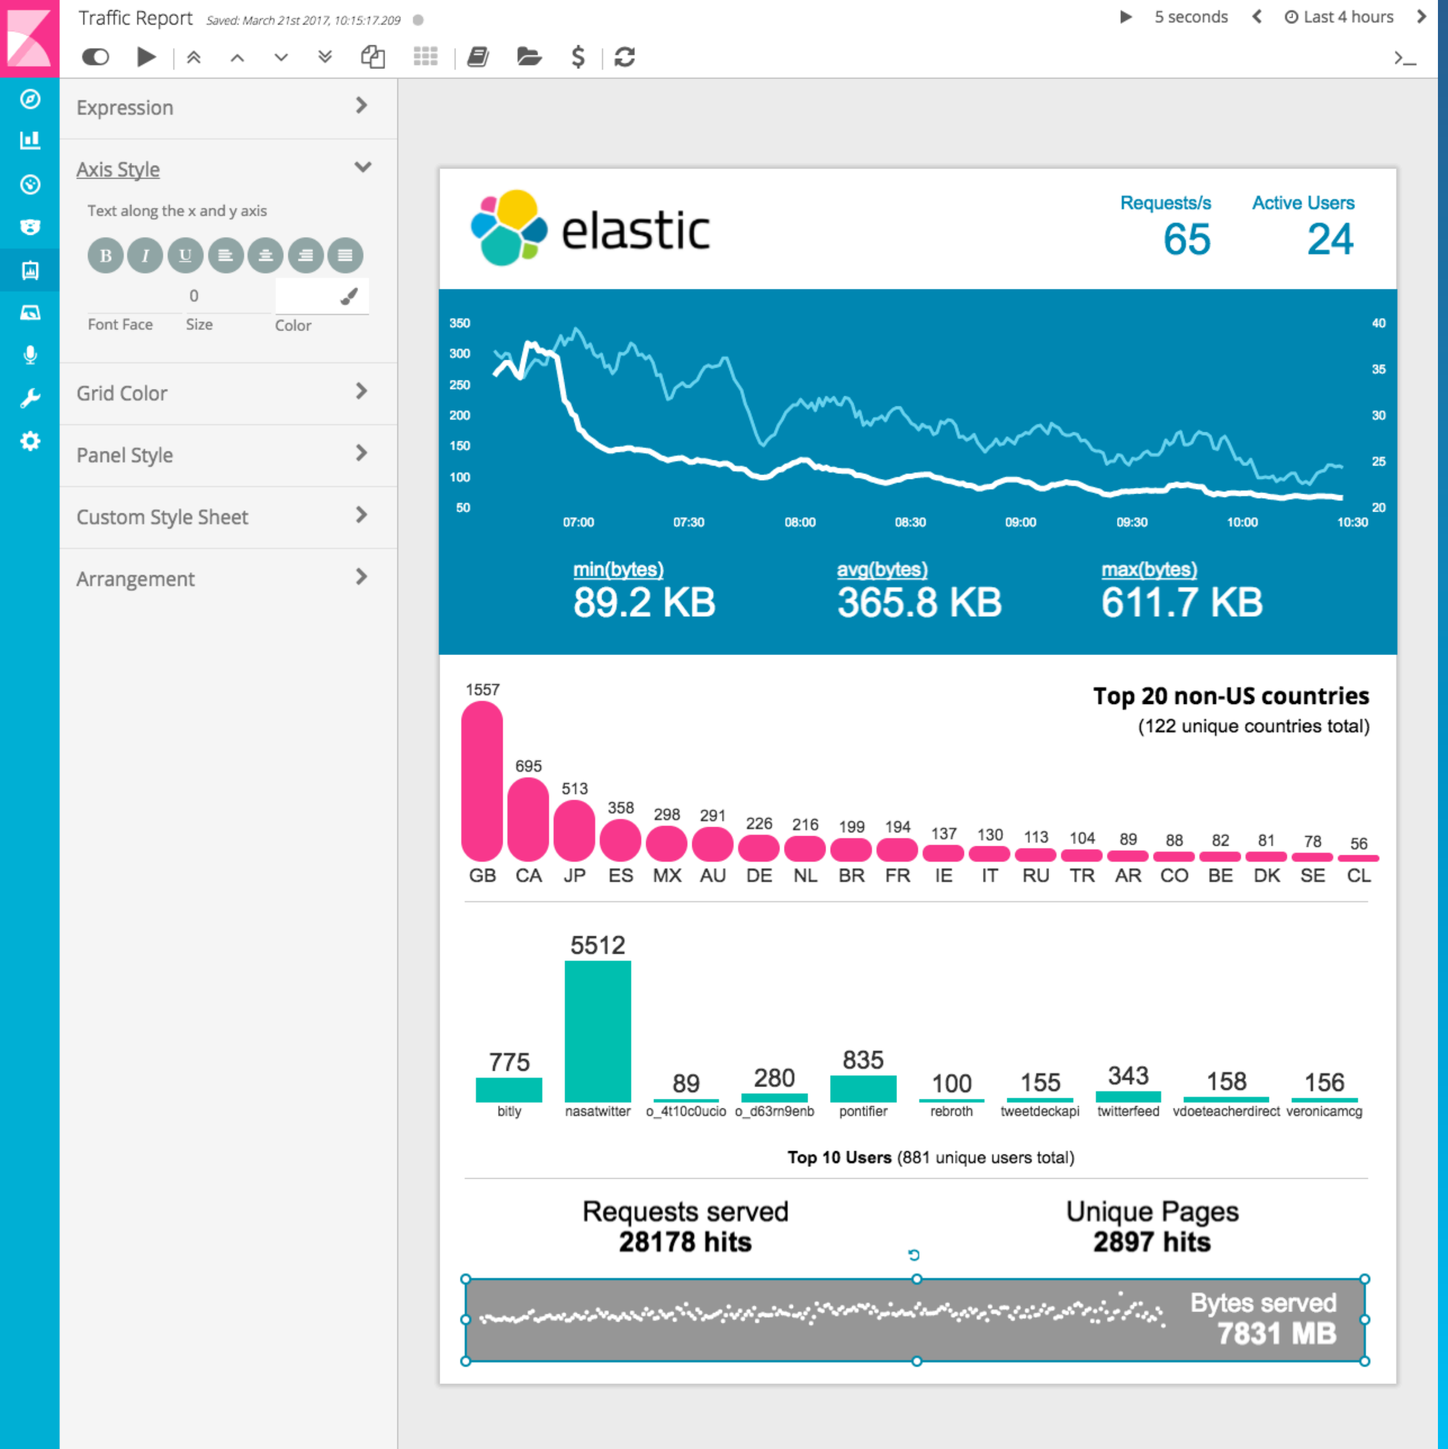

- New visualization application on top of Elasticsearch data
- Use Case:
  - live infographics
  - presentations with live data feeds
  - highly customized reports
- Currently, in the prototyping phase
- Release date: TBD

Volunteers *Saved: March 21st 2017, 10:13:15.609*

Last 1 year

Page Background

Image | Lightbox | Vertical_b... | Horizonta... | Variable S... | String Input | Number S... | Circles | Grow | Number | Repeat | Markdown | Timechart | Debug

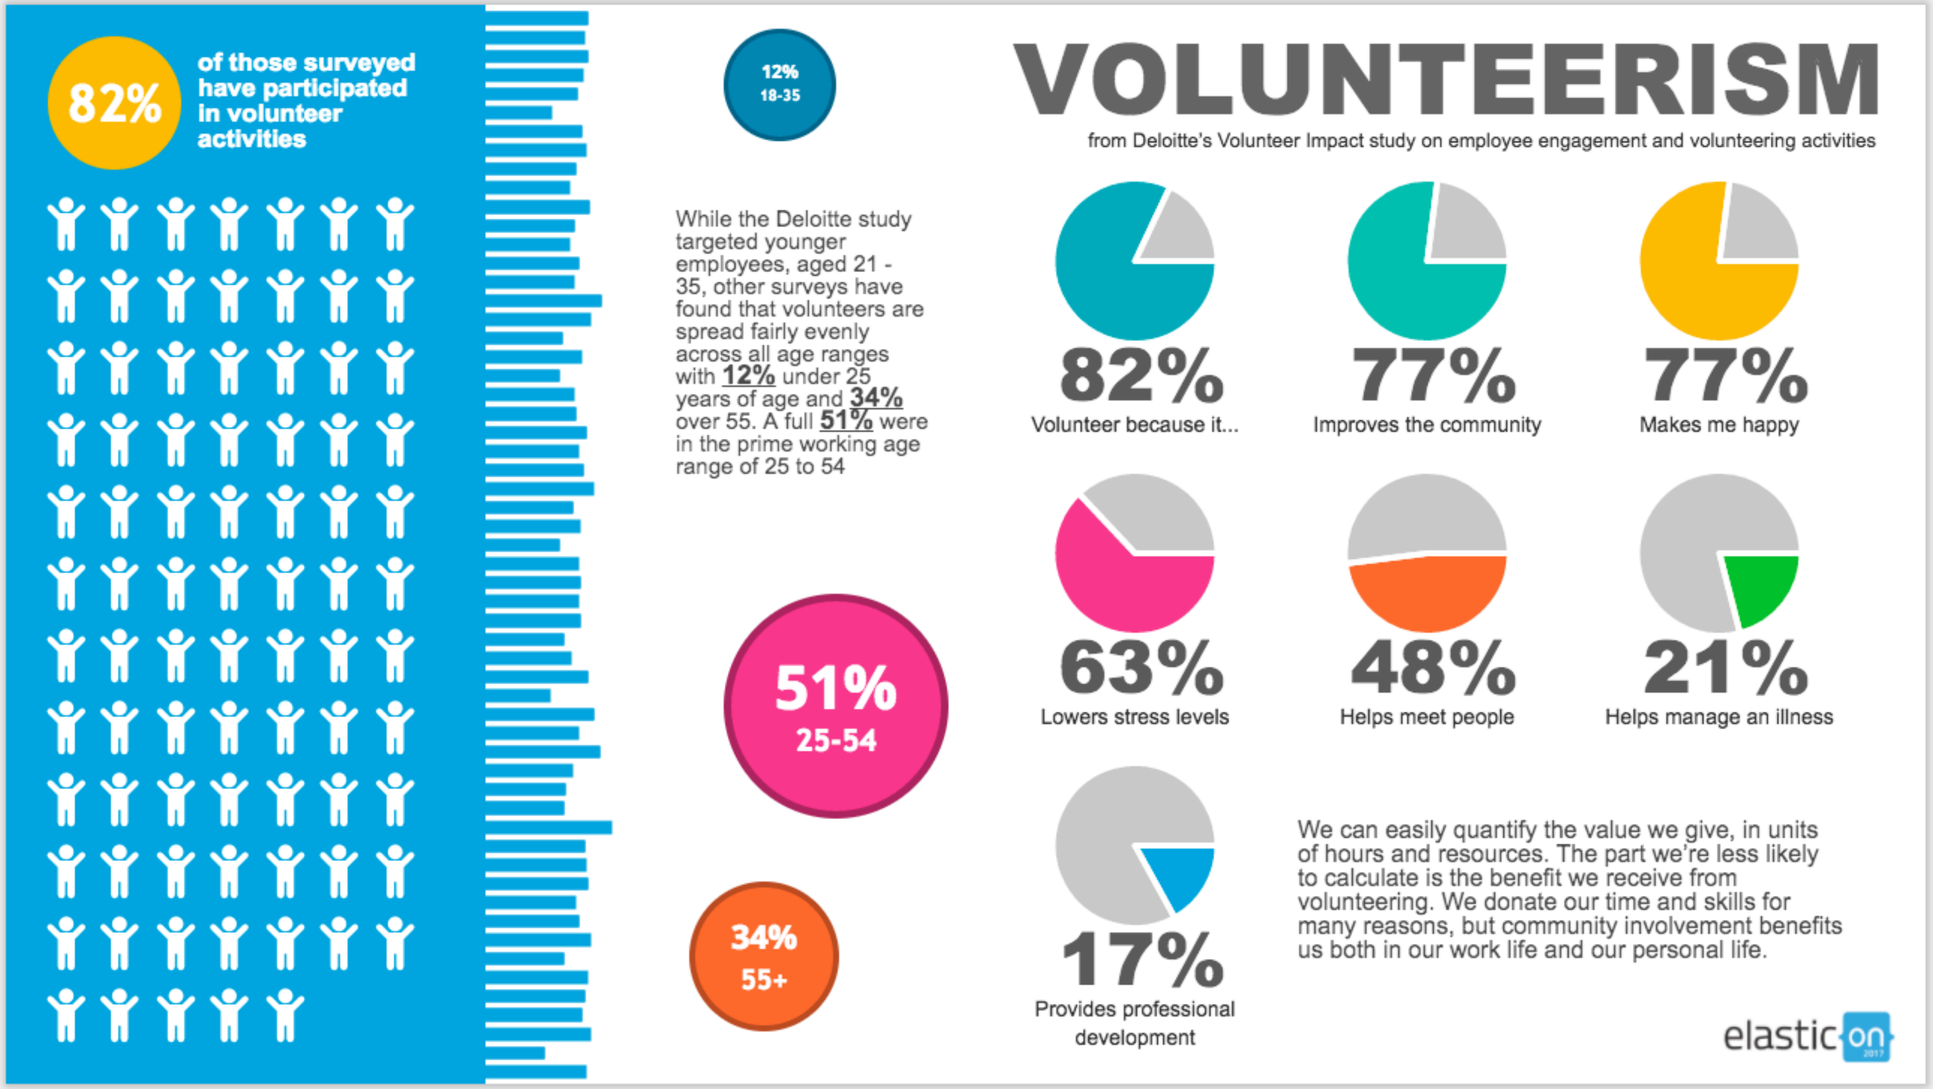

6.0.0 alpha1