Elastic Stack & What’s New Released

Medcl

About Me

- Medcl, Zeng Yong, 曾勇
  - http://github.com/medcl
  - http://weibo.com/medcl
  - http://log.medcl.net
  - @medcl
- Me & Elasticsearch
  - Follow with elasticsearch since 2010.Mar.23, v0.5.1, 3rd released version
  - Maybe first 200 user worldwide, and first 10 user in China

Elastic proprietary and confidential.
Please do not share or forward.

Introducing the Elastic Product Portfolio

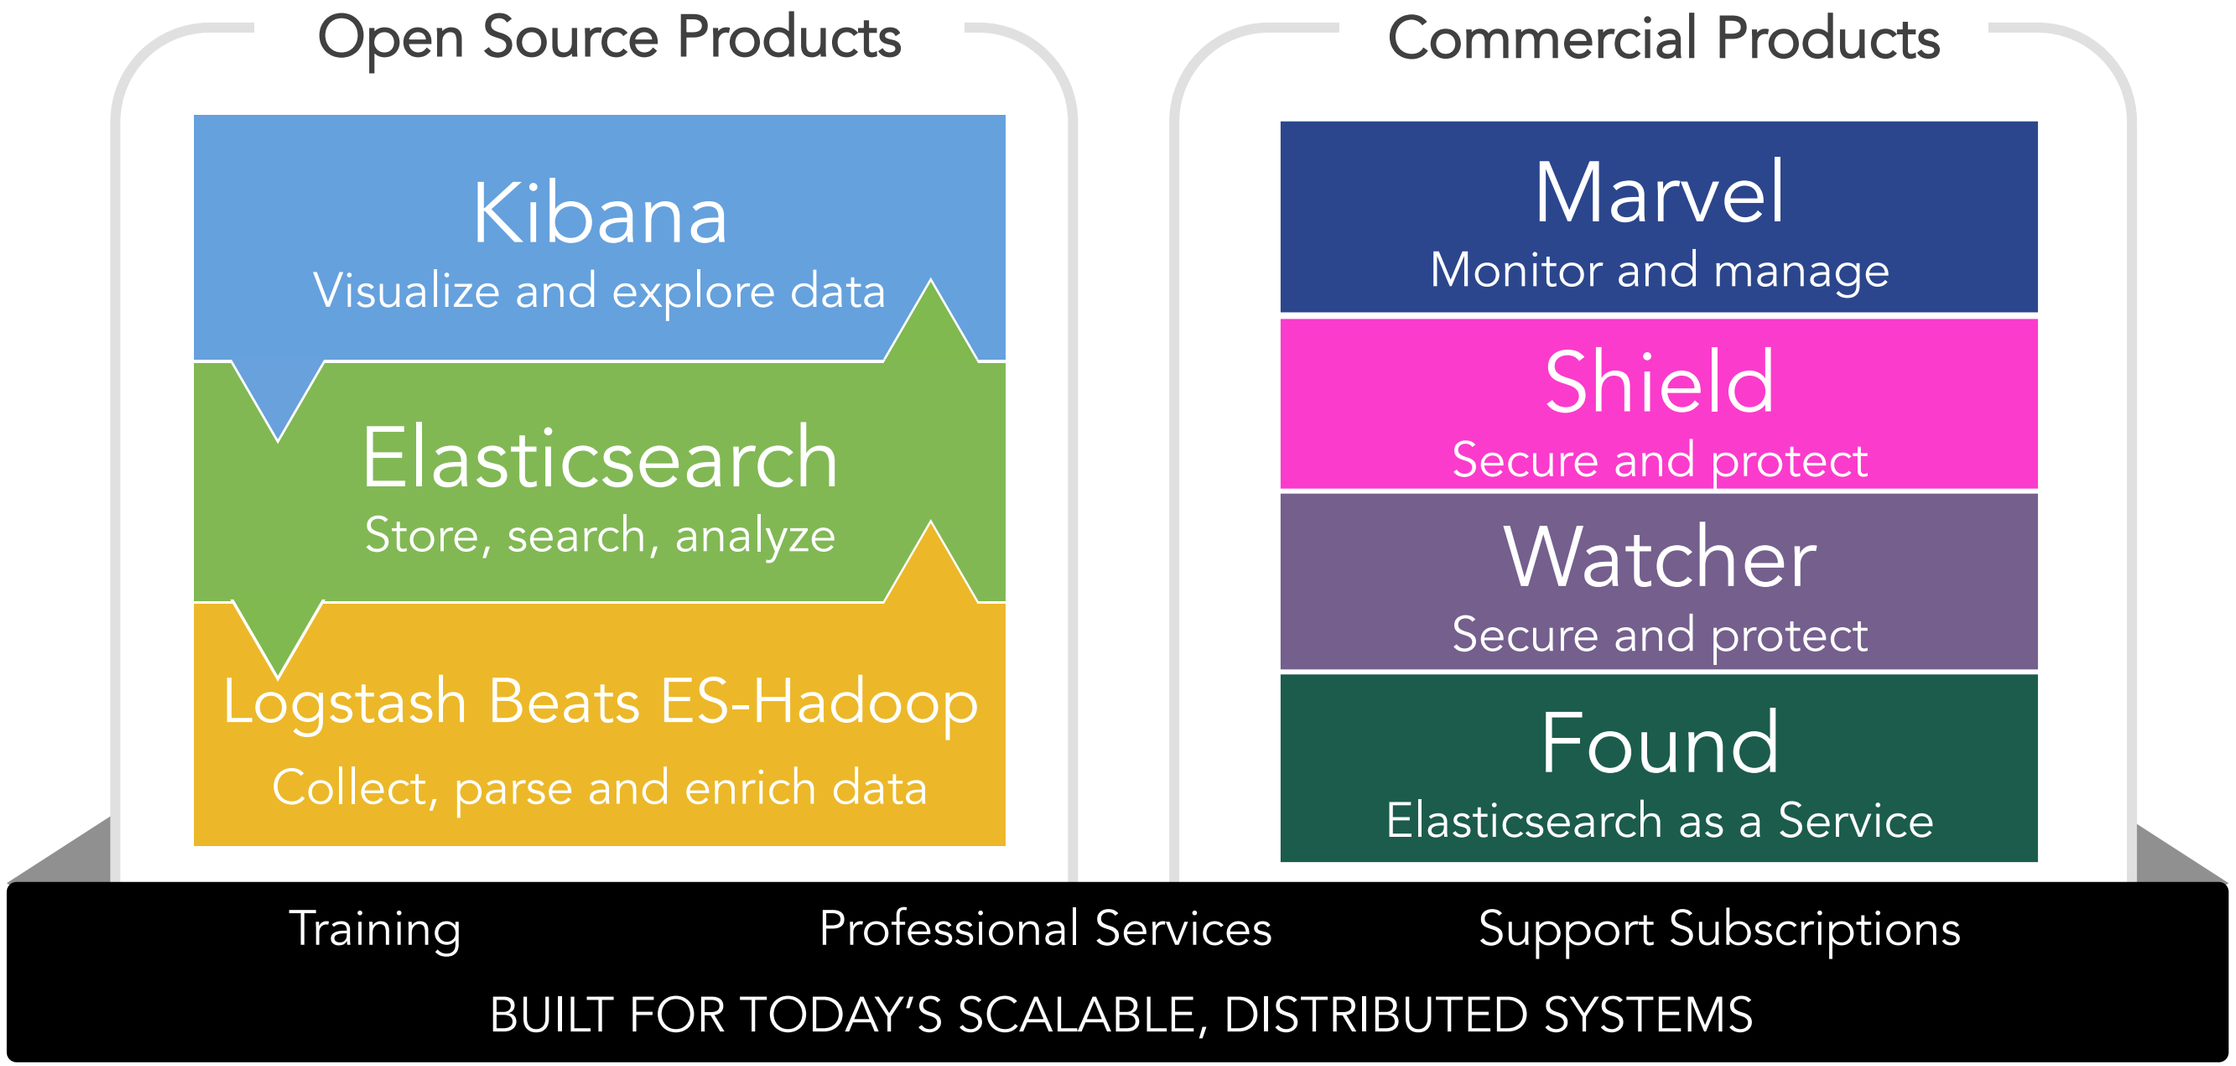

Elastic proprietary and confidential.
Please do not share or forward.

The Elastic Stack

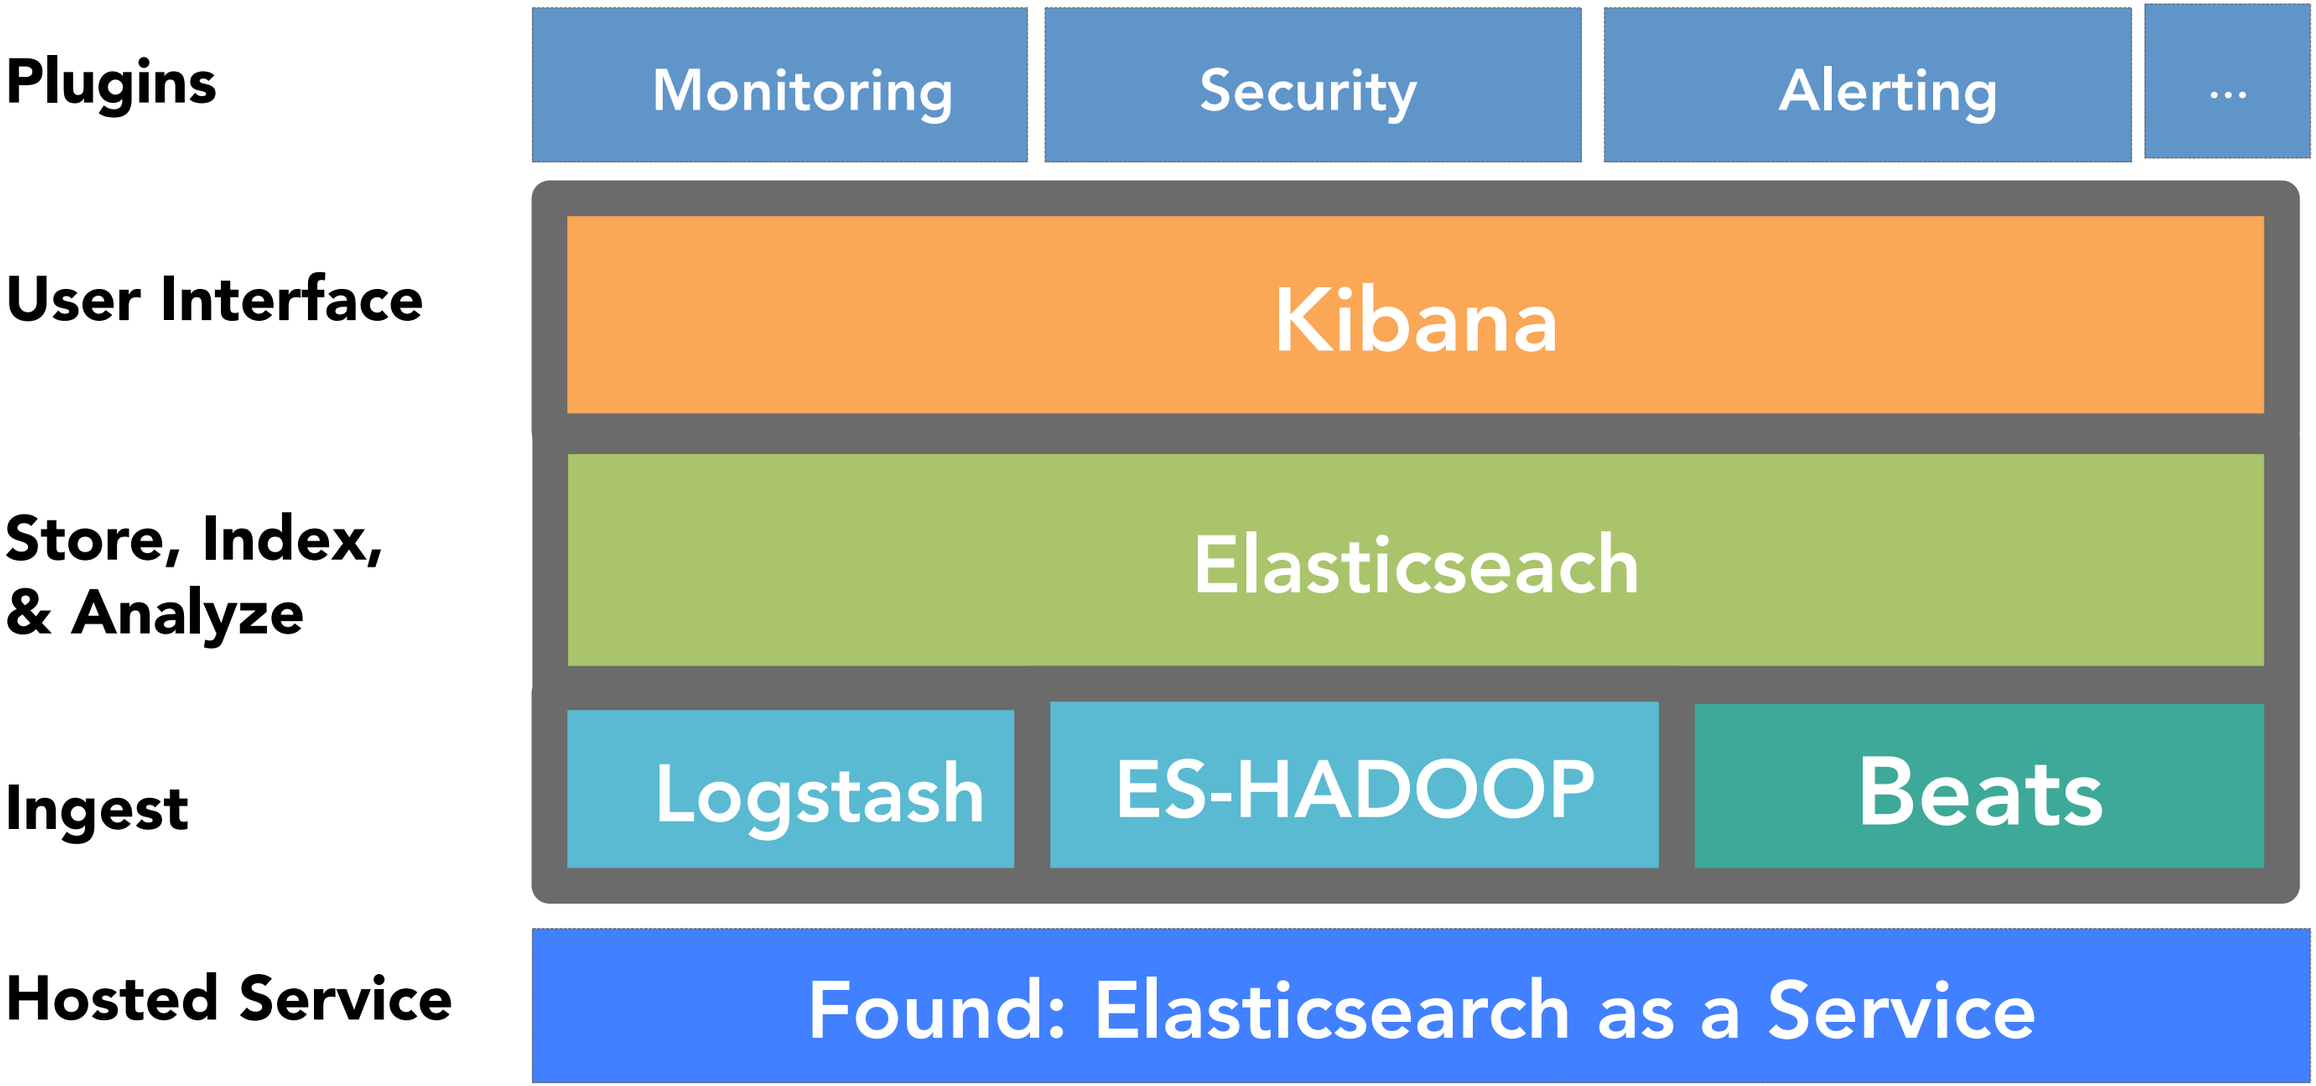

Elastic proprietary and confidential.
Please do not share or forward.

elastic

Elasticsearch

Elasticsearch 2.0

- New Features
  - Pipeline Aggregations
  - Query DSL/Doc Improvement
  - Index Compression
- Performance & resilience
  - Lucene 5.2
  - Update Cluster State with diffs
  - Doc_values by default
  - Sync-flush (1.6+)
  - Better handling for node-leave/rejoin (1.7+)
  - Durability-by-default
  - Async shard allocation (1.6+)
- Breaking Backward Compatibility
  - Facets – removed
  - Zen discovery improvements
  - Type mappings are now strict
  - Index segments created before ES .90.0 must be upgraded
  - Migration Assistant
  - Units are required in settings

Elastic proprietary and confidential.
Please do not share or forward.

Elasticsearch正成为企业级搜索引擎行业第一

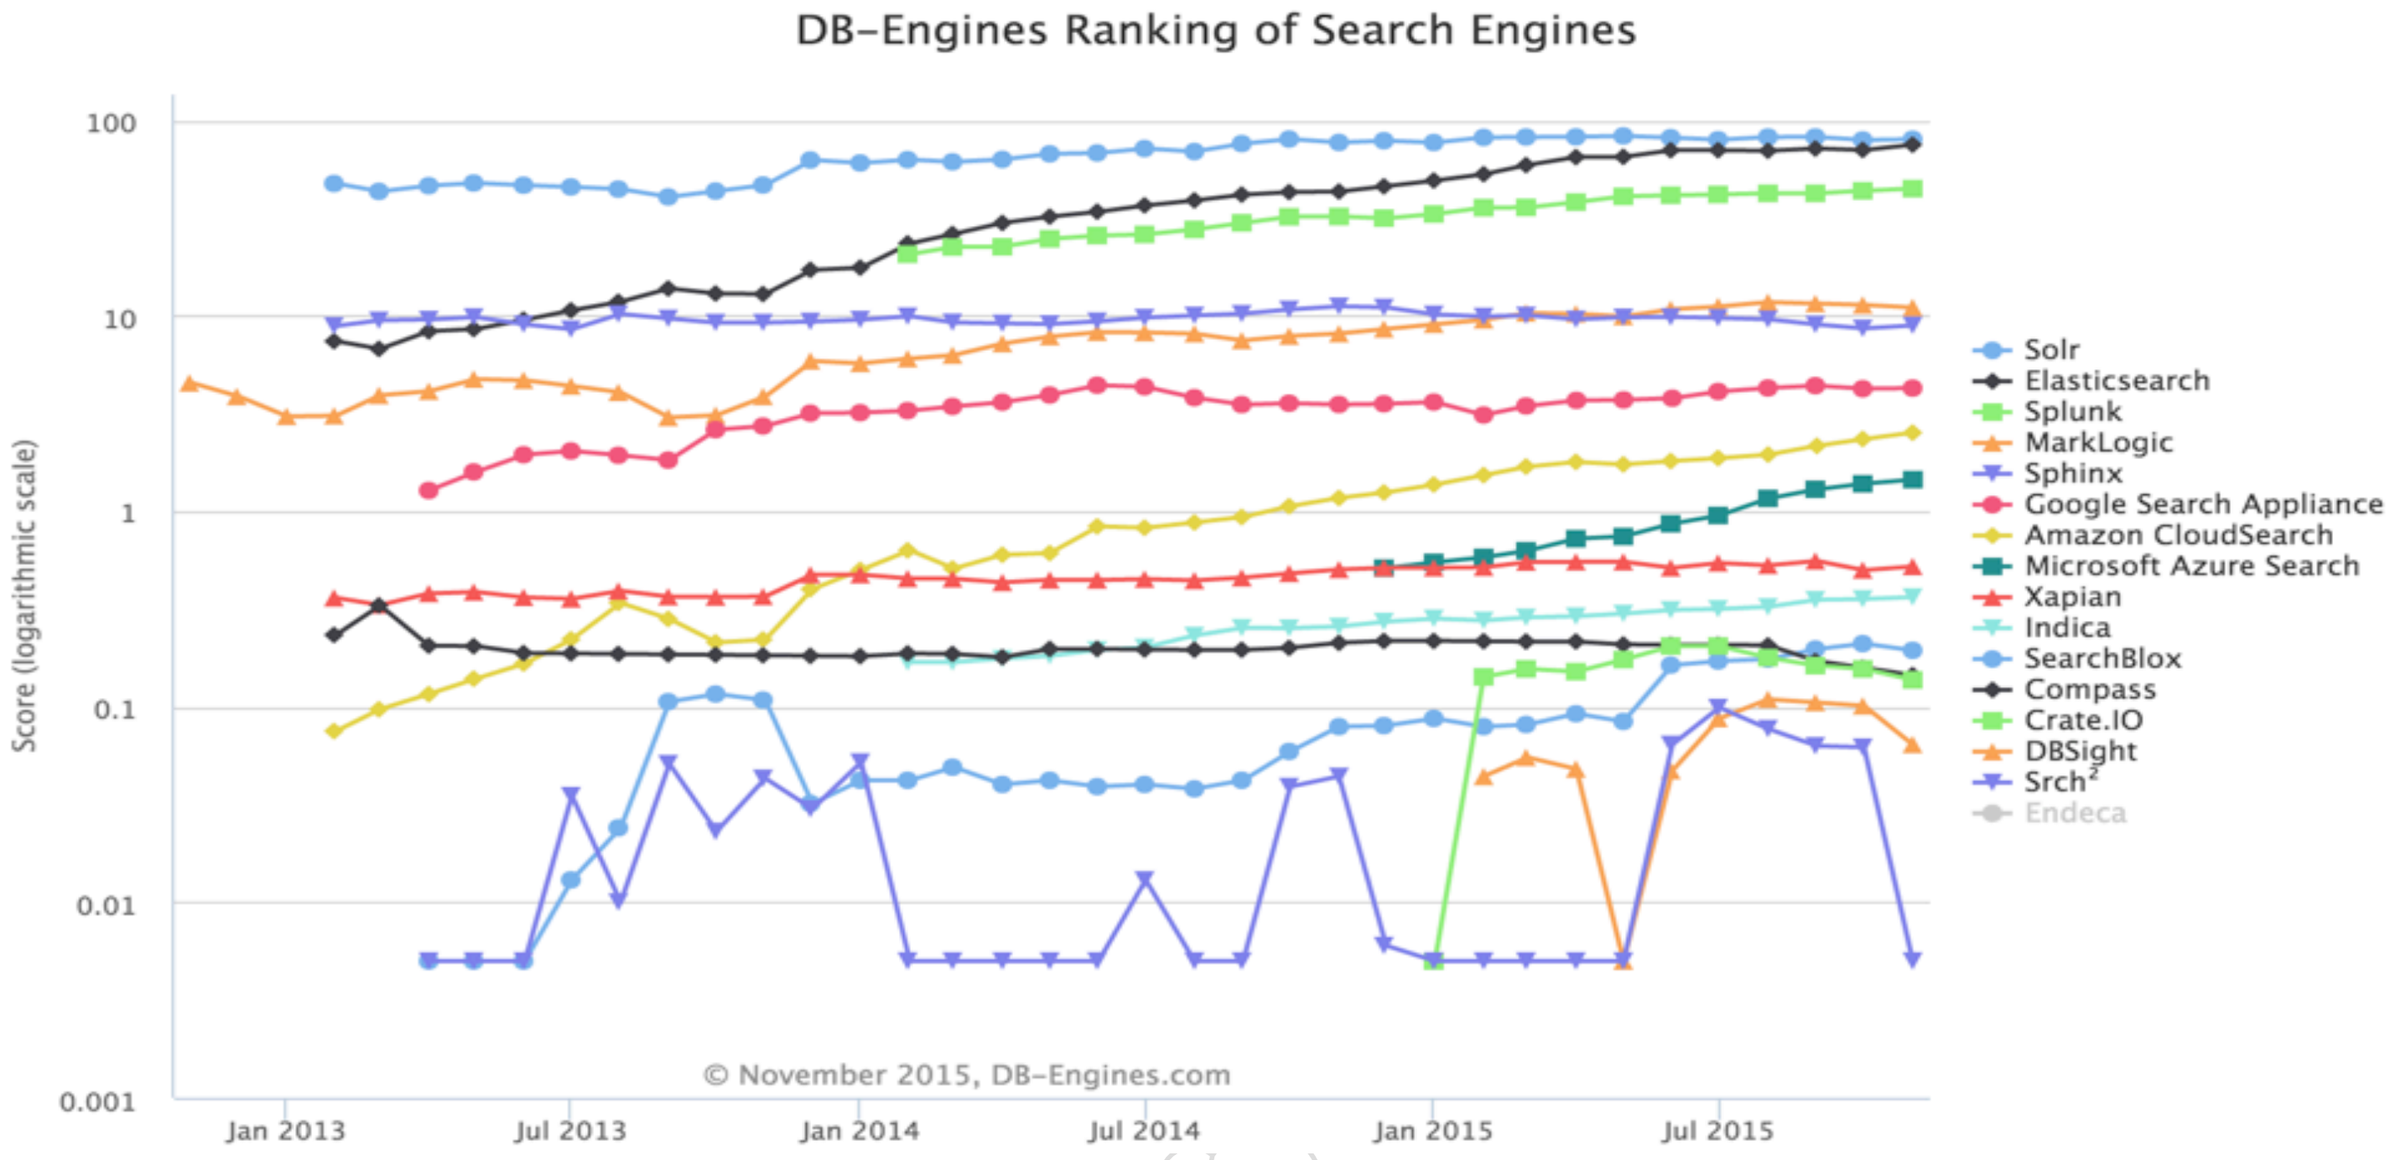

来自: http://db-engines.com/en/ranking_trend/search+engine

Elasticsearch 2.0 – New Features

- Pipeline Aggregations
  - Derivatives
  - Moving average
  - Holt Winters (prediction / anomaly detection)
  - Stats: Min/Max/avg
  - Time-series math
- Query DSL/Doc Improvement
  - All non-scoring components are cache-able (Filters vs Queries)
- Index Compression
  - 10-30% reduced index size, some indexing/merging impact
  - Dynamic setting – could set before optimization for time-series indexes

Elastic proprietary and confidential.
Please do not share or forward.

Aggregations

```
GET /person/person/_search?search_type=count
{
    "aggs": {
       "by_country": {
          "terms": {
             "field": "address.country"
          }
       }
    }
}
```

```
{ ..., "aggregations" : {
     "by_country" : {
      "buckets" : [ {
        "key" : "England",
        "doc_count" : 30051
      }, {
        "key" : "Germany",
        "doc_count" : 30004
      }, {
        "key" : "France",
        "doc_count" : 15034
      }, {
        "key" : "Spain",
        "doc_count" : 14912
      } ]}}}
```

Like facets but with more power
Can be nested to add additional dimensions
Give analytical insights into data
Allow complex visualizations
Major types: buckets and metrics
Types: terms, histogram, percentiles, etc.

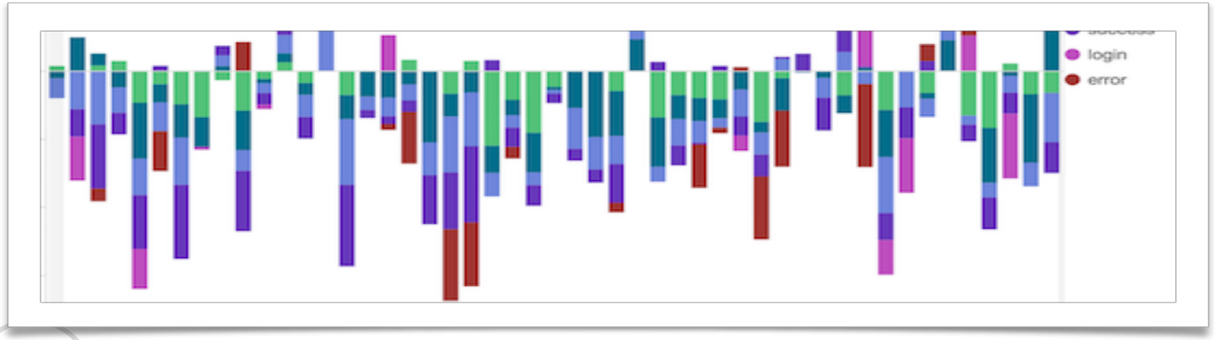

Elastic proprietary and confidential.
Please do not share or forward.

Pipeline aggregations

Work on outputs of other aggregations
Used for smoothing, prediction, etc.
Different types: avg, derivative, max, min, sum moving avg, cumulative sum, etc.

```
{
    "my_date_histo":{
        "date_histogram":{
            "field":"timestamp",
            "interval":"day"
        },
        "aggs":{
            "the_sum":{
                "sum":{ "field": "lemmings" }
            },
            "the_movavg":{
                "moving_avg":{ "buckets_path": "the_sum" }
            }
        }
    }
}
```

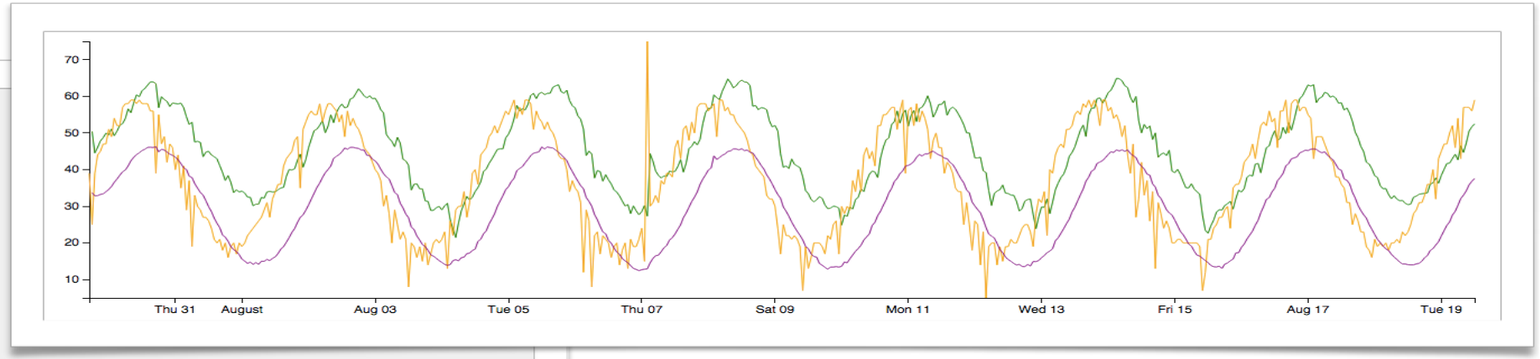

Elastic proprietary and confidential.
Please do not share or forward.

Elasticsearch 1.6-2.0 – Performance & Resilience

- Lucene 5.2
  - Memory reduction when merging
  - Improved file handling and corruption detection
- Update Cluster State with diffs
  - Reduces network traffic, improves scalability
- Doc_values by default
  - Dramatic memory reduction by default
- Sync-flush (1.6+)
  - Faster rolling cluster restarts
- Better handling for node-leave/ rejoin (1.7+)
  - Better handle network interruptions or server reboots
- Durability-by-default
  - Sync transaction log on every request
  - Small performance impact for bulk requests
- Async shard allocation (1.6+)
  - Do not wait for expensive filesystem commands during shard allocation
  - No more long pending task queue

Elastic proprietary and confidential.
Please do not share or forward.

Elasticsearch 2.0 – Breaking Backwards Compatibility

- Zen discovery improvements
  - No rolling cluster update
- Index segments created before ES .90.0 must be upgraded (via 1.6)
  - Lucene 5 does not support Lucene 3.x segments
- Facets, Rivers removed
  - Facets have been removed in favor of Aggregations
  - Support for previously-deprecated Rivers removed
- Type mappings are now strict
  - Two types may not define the same field with different settings
- Units are required in settings
  - May need to update your elasticsearch.yml, but worth it
- Migration Assistant
  - Site plugin for ES 1.x that checks your cluster for BWC issues

Elastic proprietary and confidential.
Please do not share or forward.

elastic

Kibana

Why Kibana 4?

Ease of Use

- Separate tasks
- Drag and drop
- Re-usable components

Advanced Analytics

- Elasticsearch Aggregations
- Multi-dimensional visualizations

Sharing & Integration

- Export data: CSV
- Embeddable charts and dashboards
- Mobile support

Elastic proprietary and confidential.
Please do not share or forward.

Kibana 4: Ease of Use

Separate tabs for different interaction types

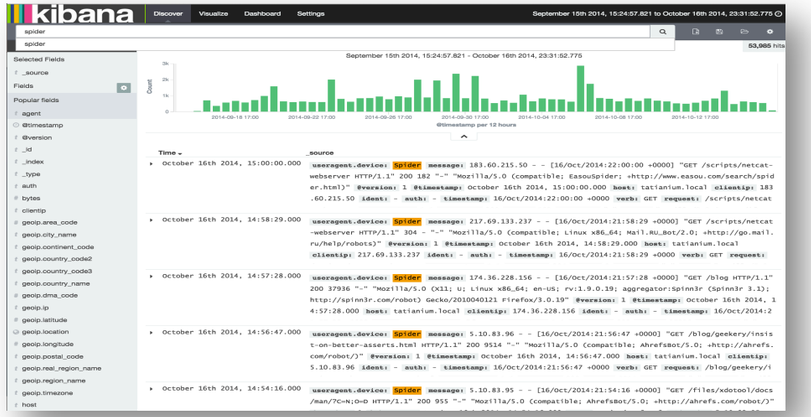

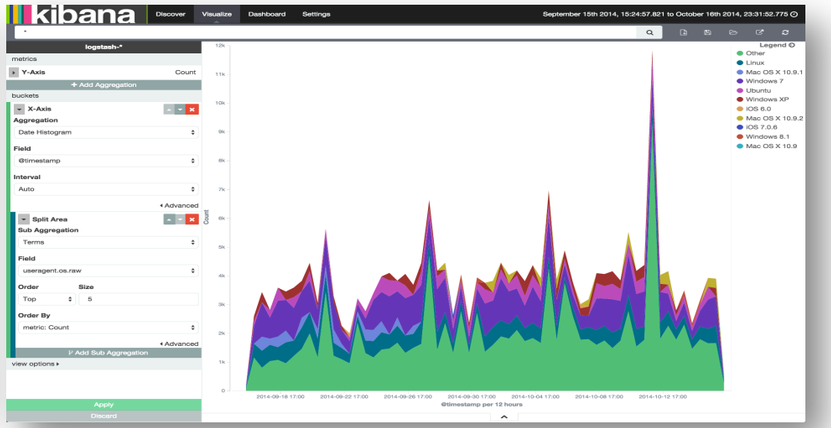

Discover

- Field discovery
- Ad-hoc search
- Top values

Visualize

- Build a single chart
- Experiment with different analytics

Dashboard

- Combine charts
- Filter and search across multiple visualizations

Kibana 4: Advanced Analytics

Support for visualizing nested aggregations

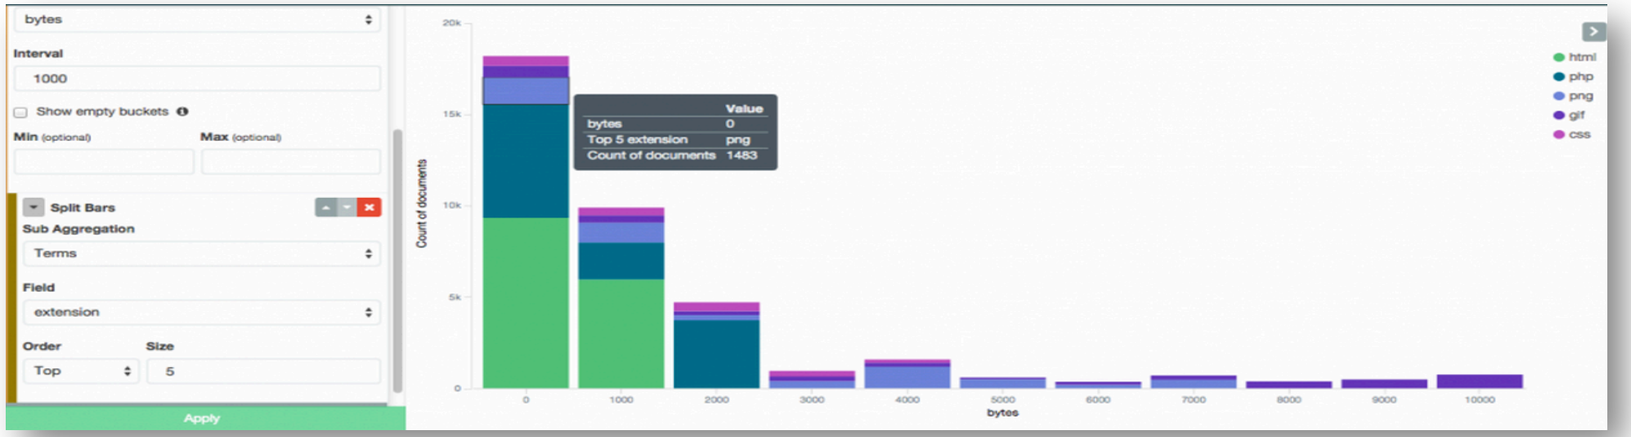

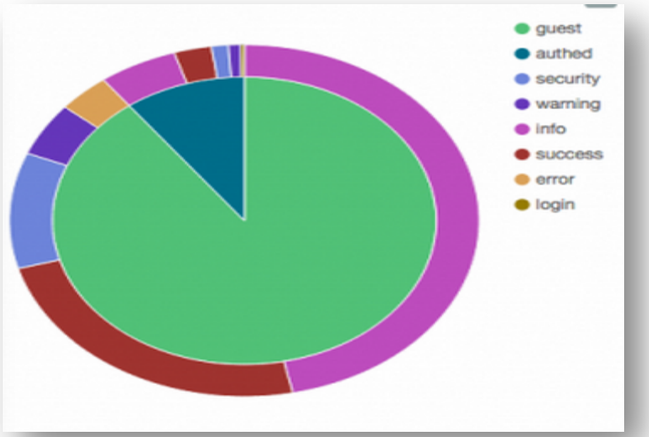

Multiple aggregations on a single chart

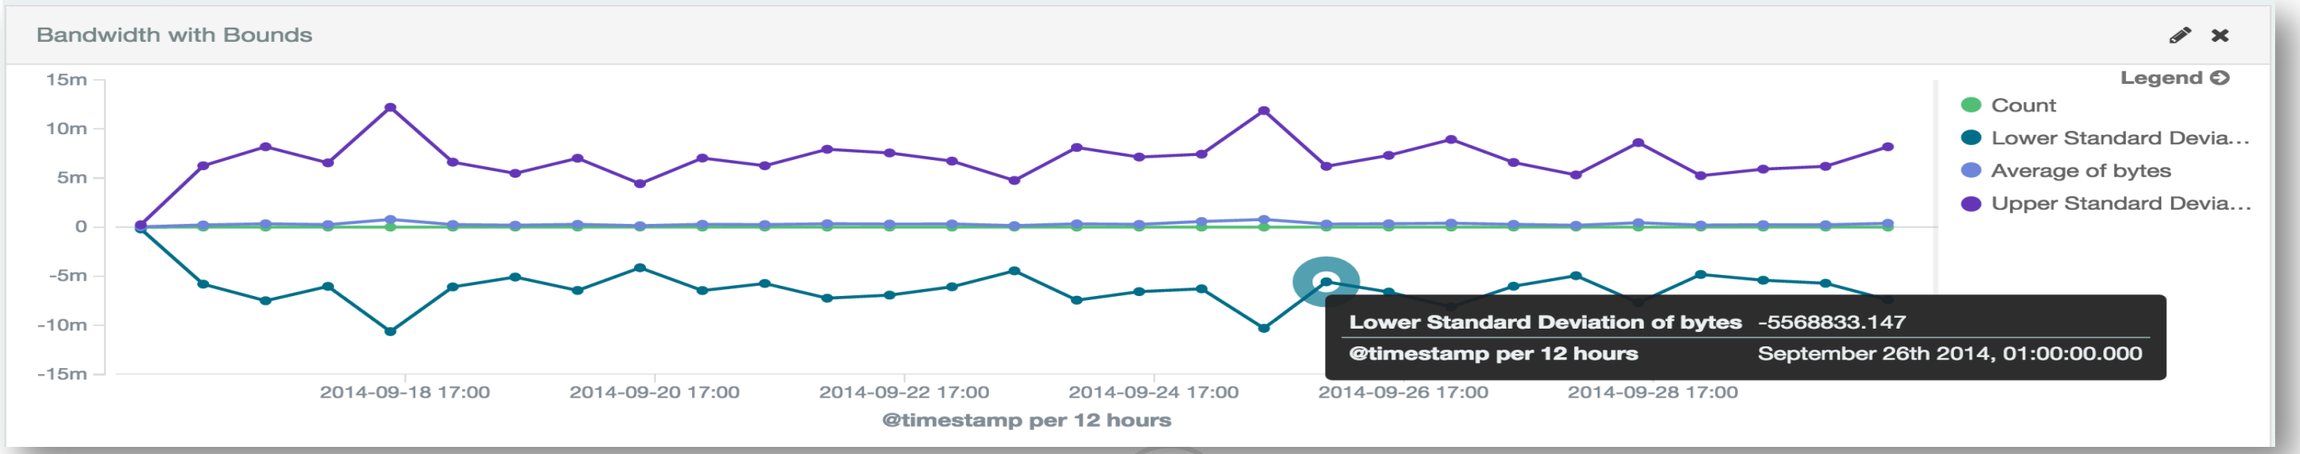

Elastic proprietary and confidential.
Please do not share or forward.

Kibana 4: Sharing and Integration

- Embeddable charts
- URL contains dashboard + filters; designed for sharing
- CSV data export
- Mobile support
- Shield support for security

Elastic proprietary and confidential.
Please do not share or forward.

Kibana 4.1-4.2: Customizability

Dark theme

Field formatters

Customizable maps

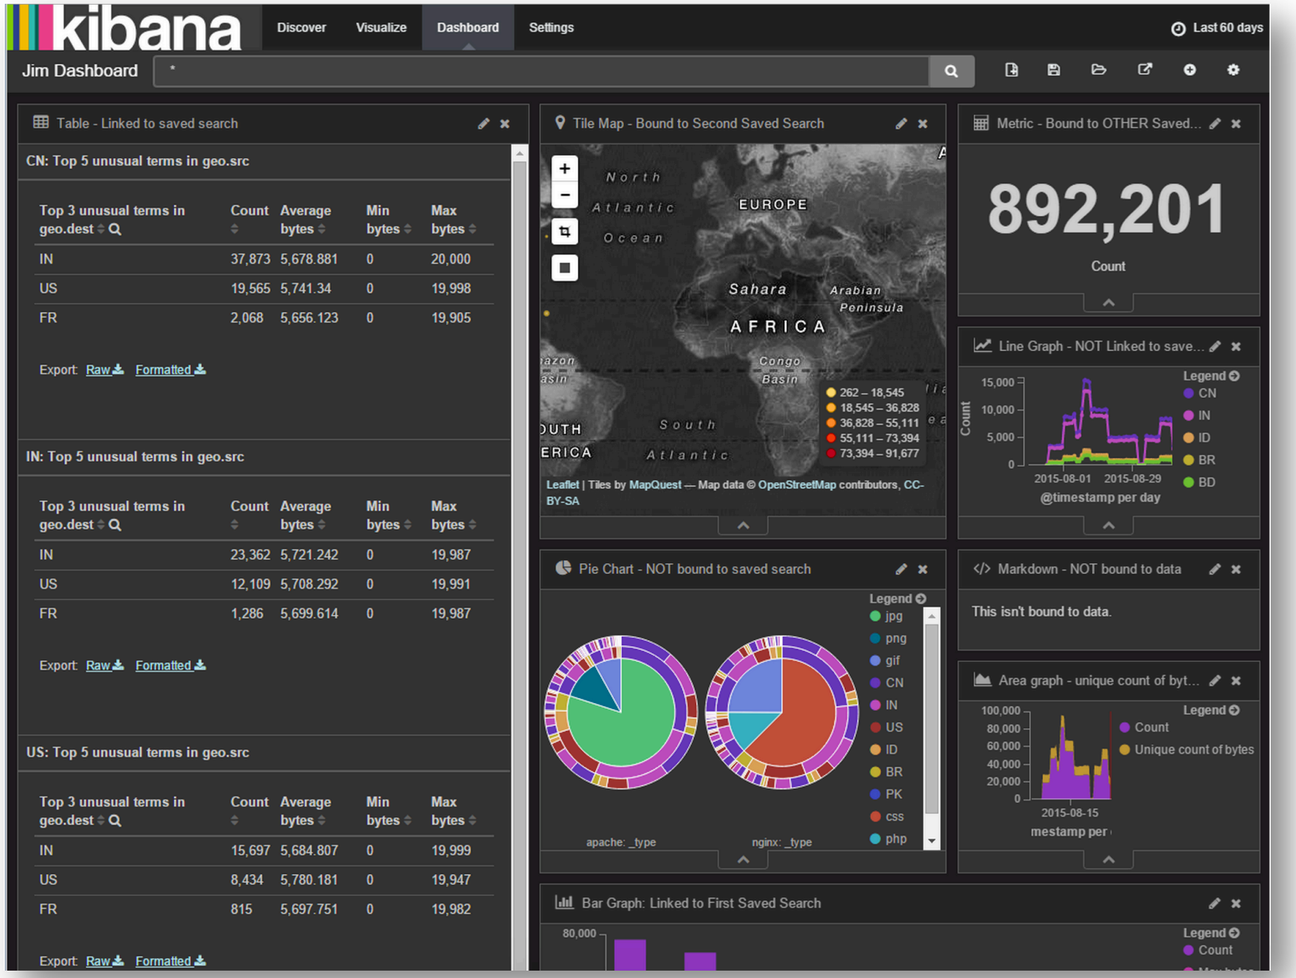

Elastic proprietary and confidential.
Please do not share or forward.

Kibana 4.1-4.2: Customizability

Dark theme

Field formatters

Customizable maps

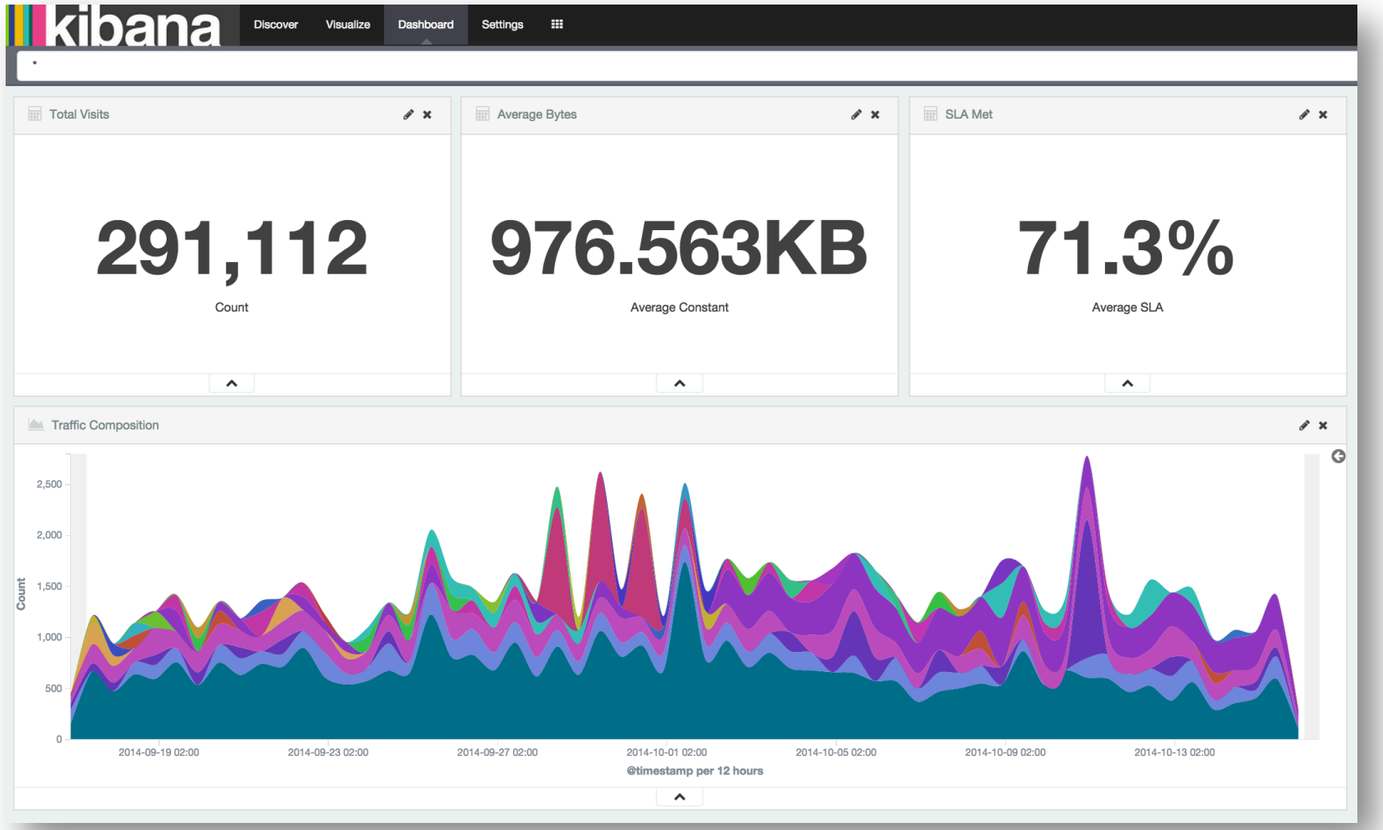

Elastic proprietary and confidential.
Please do not share or forward.

Kibana 4.1-4.2: Customizability

Dark theme

Field formatters

Offline/customizable maps

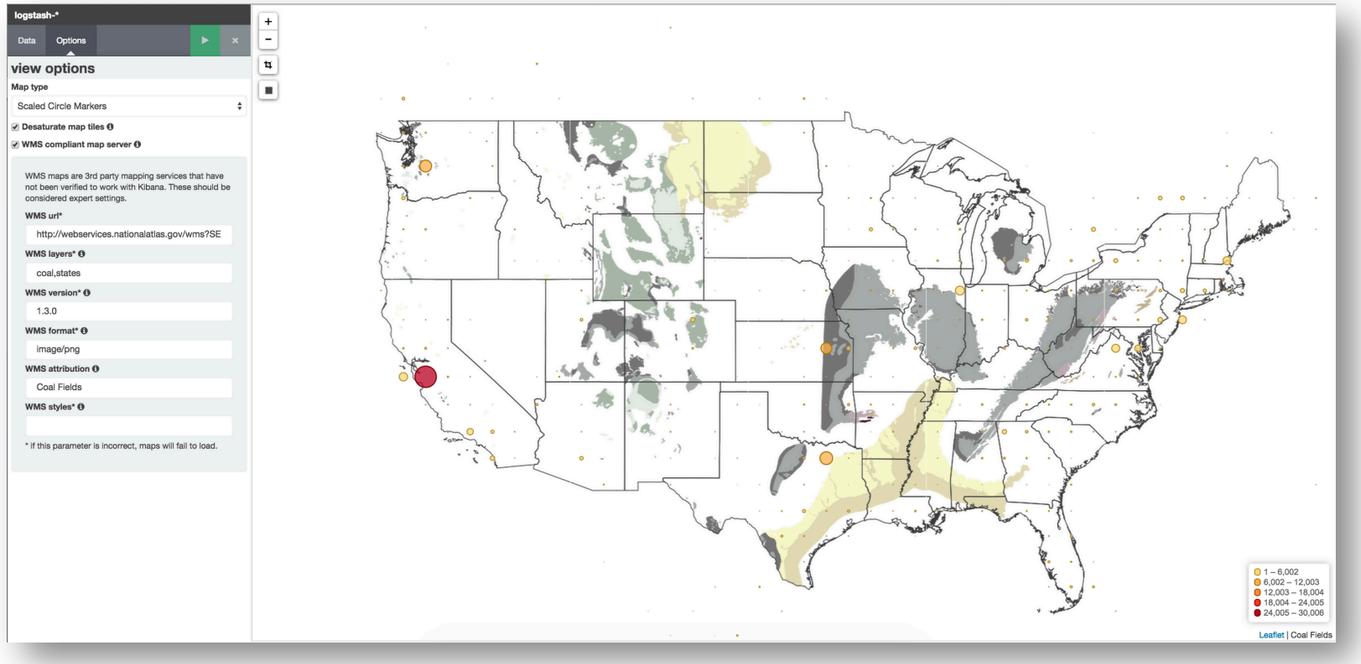

Elastic proprietary and confidential.
Please do not share or forward.

Kibana 4.1-4.2: Manageability

Kibana server status page

Configurable log levels

Saved object export

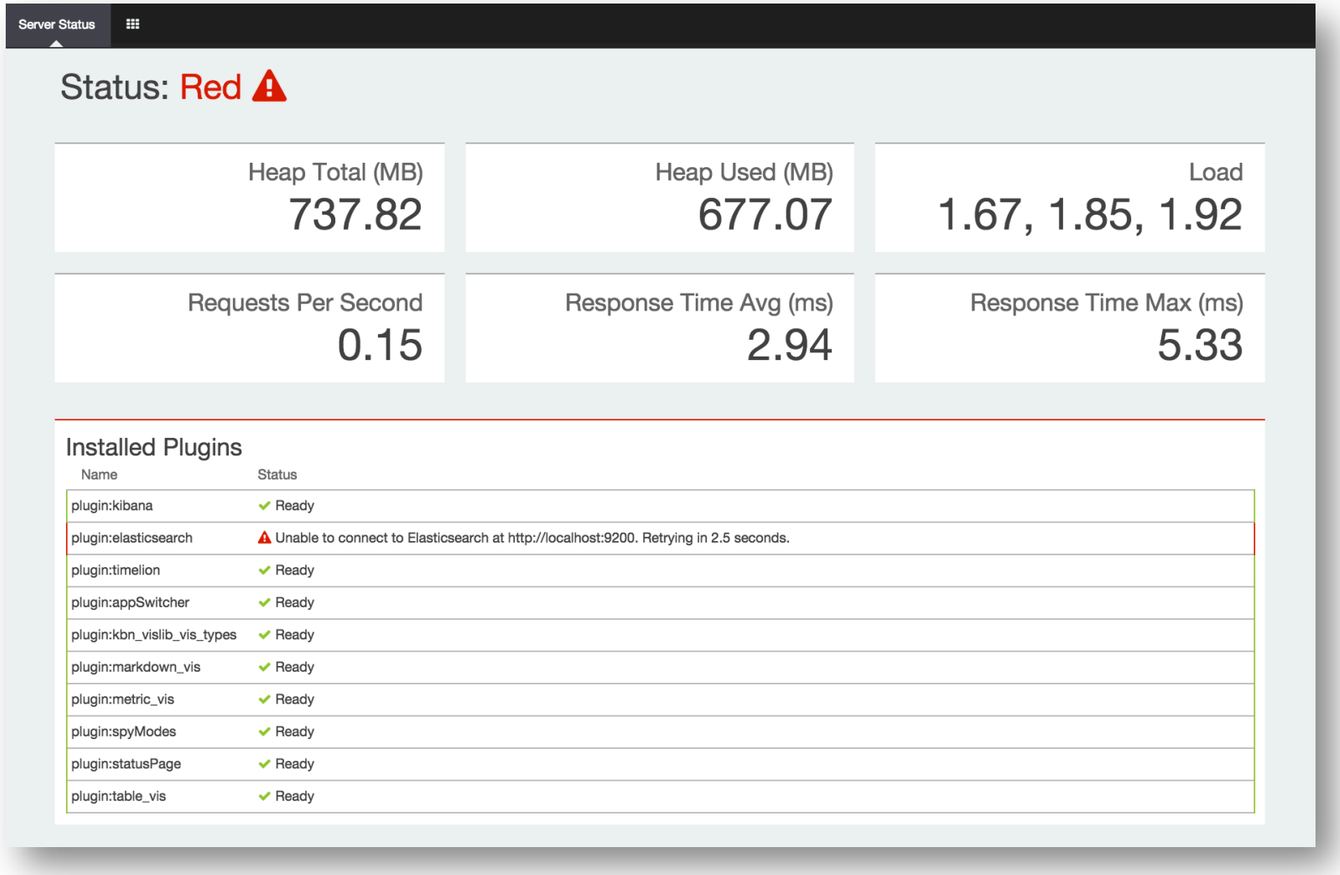

Elastic proprietary and confidential.
Please do not share or forward.

Kibana 4.1-4.2: Analytics

Bubble charts

Geo heatmap

New aggregations (e.g. IP range)

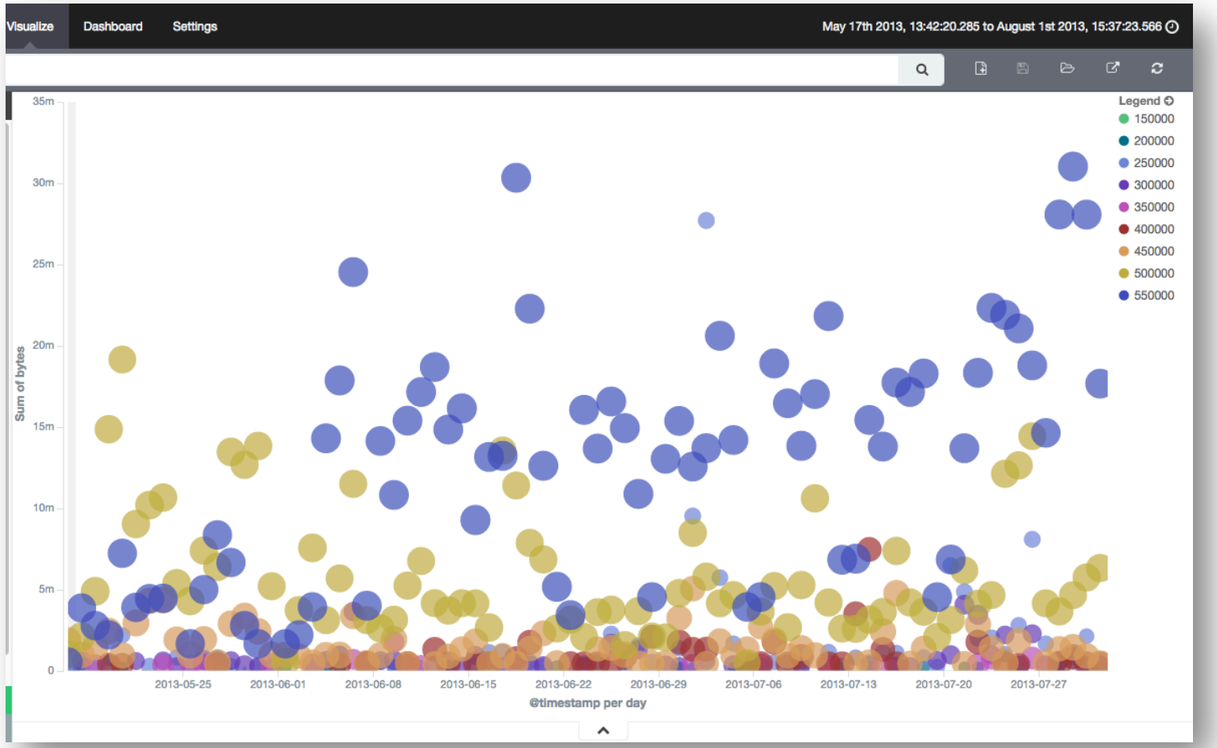

Elastic proprietary and confidential.
Please do not share or forward.

Kibana 4.1-4.2: Usability

Map filters

Pinned filters

Clickable legends

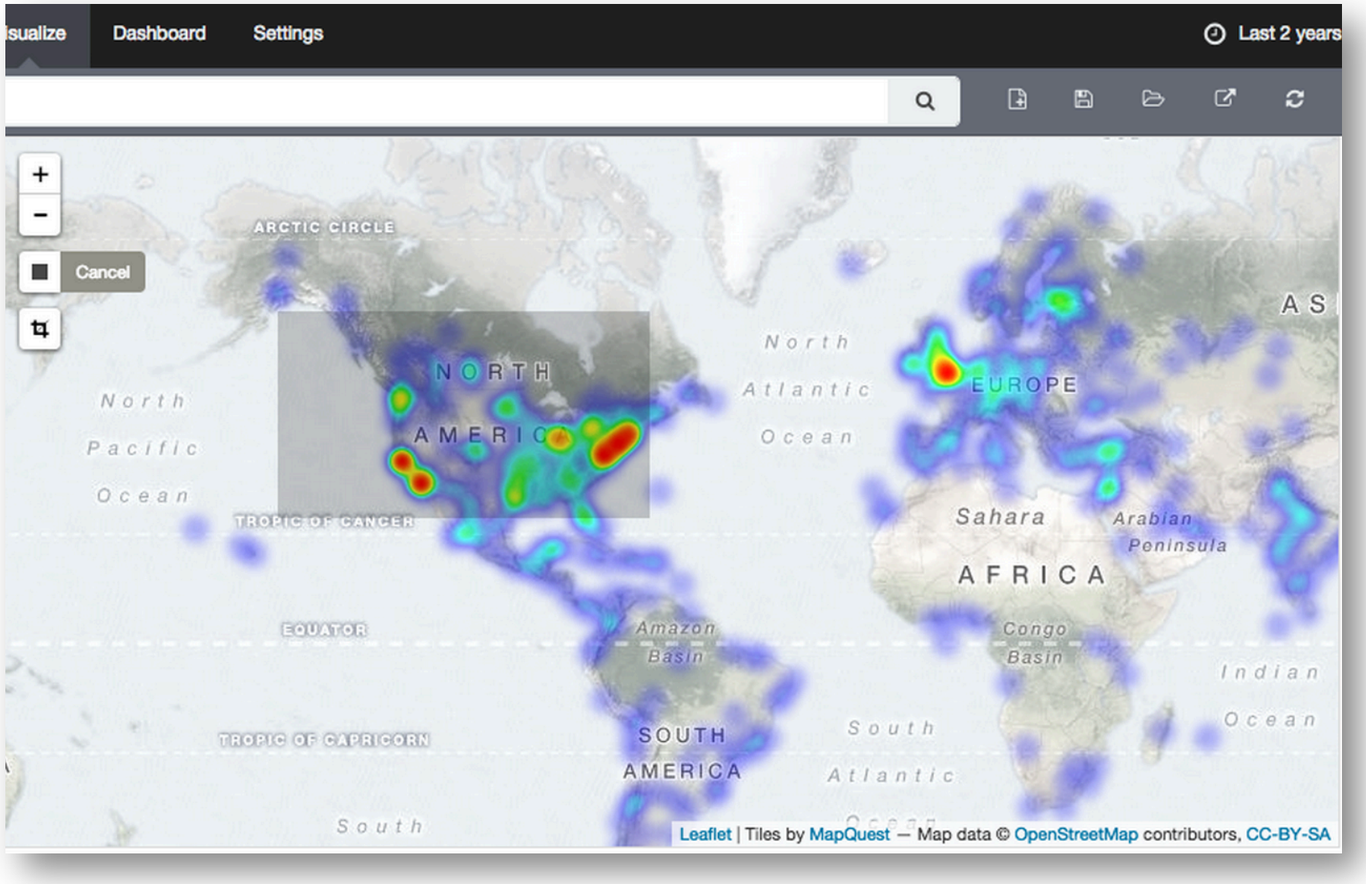

Elastic proprietary and confidential.
Please do not share or forward.

Kibana migration tips

- Kibana 3 to Kibana 4 (Kibana 3 EOL is Nov 2015)
  - Embedded web server
  - Platform-specific installation packages
  - Performance improvements
  - Dashboards will not be migrated
  - Some panel types not available, yet
- Kibana 4 to Kibana 4.2
  - Support for Elasticsearch 2.x
  - Not backward-compatible with Elasticsearch 1.x
  - Dashboards are automatically migrated

Elastic proprietary and confidential.
Please do not share or forward.

elastic

Data Ingest
(Logstash, Beats, ES-Hadoop)

Logstash

- Logstash 1.5
  - Plugin management improvements
  - Grok performance improvements (2-3x)
  - Heartbeat plugin for monitoring of LS health
- Logstash 2.0
  - Elasticsearch compatibility
  - HTTP as default transport protocol
  - Better shutdown semantics
- New logstash plugins
  - Kafka input/output
  - JDBC input
  - HTTP input
  - WebHDFS output
  - Salesforce input

Elastic proprietary and confidential.
Please do not share or forward.

Beats 1.0 Beta 1-3

- More Packetbeat protocols
  - MongoDB
  - DNS
  - Memcache
- More Beats
  - Topbeat: Shipper for CPU, memory, process resource metrics
- Improved platform support
  - Windows support (e.g. self-contained installer)
- Developer guides
  - Building Beats
  - Building Packetbeat protocol modules

Elastic proprietary and confidential.
Please do not share or forward.

Elasticsearch for Apache Hadoop

2.0

- Native, bi-directional integration with Apache Hadoop
  - Query from and write to Elasticsearch from Hadoop/HDFS
  - Allows Elasticsearch to snapshot & restore into HDFS
- Supports the use of MapReduce, Apache Hive, Apache Pig, and Cascading

2.1

- Native support for Apache Spark, SparkSQL, and Apache Storm
- Basic authentication, PKI, and SSL/TLS support for Elasticsearch Shield
- Elasticsearch on YARN (beta)

2.2 (beta)

- Elasticsearch 2.0 support
- Support for Elasticsearch aggregations

Elastic proprietary and confidential.
Please do not share or forward.

elastic

Commercial Products

Shield: Security for Elasticsearch

Shield 1.x

- Authentication
- AD & LDAP Auth
- Index/document/alias level RBAC
- SSL encryption
- Audit log

Shield 2.0

- Field and document level access control
- Custom authentication/authorization realms
- Run-as, user impersonation

Elastic proprietary and confidential.
Please do not share or forward.

Watcher 1.0

- Notifications based on content in the cluster
- Components of an Watch
  - Query
    - Full Elasticsearch query language
  - Schedule
    - Simple scheduler
  - Condition
    - Simple count or advanced script
  - Actions
    - Email, Webhook, Indexing
  - Alert History
    - Enable nested alerts
    - Visualize history in Kibana

Use-Cases

- Monitor Elasticsearch together with Marvel
- Log analysis: trigger notification when errors, anomalies, or specific values are detected in your data
- SIEM: detect and respond to attacks
- SLA monitoring: identify deviations in response time
- Push notifications via email or integrate w/ 3rd party systems

Elastic proprietary and confidential.
Please do not share or forward.

开源技术社区

技术交流&经验分享

- 源码:
  - http://github.com/elastic
- 英文社区:
  - http://discuss.elastic.co
- 中文社区:
  - http://elasticsearch.cn
- QQ群:
  - 19 06 05 846

Elastic proprietary and confidential.
Please do not share or forward.

elastic

Questions?