# From ELK to the Elastic Stack

Medcl Zeng
Developer Evangelist

# **ELK** Stack


APIs

Plugins

Visualization

# Along Came **Beats**

**ELKB**

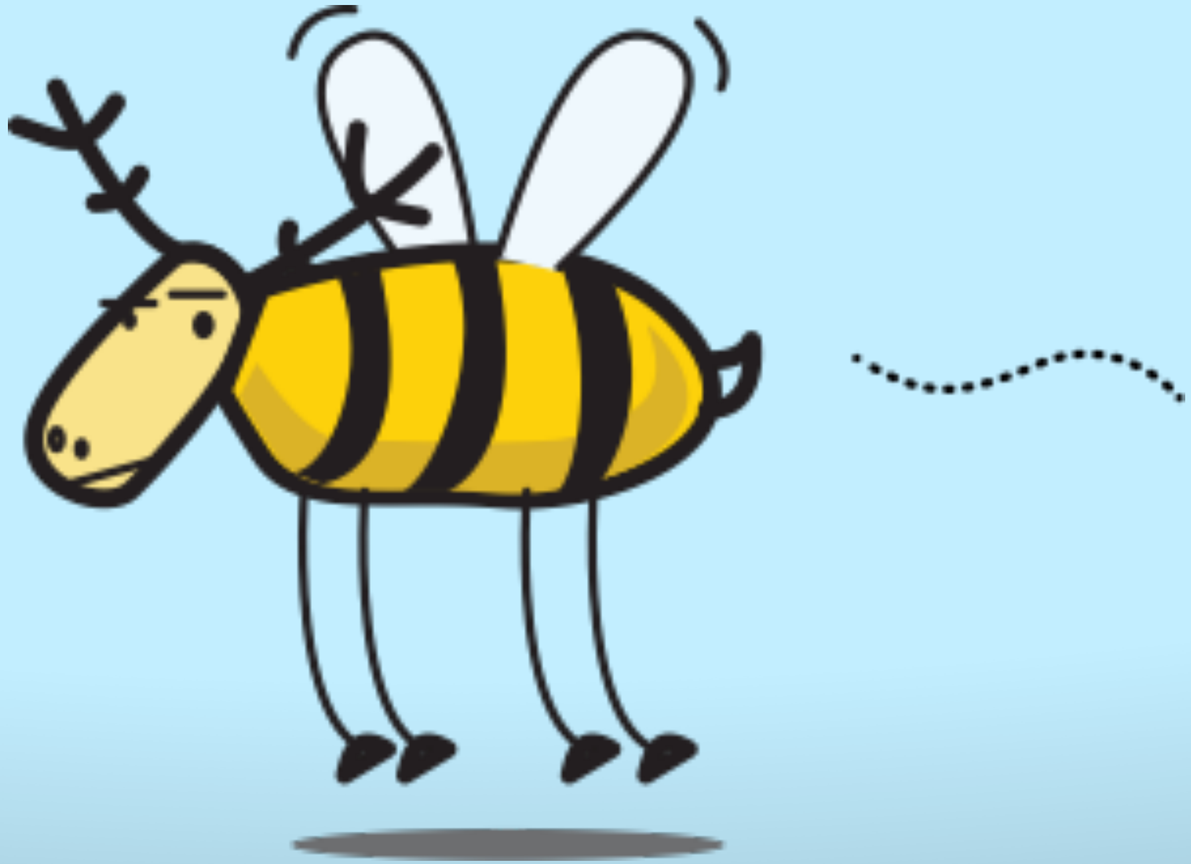

# Elastic Stack

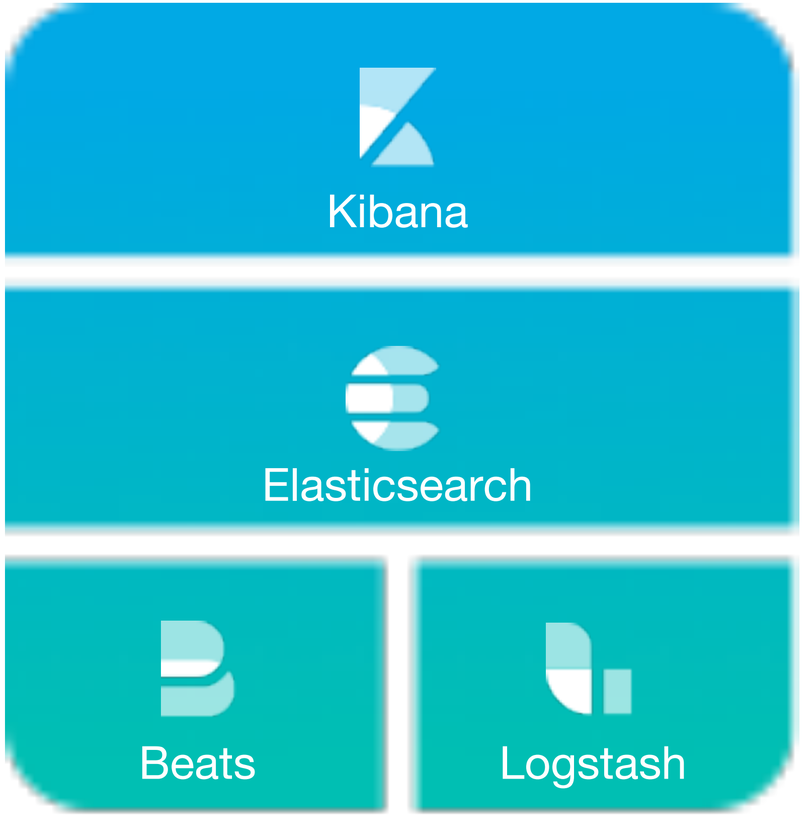

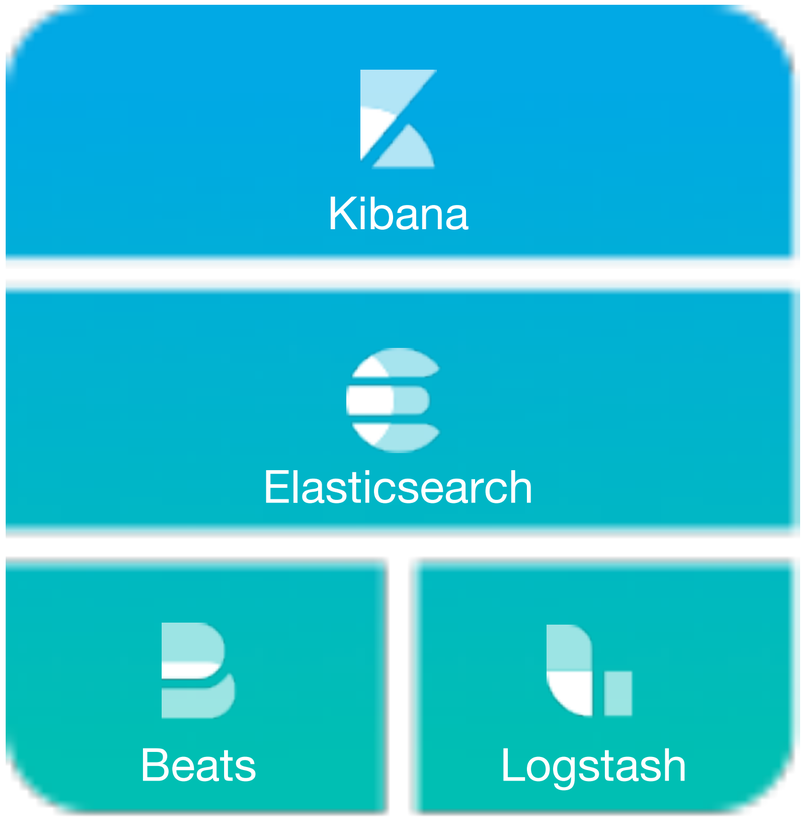
Kibana
Elasticsearch
Beats
Logstash

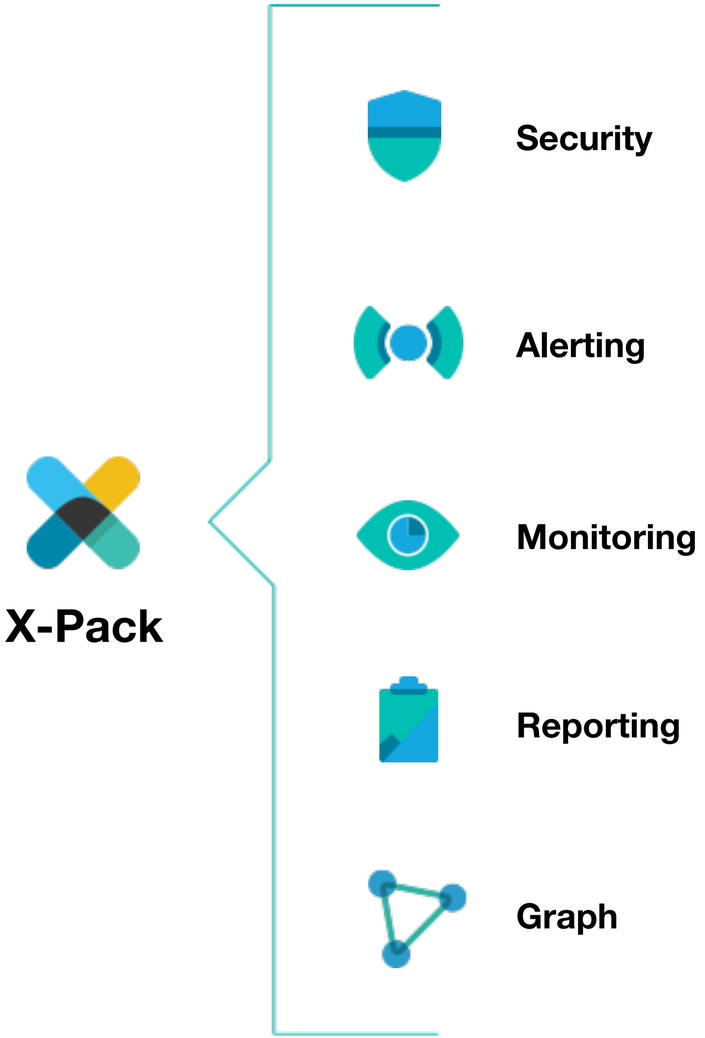
X-Pack
Security
Alerting
Monitoring
Reporting
Graph

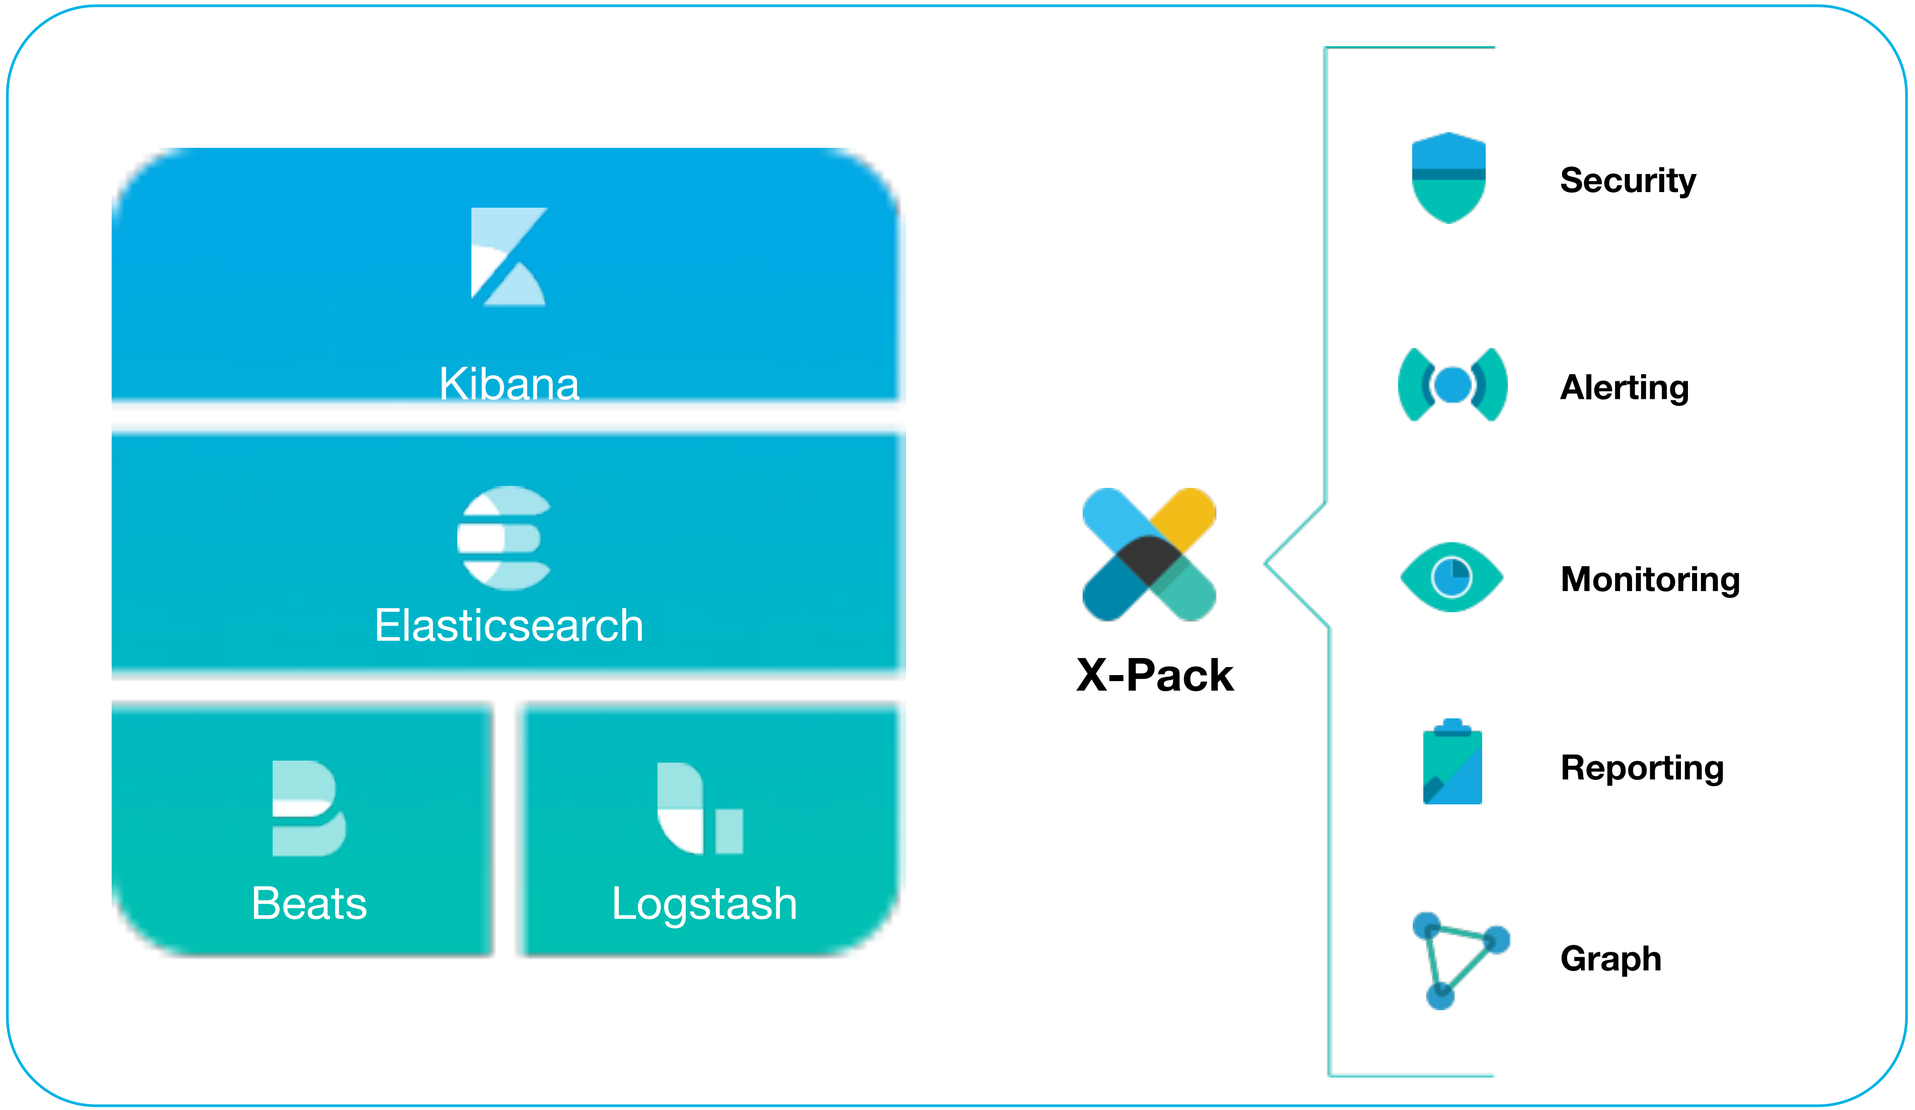
Kibana
Elasticsearch
Beats
Logstash
X-Pack
Security
Alerting
Monitoring
Reporting
Graph
Elastic Cloud

# It’s complicated

| | | | | |
|---|---|---|---|---|
| es | Nov 5, 2014 **1.4** | May 23, 2015 **1.5** | Jun 9, 2015 **1.6** | Jul 16, 2015 **1.7** |
| kibana | Feb 19, 2015 **4.0** | | Jun 10, 2015 **4.1** | |
| ls | May 14th, 2015 **1.5** | | | |
| beats | May 27th, 2015 **1.0 Beta 1** | July 13th, 2015 **1.0 Beta 2** | Sept 4 th, 2015 **1.0 Beta 3** | |

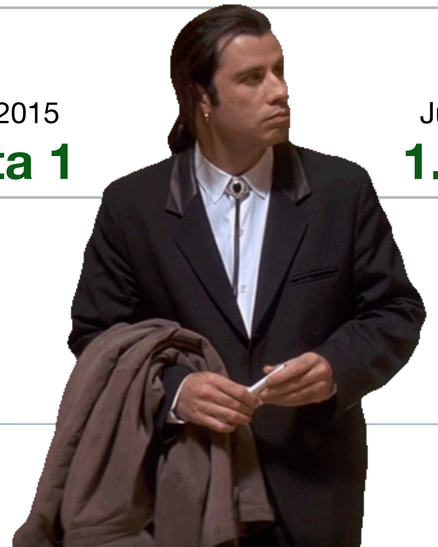

**5.0 is here.**
All new versions.
All aligned.

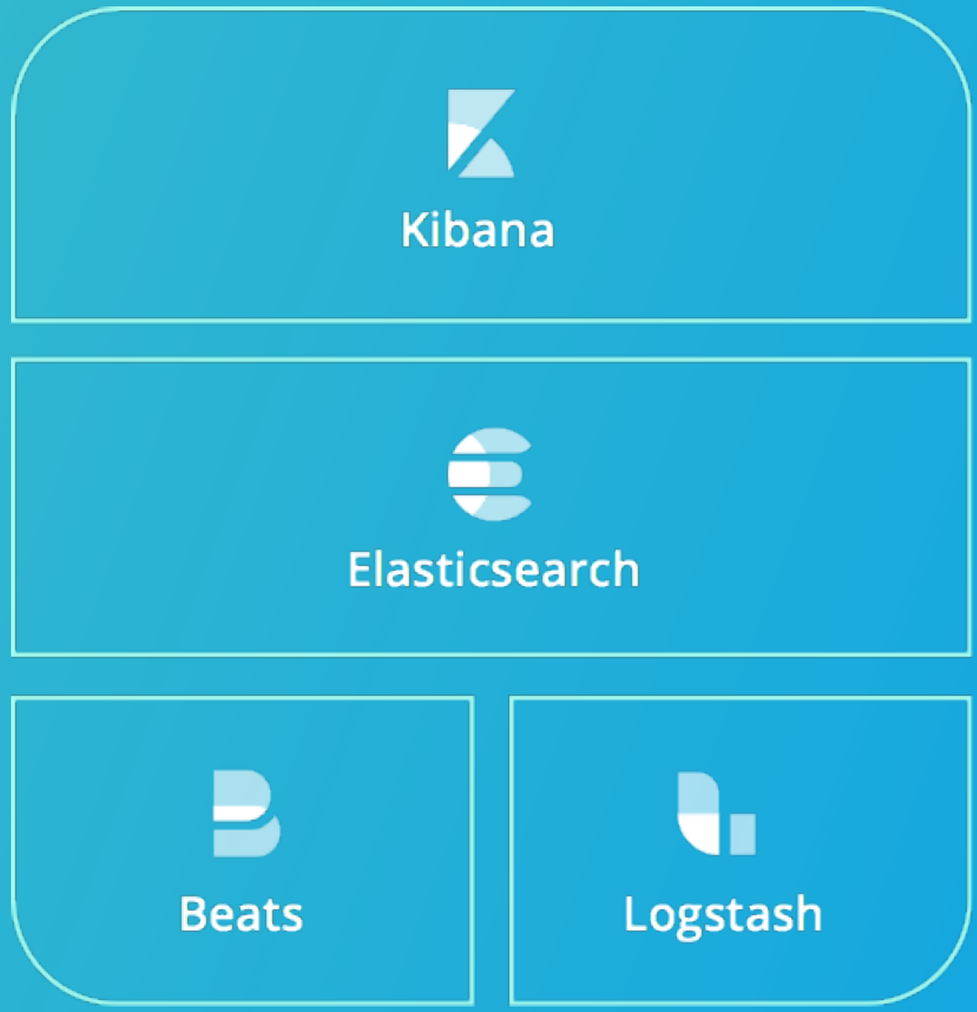

# Working beautifully together

# Elastic Cloud as a Product

## In **ANY** cloud …

amazon web services™

Google Cloud Platform

Microsoft Azure

IBM Cloud

Available in AWS today

## In **YOUR** cloud …

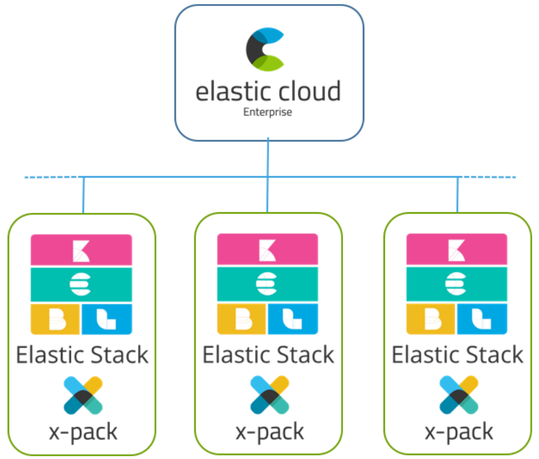


Many clusters / use cases
Single use case, as a service

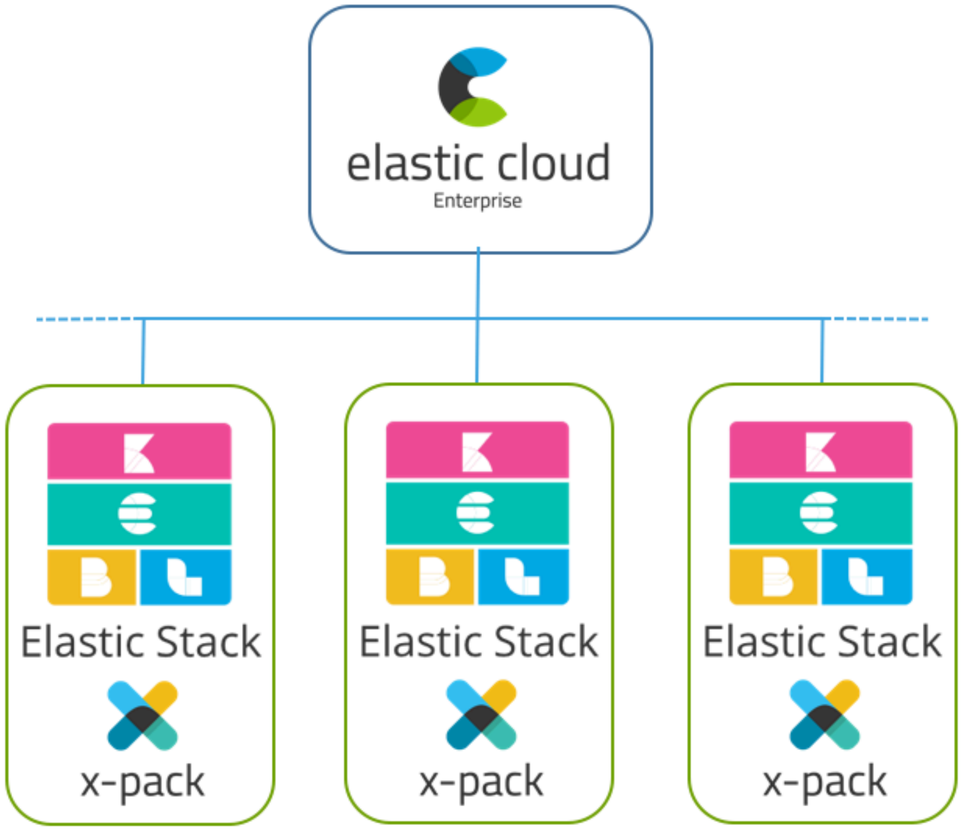


Provisioning, orchestration, and management of multiple Elastic Stacks

Same technical foundation as the Elastic Cloud service

Expected GA Q1 2017

# Welcome Prelert

## Behavioral analytics and unsupervised machine learning

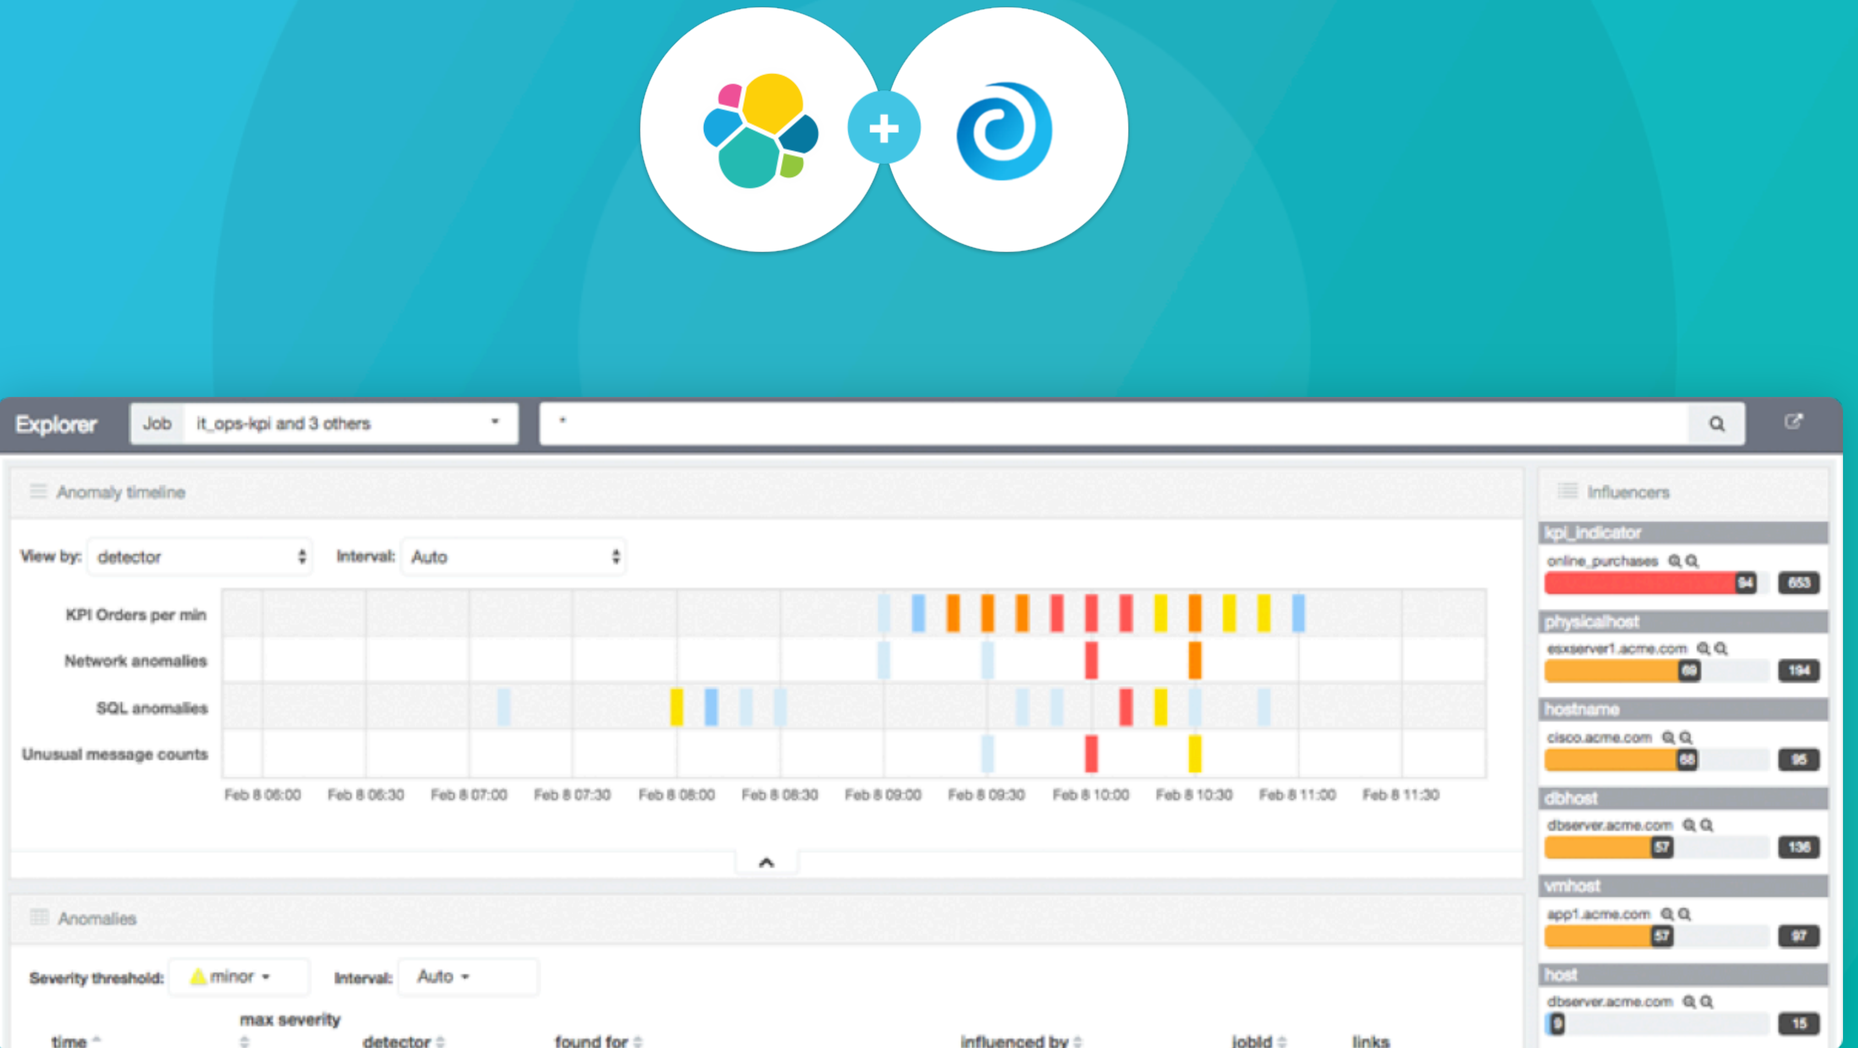

# Elasticsearch 5.0:
# What You Need to Know

# Better support for Numb3rs

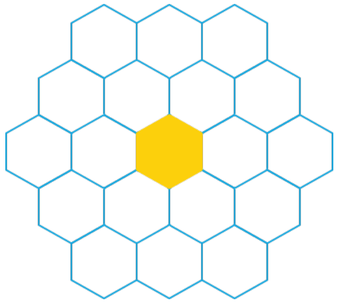

## Faster & reduced memory/disk for many use cases

- BKD Trees
- Lower heap usage
- IPv6 Support

```
PUT iptest
{
  "mappings": {
    "my_type": {
      "properties": {
        "ip_addr": {
          "type": "ip"
        }
      }
    }
  }
}

PUT iptest/my_type/1
{
  "ip_addr": "2001:db8:85a3:8d3:1319:8a2e:370:7348"
}

GET iptest/_search
{
  "query": {
    "term": {
      "ip_addr": "2001:db8::/48"
      }
    }
  }
}
```

```
{
  "took": 1,
  "timed_out": false,
  "_shards": {
    "total": 5,
    "successful": 5,
    "failed": 0
  },
  "hits": {
    "total": 0,
    "max_score": null,
    "hits": []
  }
}
```

# Better support for Numb3rs

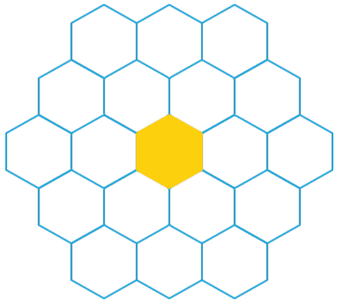

## Faster & reduced memory/disk for many use cases

Scaled / Half float

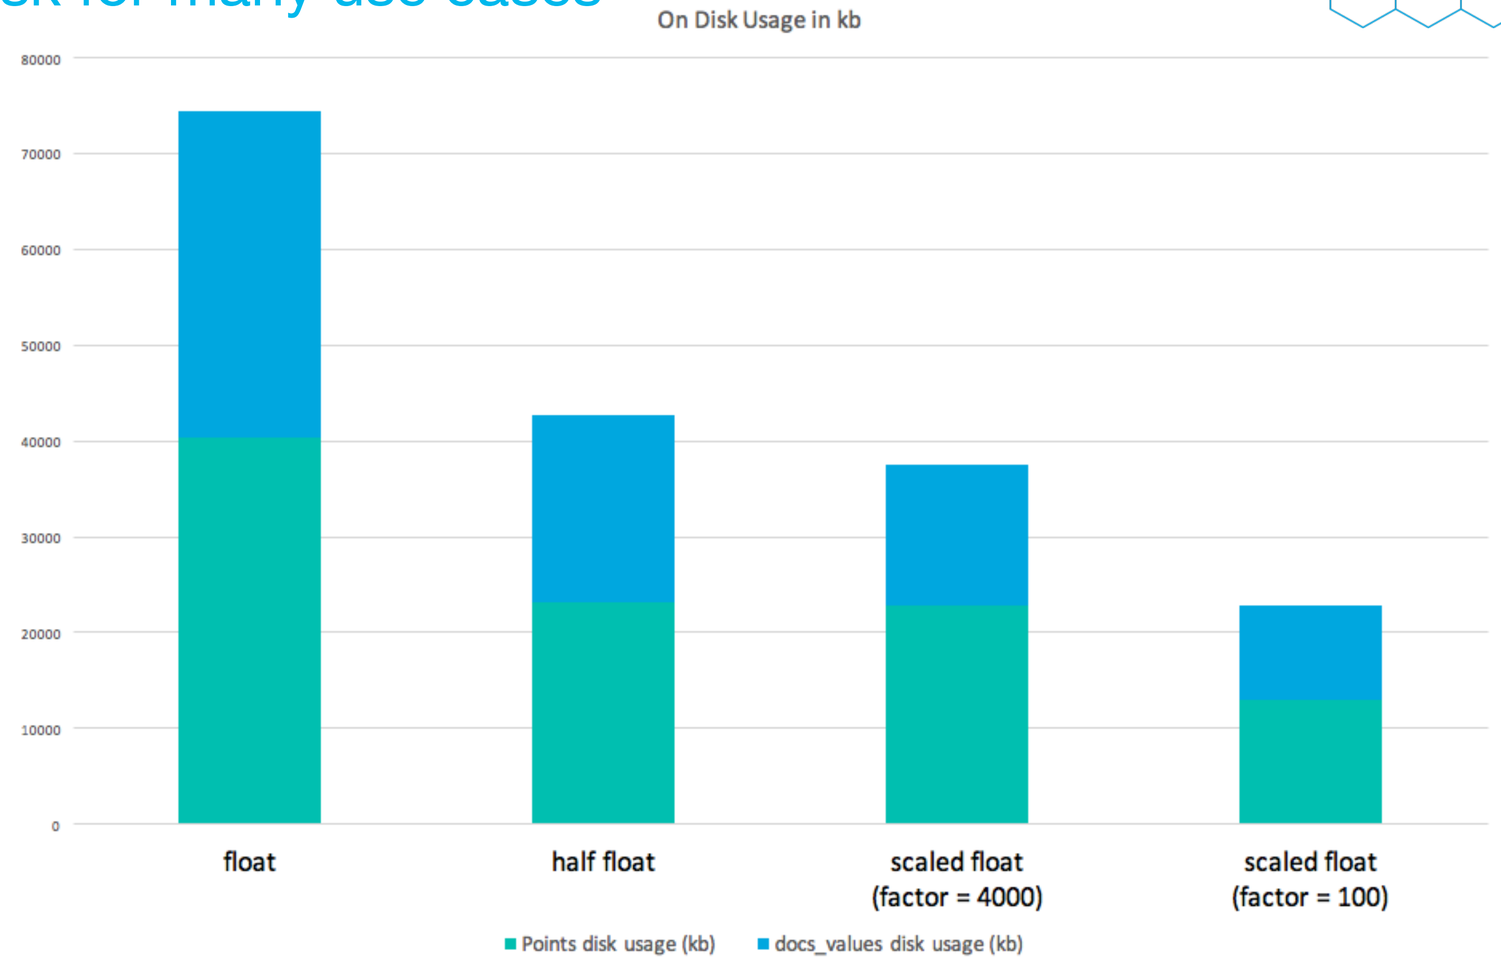

# Improved Indexing Time Performance

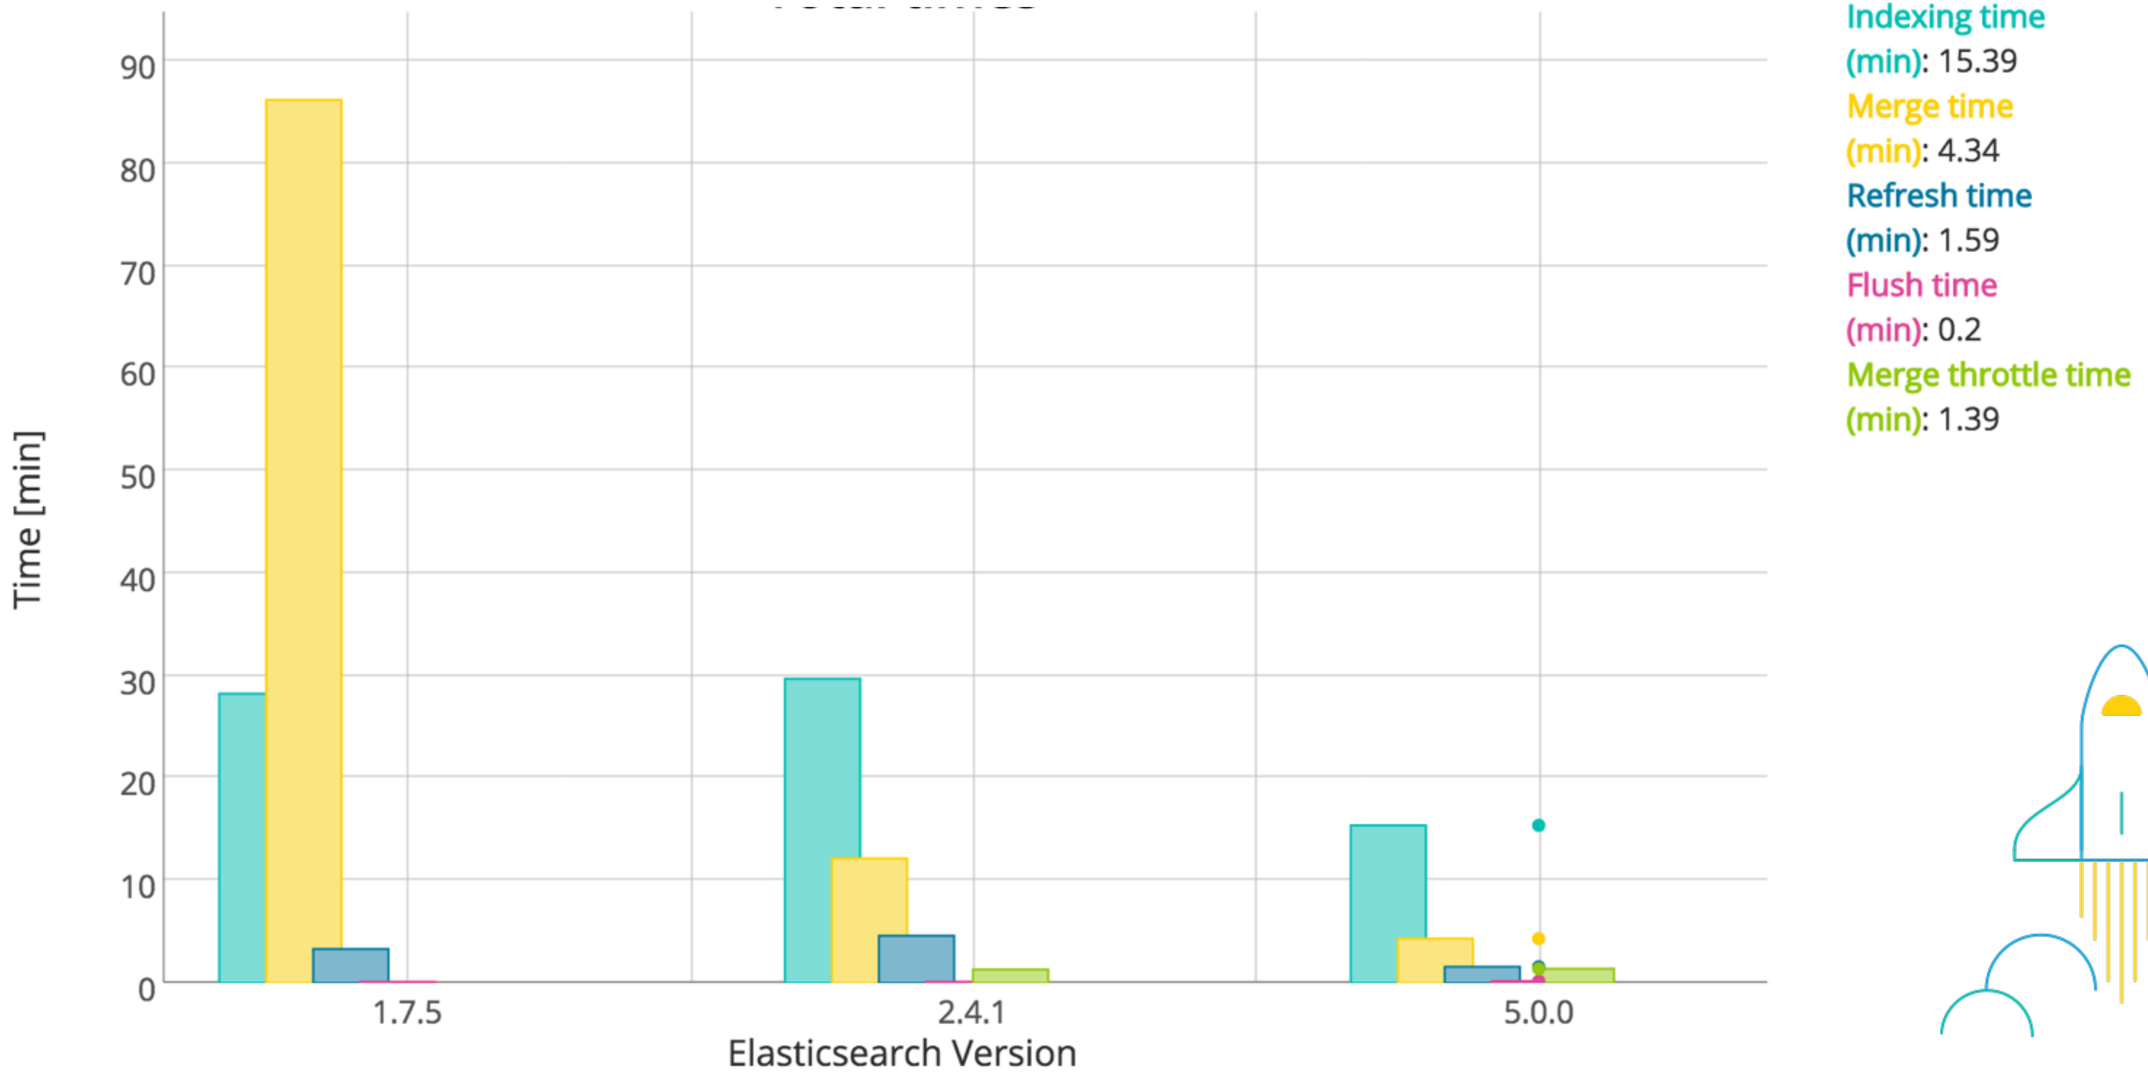

# Fast, Safe Scripting Language

## Say “Heya” to Painless

- Secure and production-safe
- Significantly faster than Groovy
- Familiar syntax
- Can be used in various places:
  - Ingest node pipeline, function scoring, scripted result filtering, watch conditions, and more

# Simplified Architecture

- Automatic time-series index management
  - **Rollover APIs**

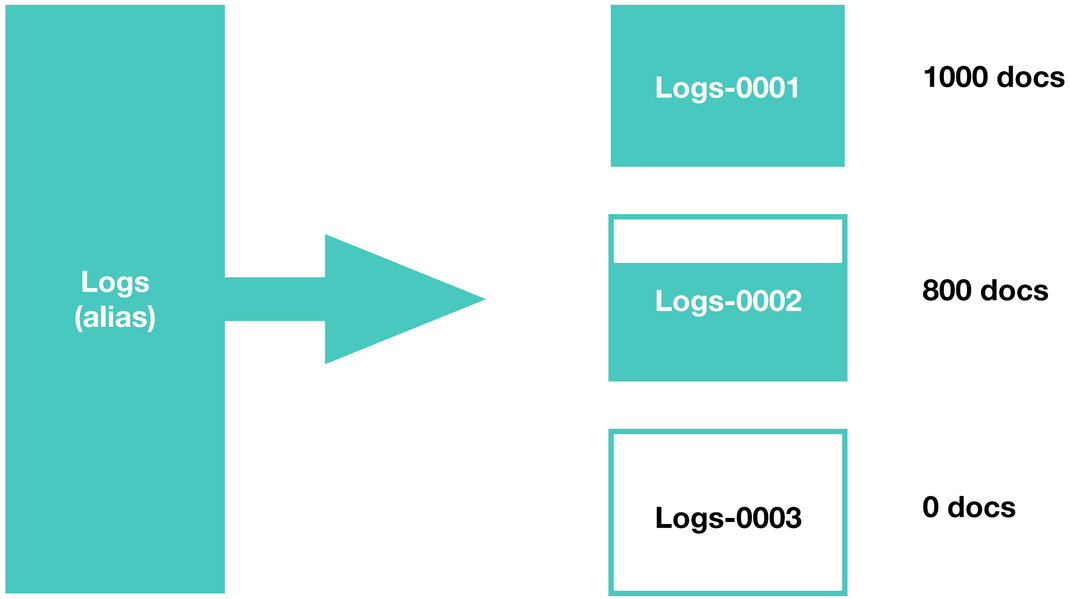

# Simplified Architecture

- Automatic time-series index management
  - **Shrink APIs**

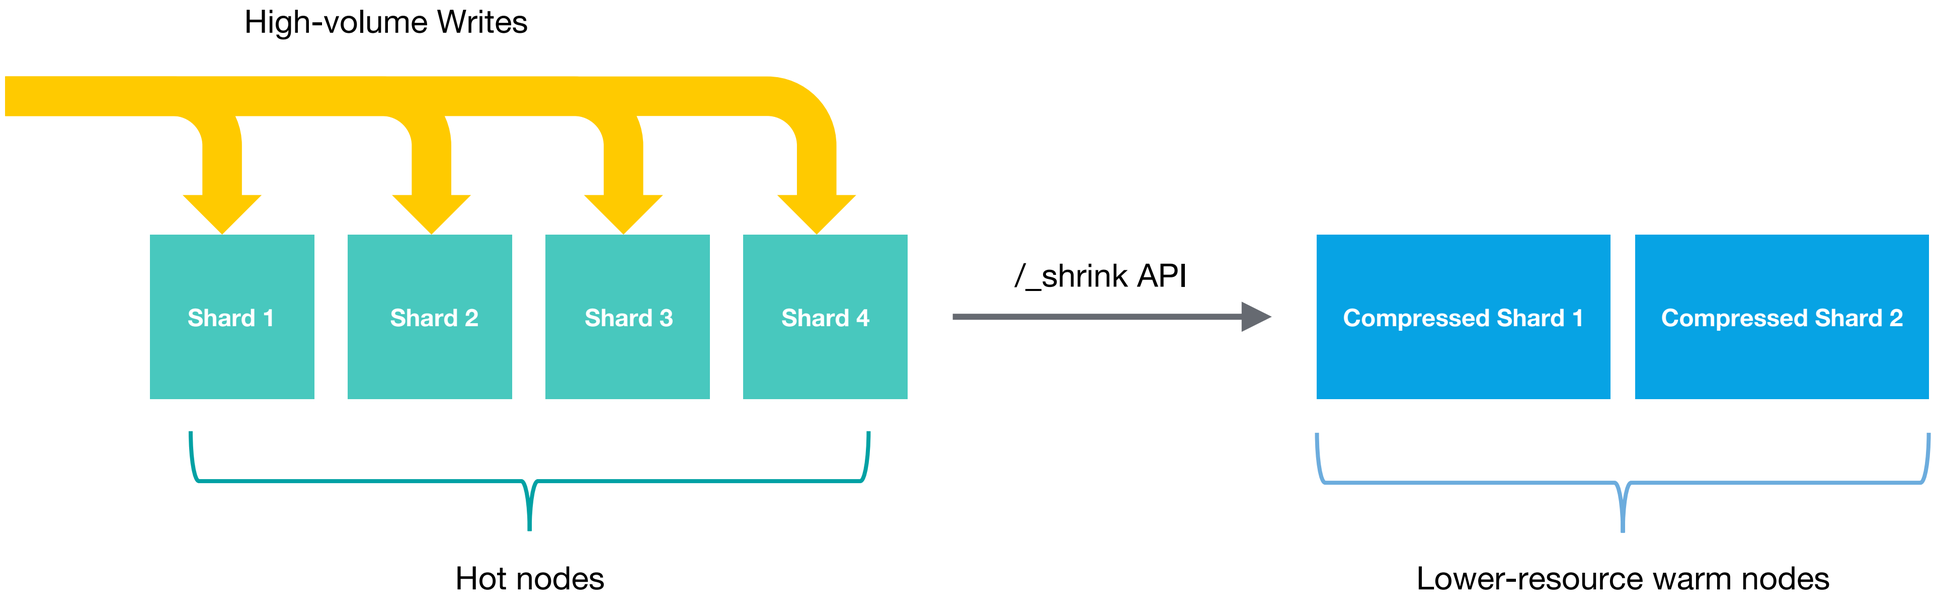

# Simplified Architecture

- Simplified experience for interactive pages
  - **Wait-for-Refresh**
- Simplified getting started experience
  - **Ingest Node: More to come on this today**

# Resiliency and Safety Improvements

- We saw some common problems when getting started or new users on a multi-tenant environment
  - **Bootstrap checks**
  - **Safeguards**

# Resiliency and Safety Improvements

- Memory management is important
  - **Keyword type**
  - **Circuit breakers**

# Resiliency and Safety Improvements

- Understanding and preventing a terrible Friday afternoon
  - **2 phase cluster state commit**
  - **safe primary relocations**

# Faster, more normalized DSL

- Completion Suggester v2
- Percolation is now a normal query
- Profile API expansion to include aggregations and not just queries

# Beyond 5.0

- Higher timestamp resolution (great for logging use cases)
- More improvements on resiliency
- Build on BKD: range fields, geo
- Increased performance for append-only time series use cases
- Native RESTful Java client
- We released 5.1.1 this week

# Ingest: Beats & Logstash 5.0

# Ingest data from any source, in any format

Beats

Logstash

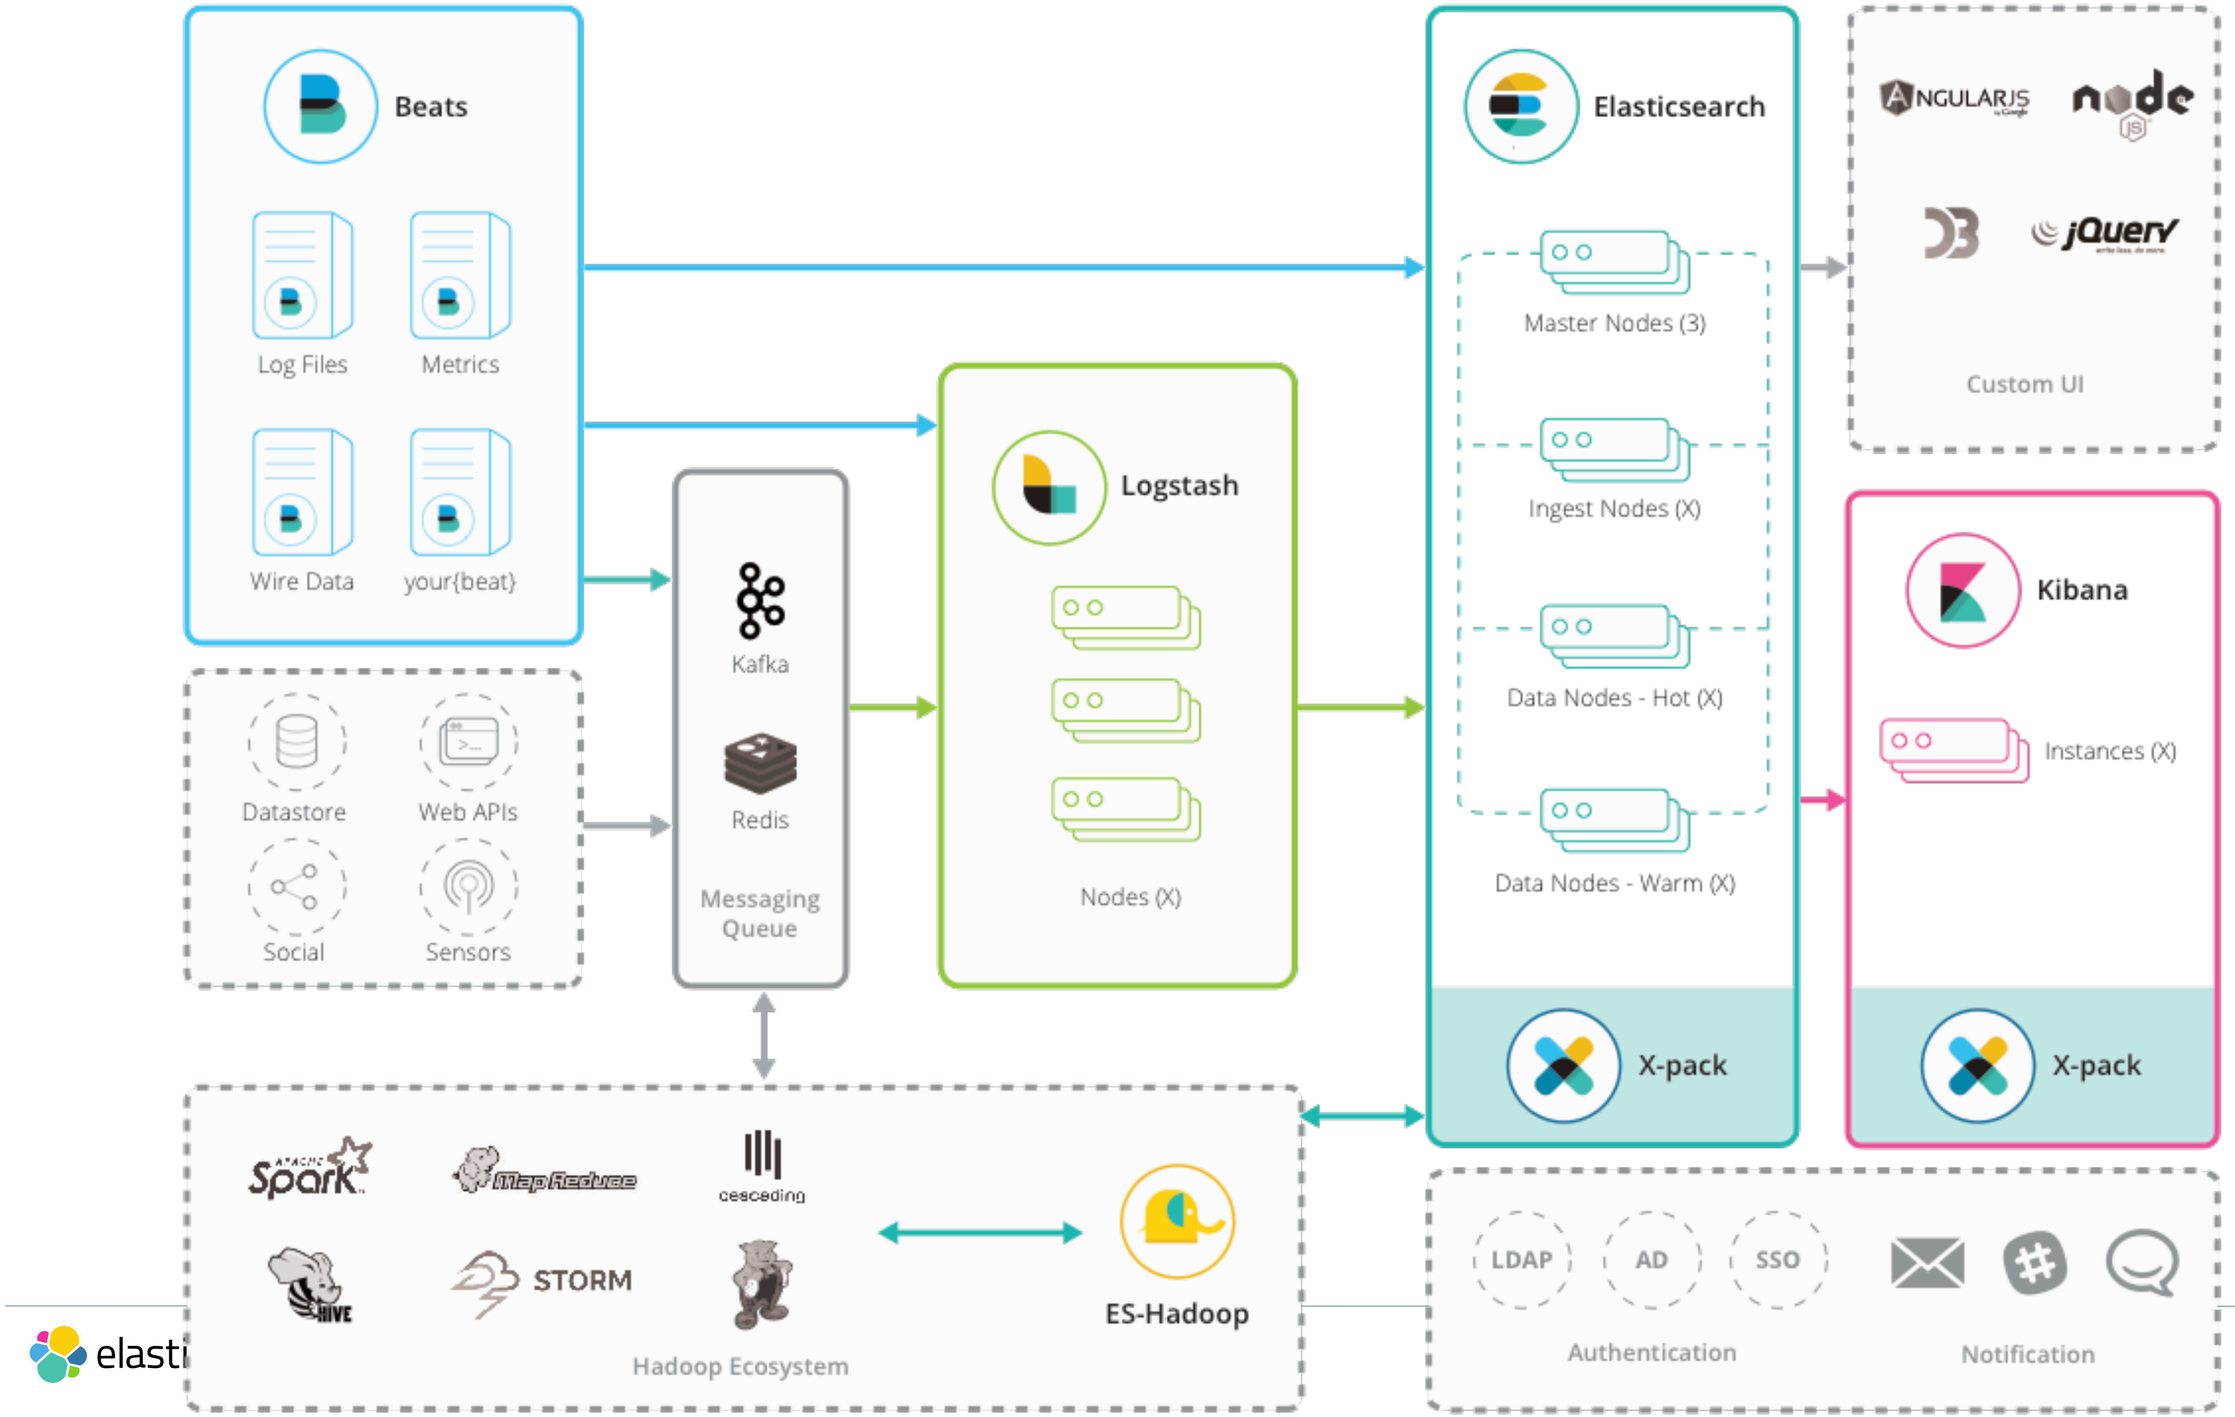
Beats
Log Files
Metrics
Wire Data
your{beat}
Datastore
Web APIs
Social
Sensors
Kafka
Redis
Messaging Queue
Logstash
Nodes (X)
Elasticsearch
Master Nodes (3)
Ingest Nodes (X)
Data Nodes - Hot (X)
Data Nodes - Warm (X)
X-pack
Custom UI
Kibana
Instances (X)
X-pack
Spark
MapReduce
cascading
HIVE
STORM
ES-Hadoop
Hadoop Ecosystem
LDAP
AD
SSO
Authentication
Notification

# Say Heya to Ingest Node

Process incoming data directly in Elasticsearch

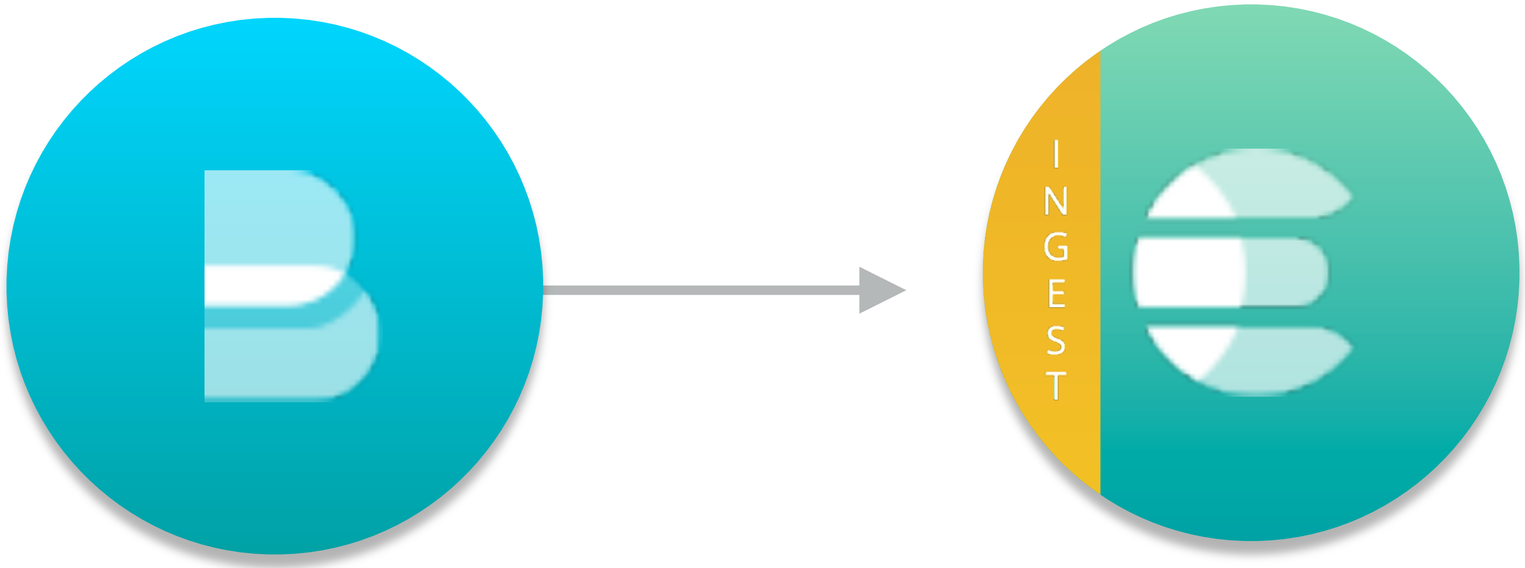

# New in 5.0


Streamline network & storage

**Beats Processors**

Count and bytes on the TCP/IP layer not application layer

**Packetbeat**

No more double Logstash

**Kafka output for Beats**

# New in 5.0


## Metricbeat is here

~~Topbeat~~

# Logstash: Goodbye Black Box!

logstash:9600/_node

Node Info
Node Stats
Plugins
Hot Threads

**Monitoring API**

Debug active pipelines
with new logging API

Component level
logging granularity

**Log4j2 Internal Logging**

# Logstash: Performance++

20%+ increase in overall pipeline performance

**New Java Event**

50% performance boost ingesting from Beats

**Beats Input Java Rewrite**

# Logstash: Plugin Features

Kinesis Input
Protobuf Codec
Dissect Filter

IPv6 Support with GeoIP2

**New Plugins**

Kafka 0.10 Support

Basic Auth & SSL/TLS

**Kafka Support++**

Developers can generate new plugins in seconds

**Plugin Generator**

# Elasticsearch-Hadoop 5.0

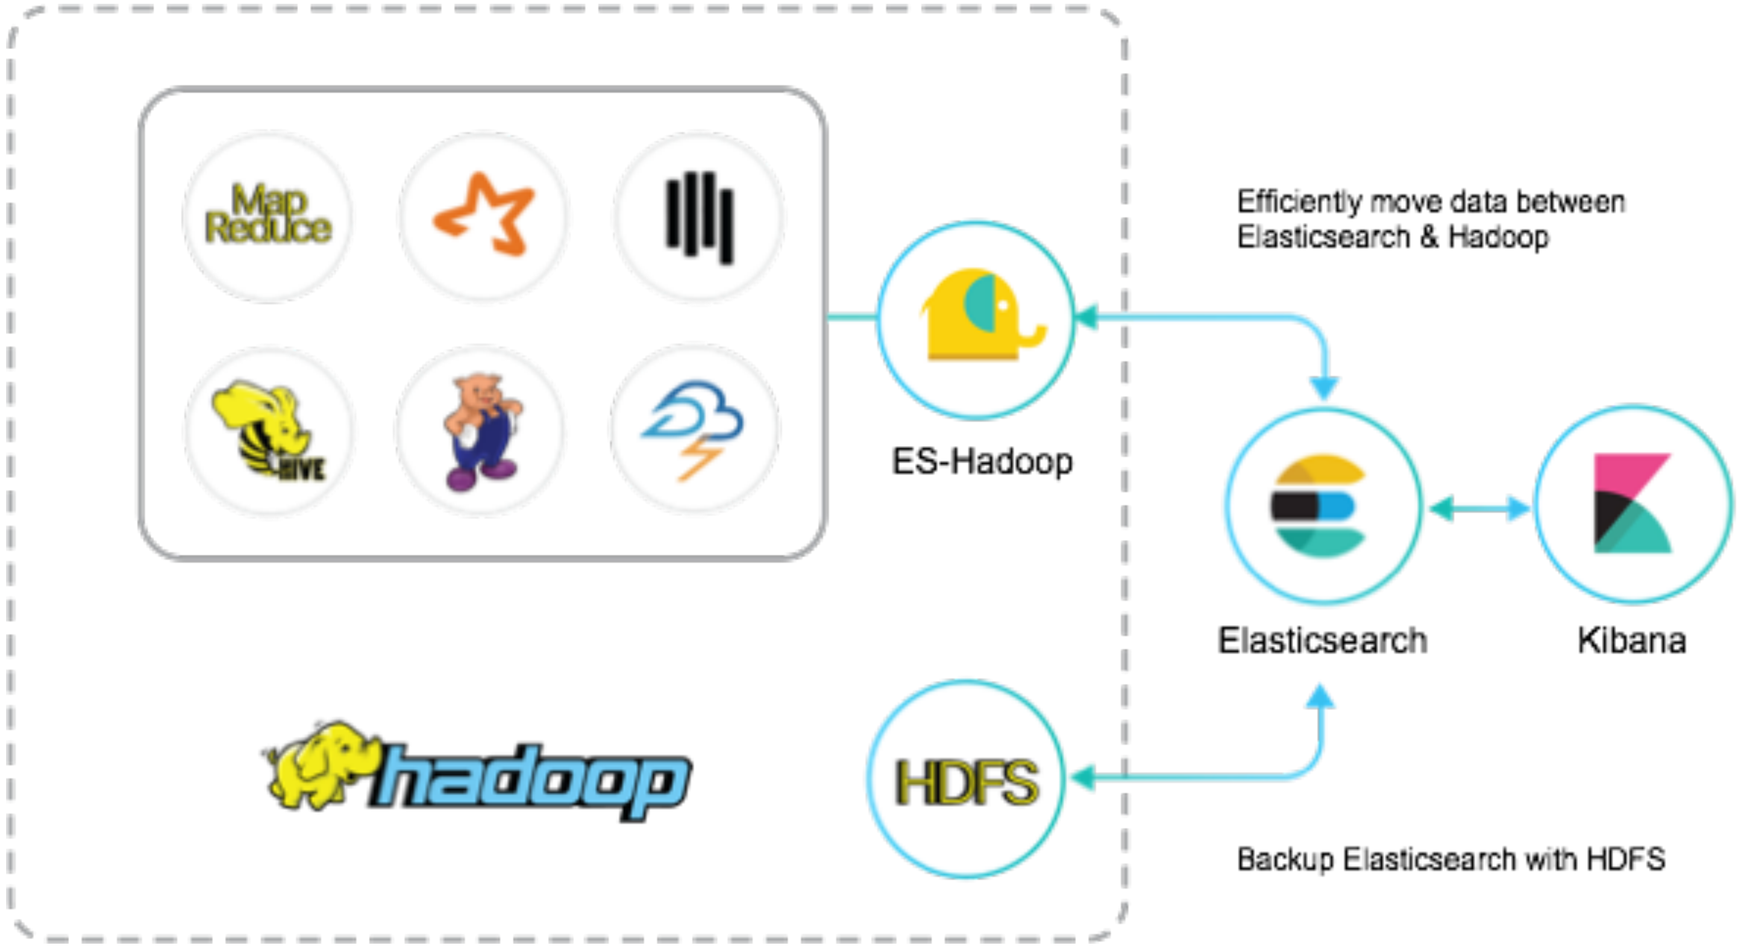


Spark 2.0 & Better Streaming Support

Ingest Node Pipeline Integration

Elasticsearch 5.0 Parallel Reader

# Beyond 5.0 (Beats)

- More modules in Metricbeat
- More Beats
- Even easier getting started experience
- Centralized configuration & monitoring

# Beyond 5.0 (Logstash)

- Logstash persistence (disk-based queuing)
- Monitoring UI & centralized configuration
- Multiple pipelines, one JVM
- Error event routing