

Elastic Platform Technical Deep-Dive

May 27, 2015

Alan Hardy
Solutions Architect

Pere Urbon-Bayes
Software Engineer

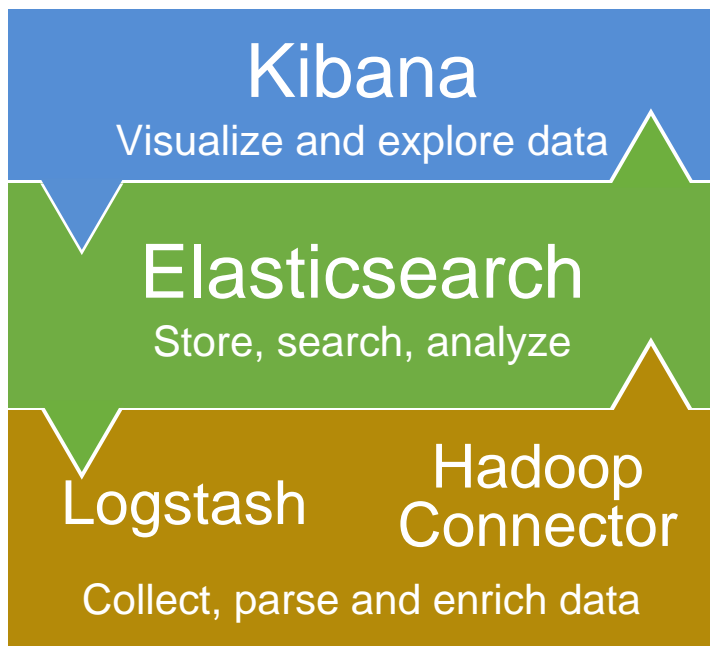


Agenda

- Elastic Platform
- Deployment Architecture
- Logstash Deep-Dive
- Elasticsearch Deep-Dive
- Kibana Deep-Dive & Demo
- Real-World Use Cases
- Q&A

The Elastic Platform

Open Source Products



Commercial Products



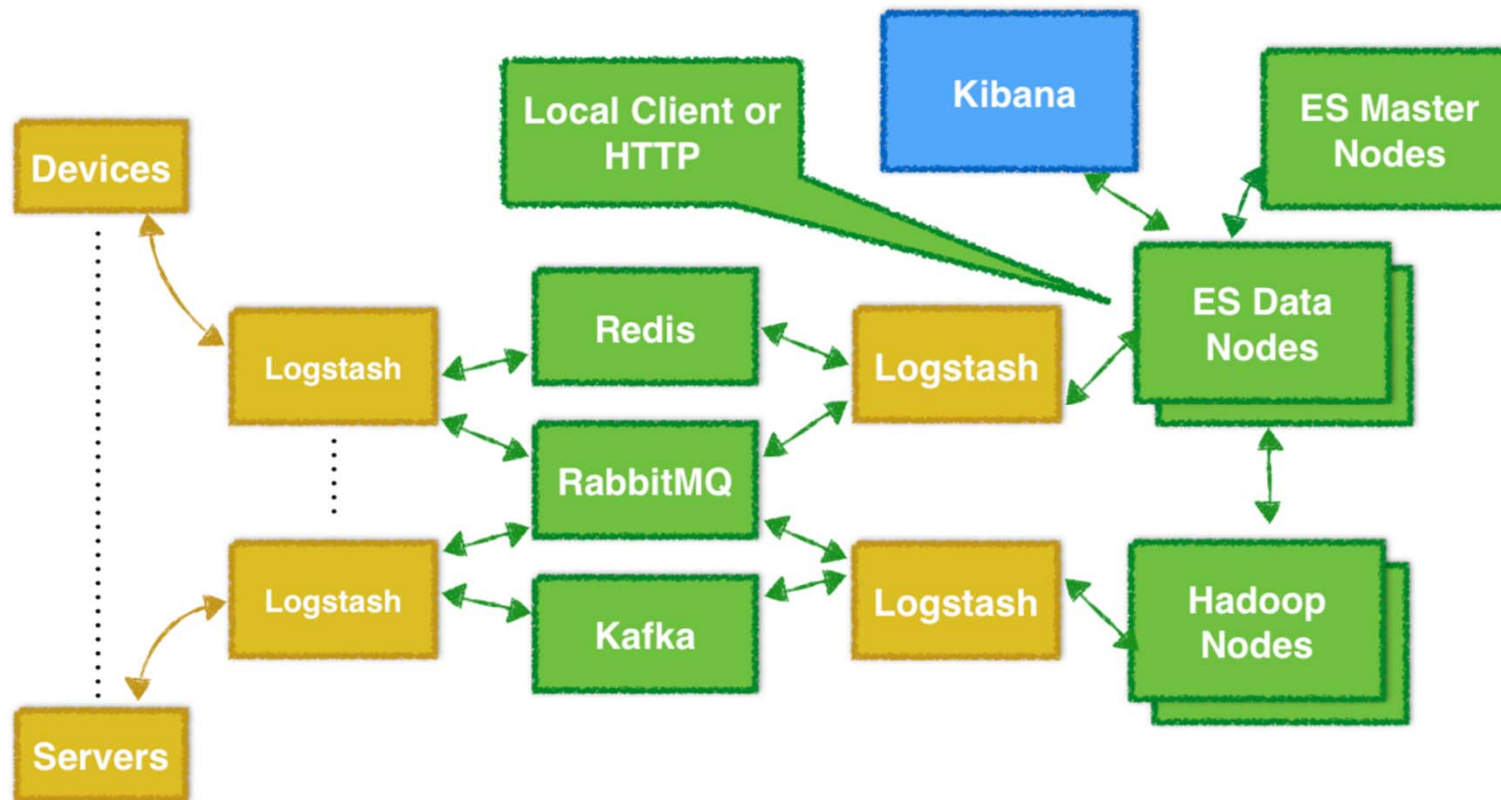
Training

Professional Services

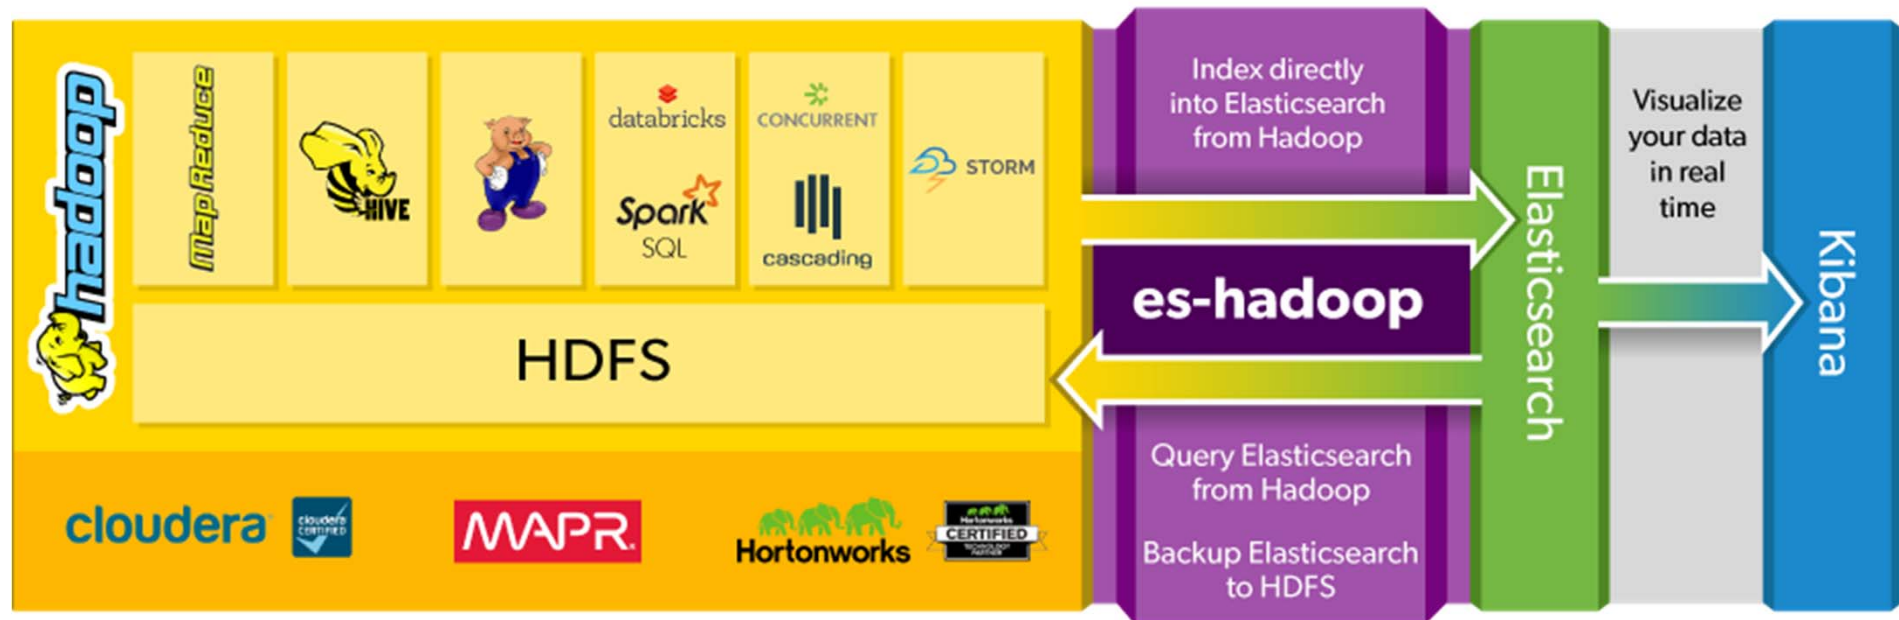
Support Subscriptions

FOR BUILDING SCALABLE, DISTRIBUTED SYSTEMS

Architecture: Logging



Architecture: Elasticsearch Hadoop ecosystem



Logstash Deep-Dive

Pere Urbon-Bayes
Software Engineer



Logstash - The shipper with a moustache



Being on call

Live on call:

Wake up!! it's 3AM.

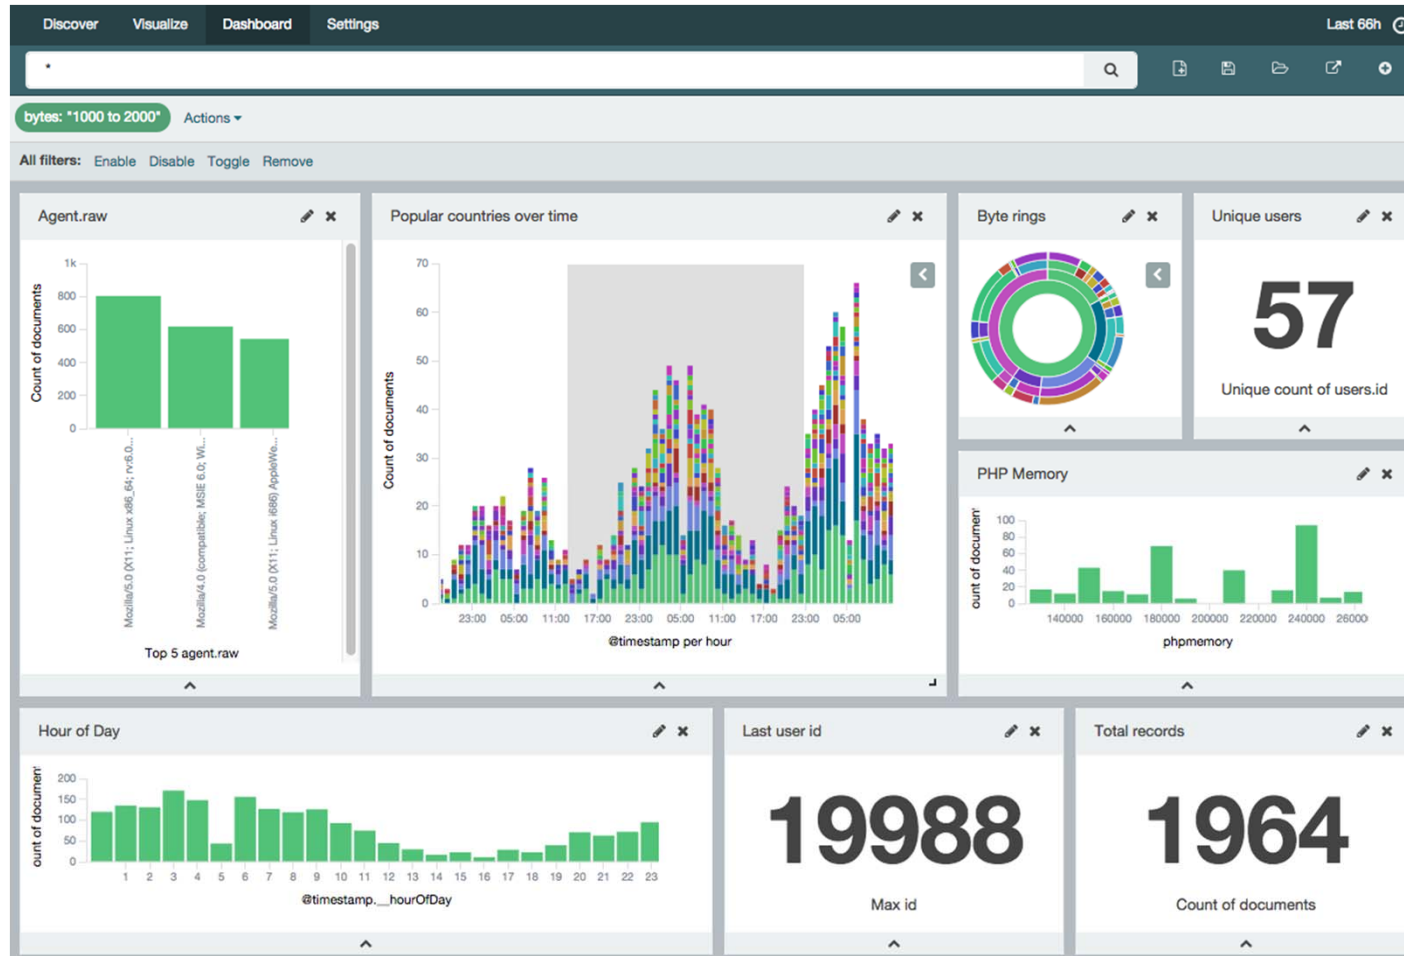




Managing Logs

Logs need to be delivered and stored somewhere,
so we can analyse them easily.

Understanding logs



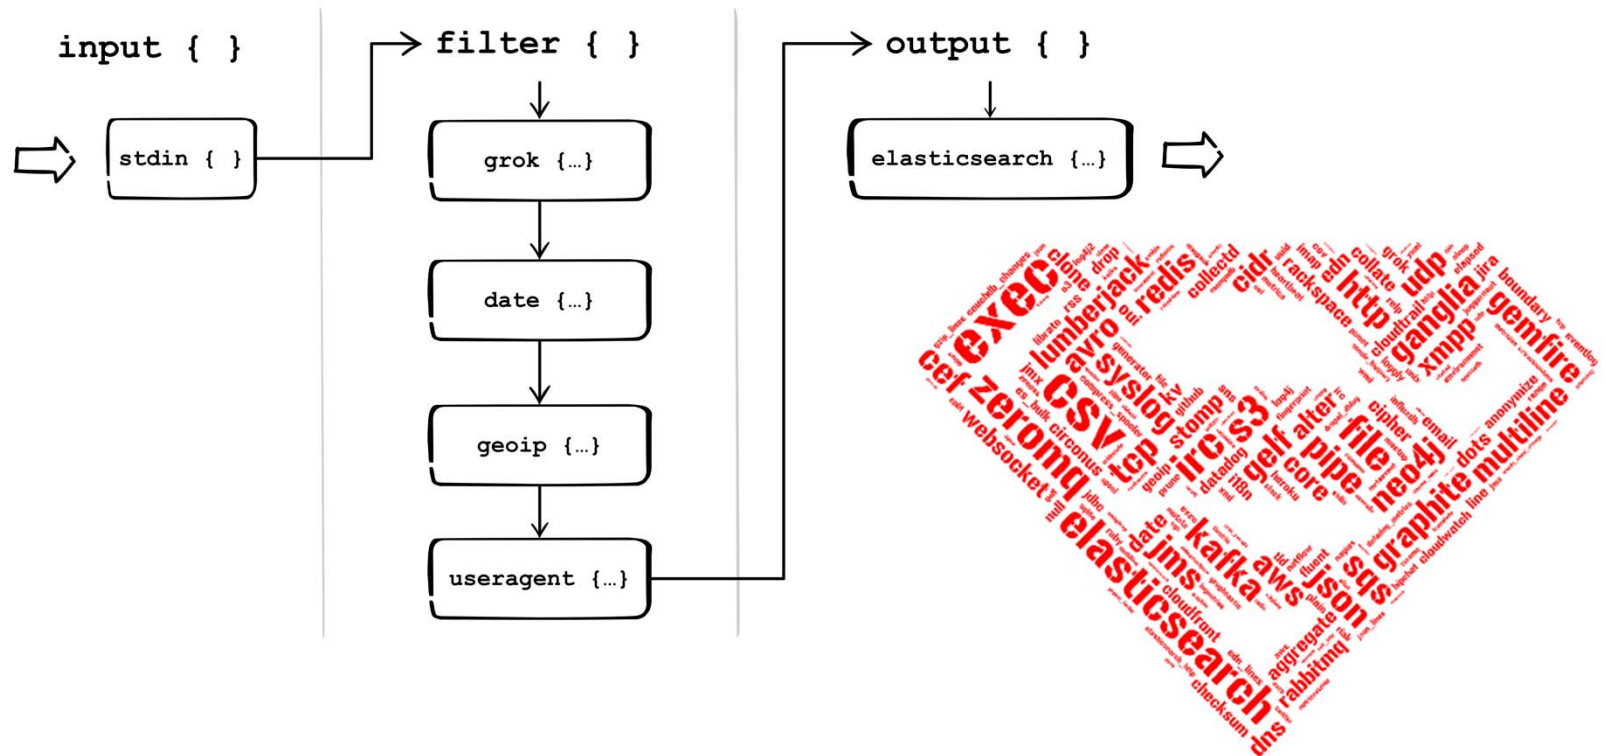
All about Logstash



What logstash can do for you....

- Receiving your events/log data
- Processing and normalising them
- Transporting them to a final destination

How Logstash works



How to configure Logstash

```
input { stdin {} }
```

```
filter {  
  grok {  
    match => {  
      "message" => "%{HTTPDATE:timestamp} %{IP:ip} < %{DATA:msg}>"  
    }  
  }  
}
```

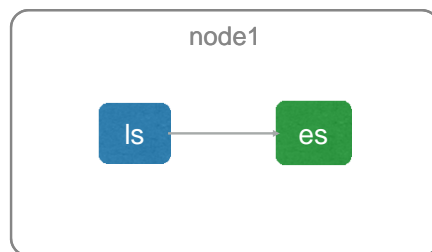
```
output {  
  stdout {  
    codec => rubydebug  
  }  
}
```

```
output {  
  if [action] == "alert" {  
    pagerduty {}  
  }  
}
```

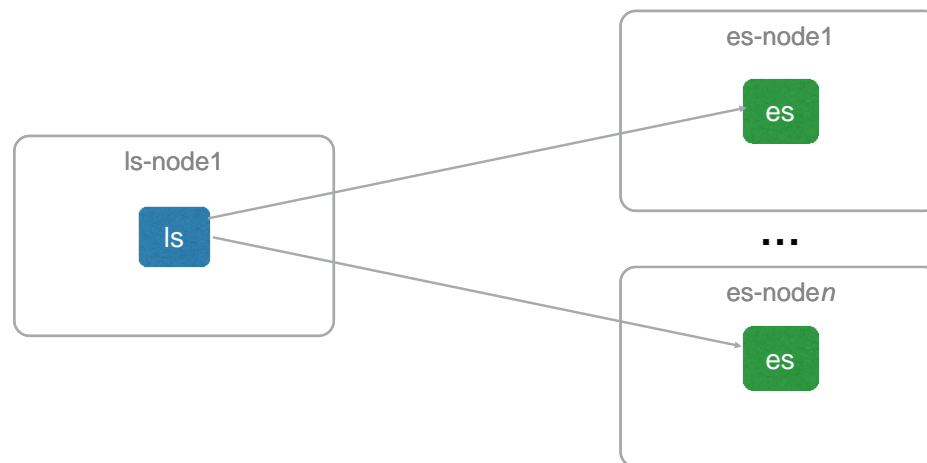

Logstash at scale



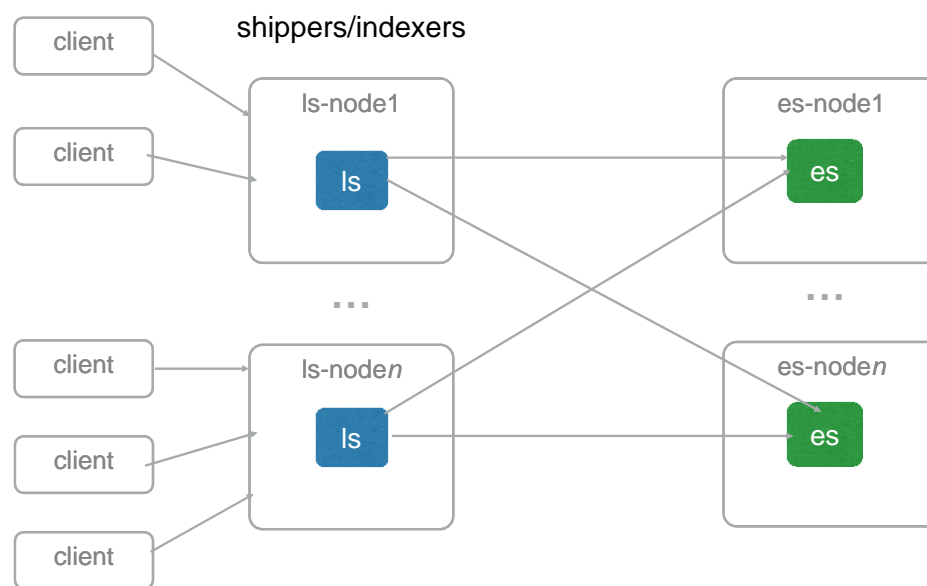
The one node experience



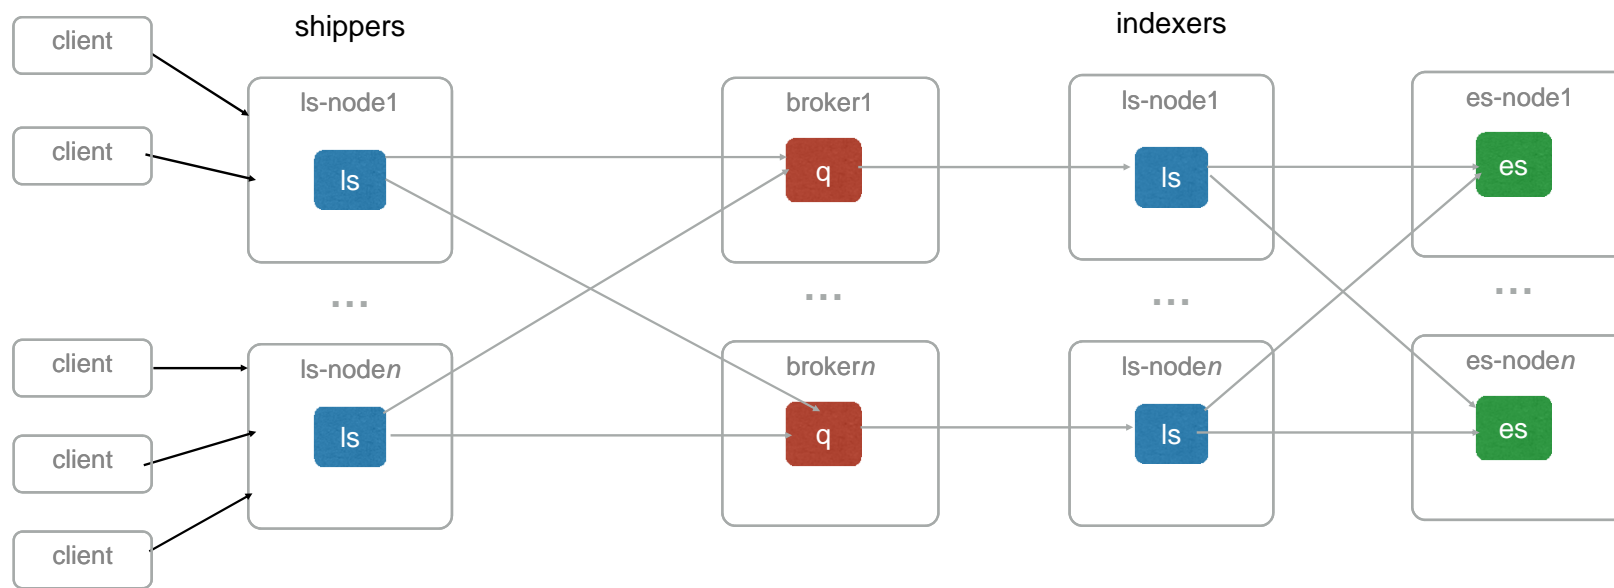
A single LS node



Multiple LS instances



Message Brokers



Baby steps on scaling



redis



In Logstash land: *Redis*



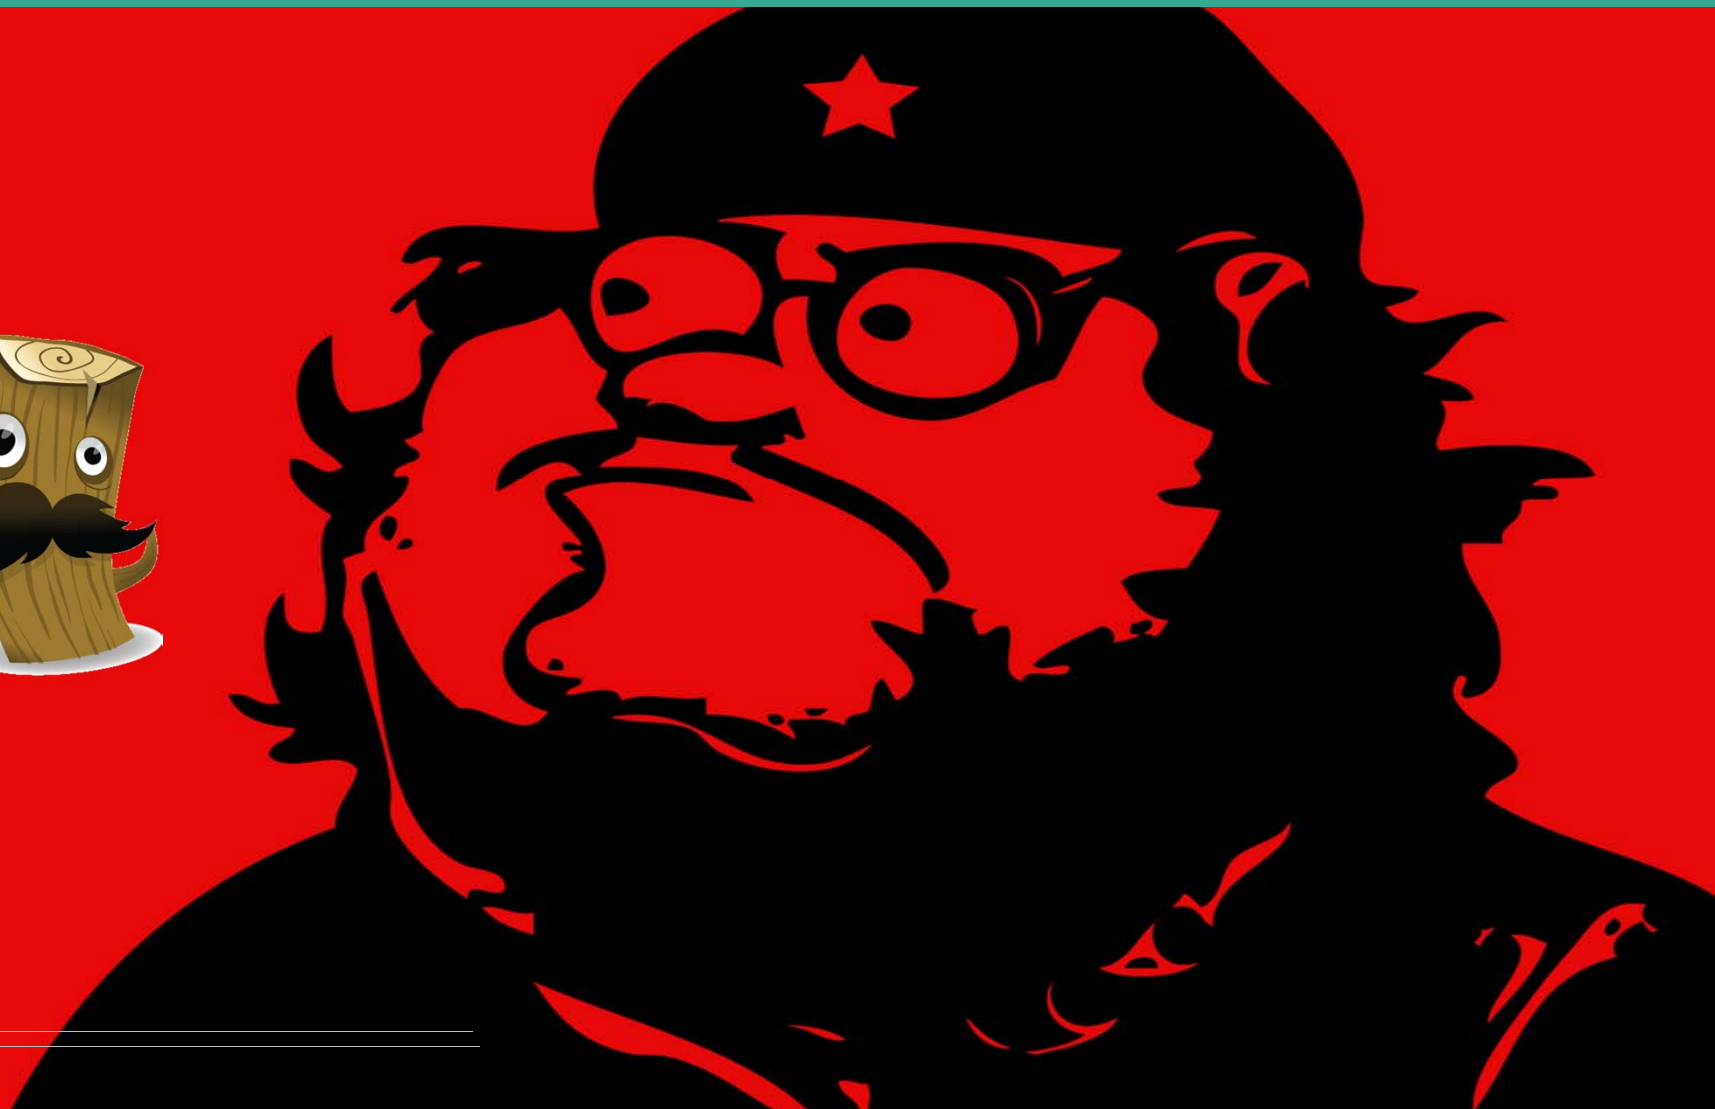
```
input {
  redis {
    name => "default"
    host => "127.0.0.1"
    port => 6379
    db => 0
    timeout => 5
    password => ""
    queue => ""
    key => ""
    data_type => ["list", "channel", "pattern_channel"]
    batch_count => ""
  }
}
```

- # The `name` configuration is used for logging
- # The hostname of your Redis server
- # The port to connect on.
- # The Redis database number.
- # Initial connection timeout in seconds.
- # Password to authenticate with.
- # The name of the Redis queue (deprecated)
- # The name of a Redis list or channel.

**There is an output counterpart
with similar configuration options**

Viva la resistance: *A more resilience way of scaling*

 RabbitMQ



In Logstash Land: *RabbitMQ*



```
input {
  rabbitmq {
    host => ""
    port => 5672
    user => "guest"
    password => "guest"
    vhost => "/"
    ssl => false
    verify_ssl => false
    debug => false
  }
}
```

RabbitMQ server address
RabbitMQ port to connect on
RabbitMQ username
RabbitMQ password
The vhost to use.
Enable or disable SSL
Validate SSL certificate
Enable or disable logging (deprecated)

There is an output counterpart
with similar configuration options

Logstash in the wild



Logstash vs Fluentd vs Flume

	Logstash	Fluentd	Flume
#plugins	180+	440+	25+
Predefined parsing formats	Yes	Yes	No
Variable Buffering	Not yet	Yes	Yes
Persistent queues	Not yet	Yes	Yes
Recipes	Yes	Yes	Not for NG
Configuration reloading	Not yet	Yes	Yes
Conditionals in config	Yes	No	No
Event enrichment	Yes	Yes	No

Logstash vs Fluentd vs Flume

	Logstash	Fluentd	Flume
Monitoring API	Not yet	Yes	No
JMX integration	Partially	No	Yes
Custom reporting	No	No	Yes
<i>Hadoop integration</i>	No	Partially	Yes
OOB Elasticsearch integration	Yes	No	No
RDBMS integration	Partially (Input)	Partially	Partially (Channel)
Lighweitght shipper	Yes	No (himself)	No (himself)

Elasticsearch Deep-Dive

Alan Hardy
Solutions Architect



Elasticsearch



**Store, Search
and Analyze**

Distributed, scalable, and resilient

Designed for scale-out; high availability

Developer friendly

API-first; schemaless, native JSON, client libraries for any language

Real-time Search & Analytics

Real-time aggregations, geospatial, full-text search; query structured and unstructured data

“node”

running instance of elasticsearch

≈ one server

“shard”

holds just a slice of the data

physical worker unit

lives on one node

(a full Lucene search engine)

“index”

logical namespace

points to one or more shards

`shard = hash(_id) % no_of_shards`

Terminology

many segments → **one shard**



many shards → **one index**



scale out, not up



Create an Index

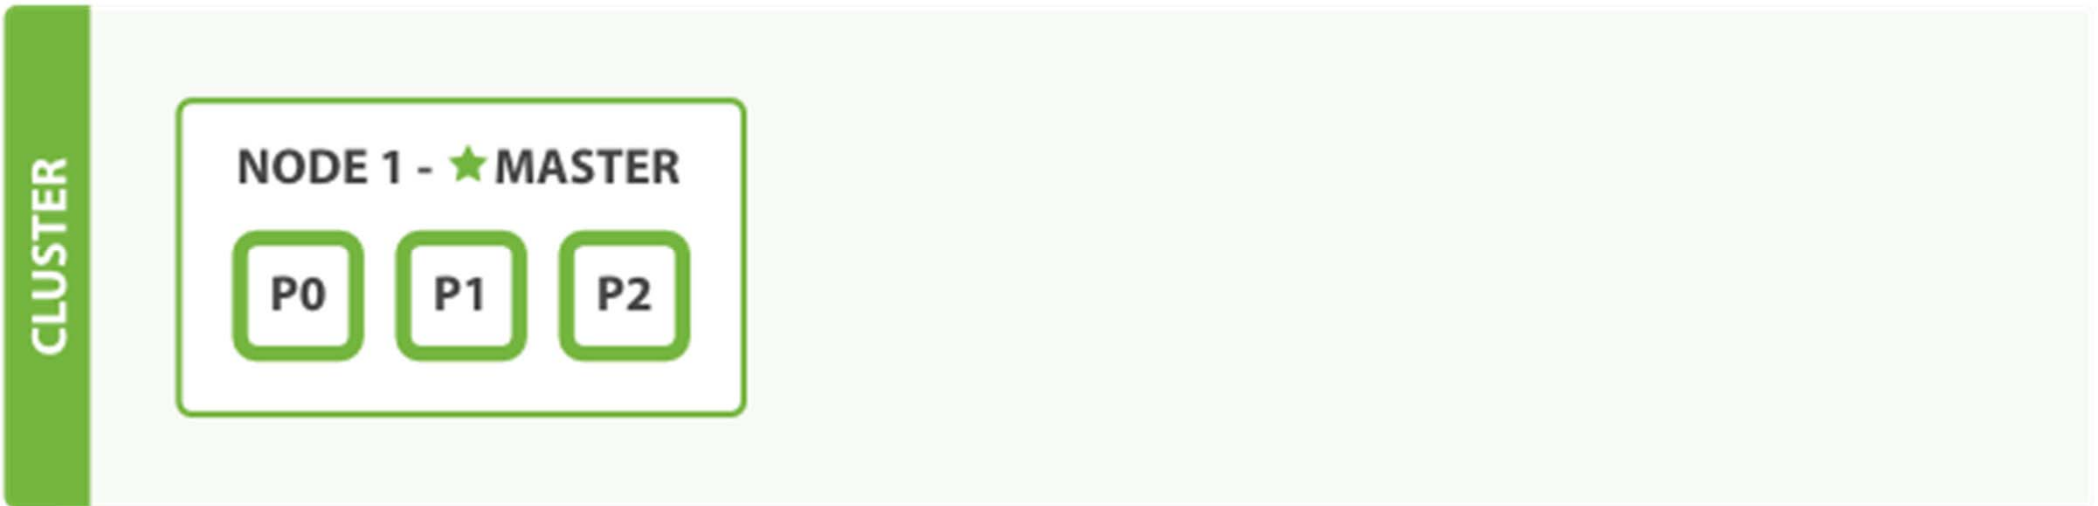
To add data we need an **index** (one or more shards)

A shard can be either a **primary** shard or a **replica** shard

A document belongs to a single **primary** shard

```
curl -XPUT 'http://localhost:9200/logs'
{
  "settings" : {
    "number_of_shards" : 3,
    "number_of_replicas" : 1
  }
}
```

Single node cluster



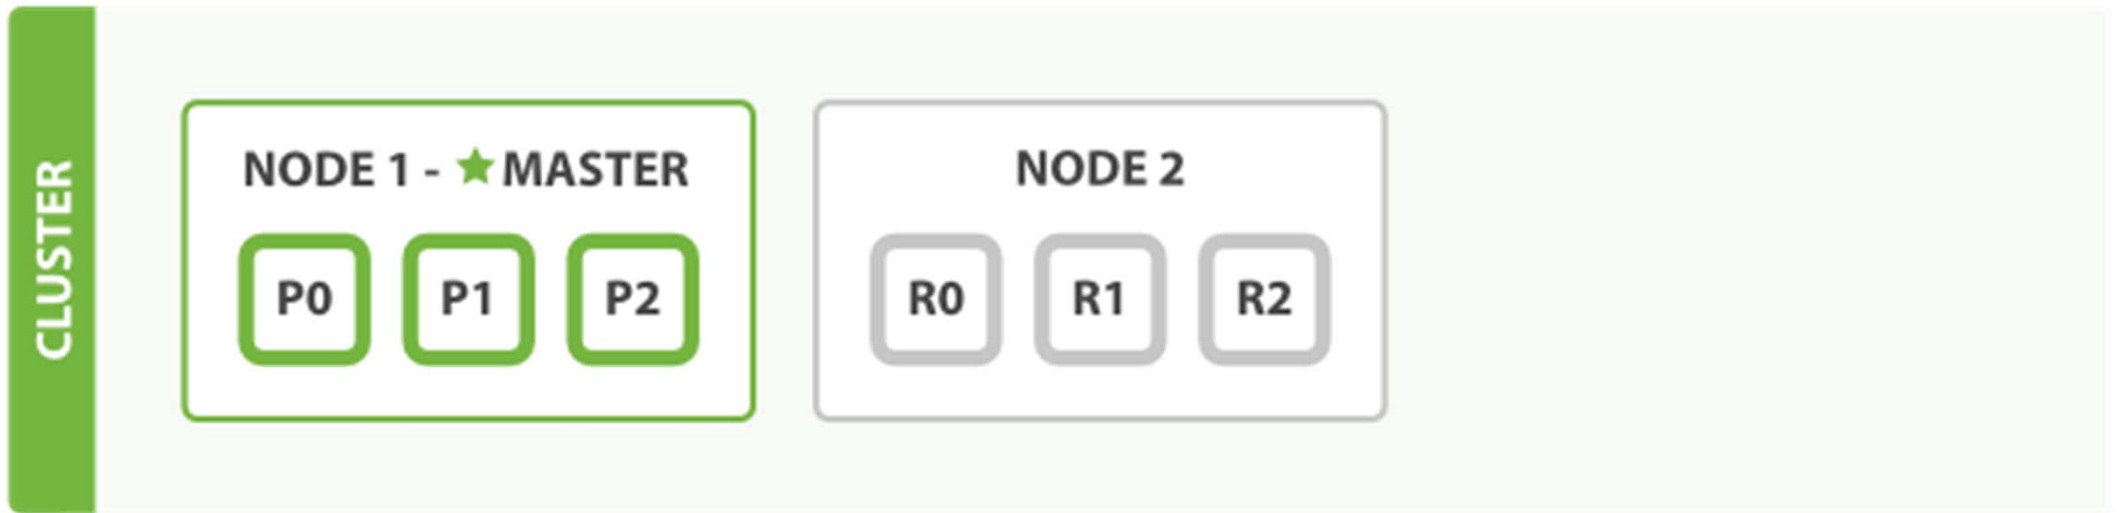
one **node** with three primary shards

creates a **cluster** of one node

node is elected to **master** role within the cluster

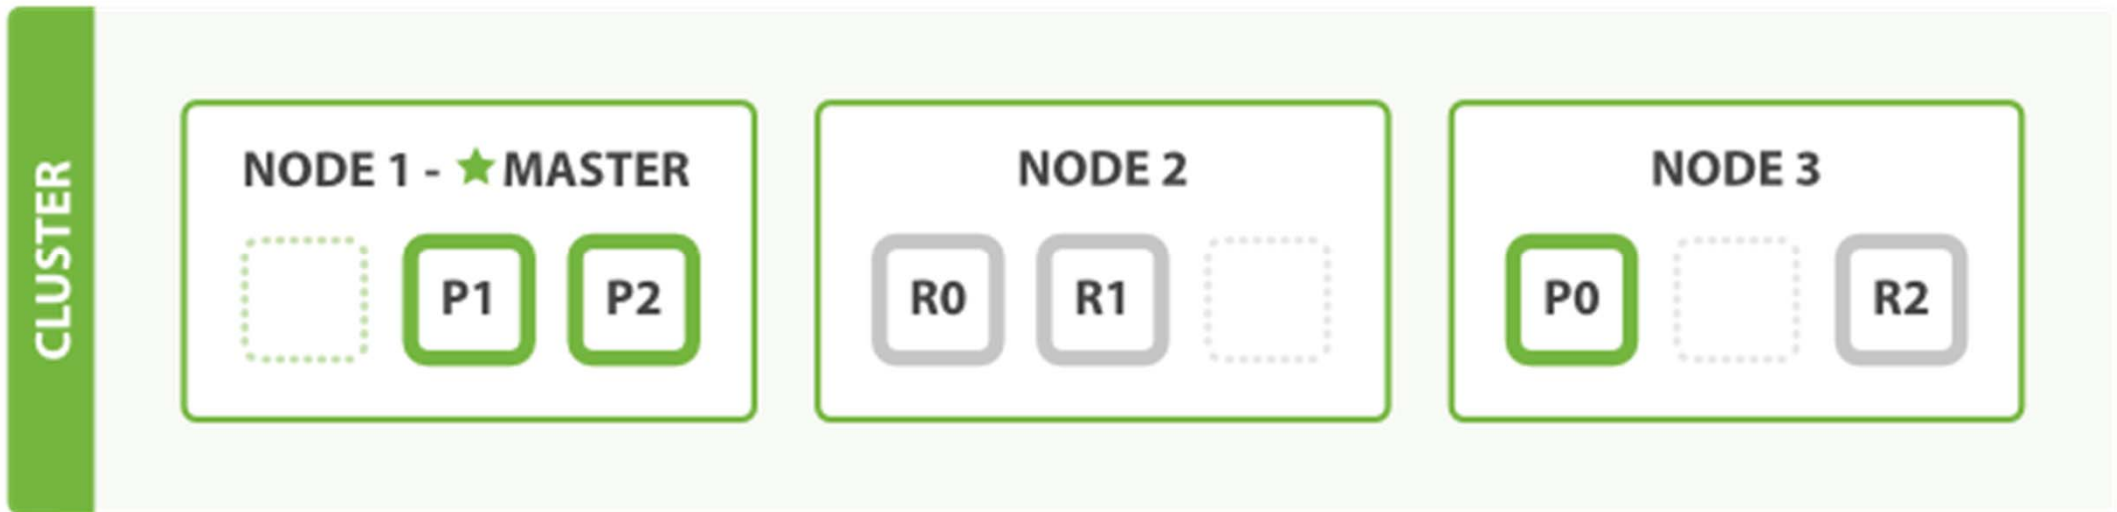
replica shards not allocated

Add Resiliency



second **node** started with same **cluster.name**
node joins cluster (**discovery** unicast/multicast)
replica shards **automatically** allocated to second node

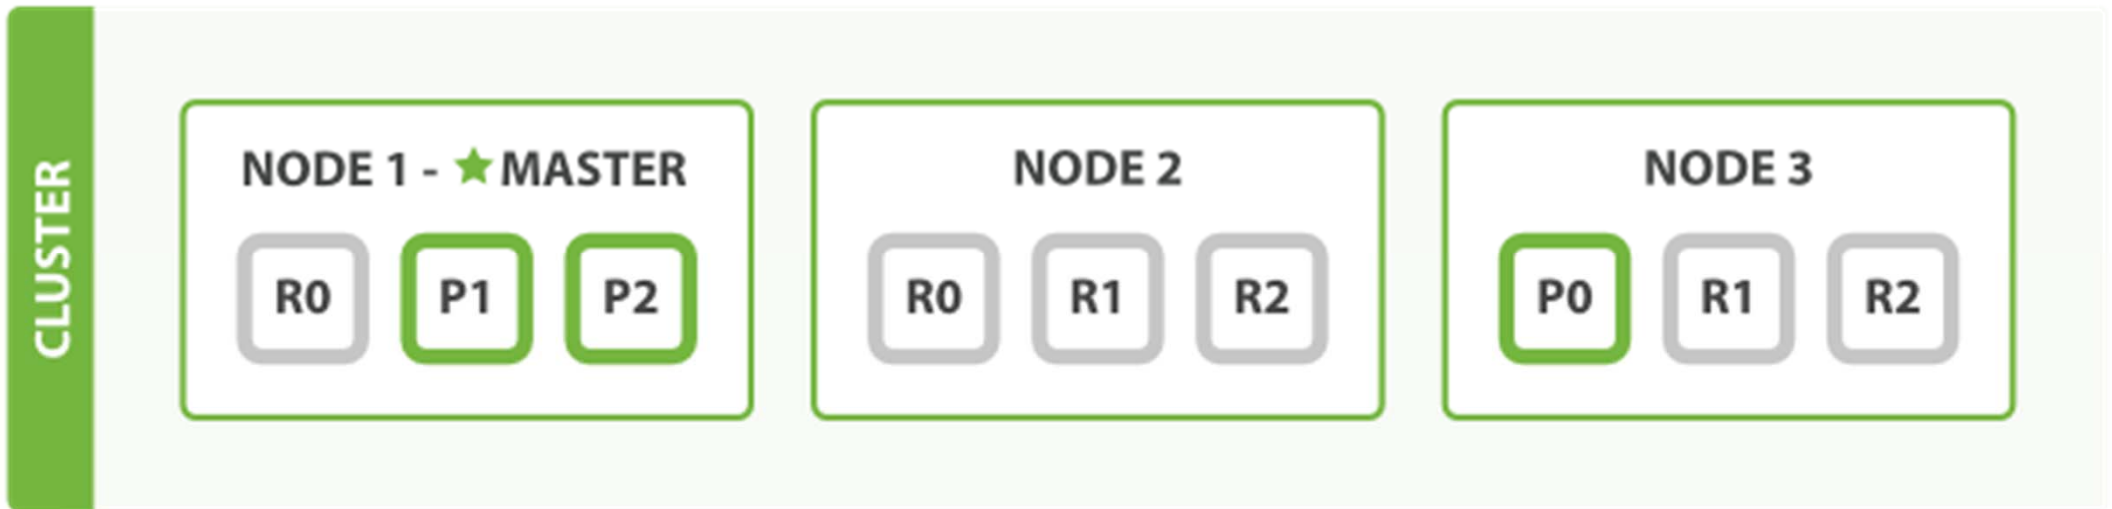
Scale Horizontally



add another node

elasticsearch **automatically** balances data

Scaling out more (number_of_replicas: n)



number of **primary** shard **fixed** at index creation
can **dynamically** increase the number of **replica** shards
more copies of you data means higher **read** throughput

Coping with failure



previous master node fails

triggers a new **master node election**

new master instantly **promotes** replicas to primary

Distributed

- **Replication**: Data duplication
 - read scalability
 - high-availability
- **Sharding**: Data partitioning
 - split logical data over several machines
 - write scalability
 - control data flow







query dsl

query dsl

**flexible, powerful
query language**

queries

- **relevance**
- **full text**
- **not cached**
- **slower**

filters

- **boolean yes/no**
- **exact values**
- **cached**
- **faster**

Filter first, then query remaining docs

GET /_search

```
{  
  "query": {...}  
}
```


GET /_search

```
{  
  "query": { "match": { "title": "search" } }  
}
```

query dsl: filtered query

GET /_search

```
{  
  "query": {  
    "filtered": {  
      "query": {...},  
      "filter": {...}  
    }  
  }  
}
```

query dsl: filtered query

GET /_search

```
{  
  "query": {  
    "filtered": {  
      "query": { "match": { "title": "search" } },  
      "filter": { "term": { "status": "active" } }  
    }  
  }  
}
```

other filter types

WHERE field CONTAINS "value"

term filter

```
"term": {  
  "title": "brown"  
}
```

WHERE field IN ["val",...]

terms filter

```
"terms": {  
  "title": ["quick", "pets"]  
}
```

other filter types

**WHERE field >= x
AND field < y**

range filter

```
"range": {  
  "content": {  
    "gte": 10,  
    "lt": 80  
  }  
}
```

```
"range": {  
  "date": {  
    "gte": "2014-01-01",  
    "lt": "2041-02-01"  
  }  
}
```

boolean filter types

```
"bool": {  
  "must": [ <filters> ],      AND  
  "should": [ <filters> ],   OR  
  "must_not": [ <filters> ]  NOT  
}
```

query dsl: full example

```
{
  "filtered": {
    "query": { "match": { "title": "full text search" }},
    "filter": {
      "bool": {
        "must": { "range": { "created": { "gte": "now - 1d / d" }}}},
        "should": [
          { "term": { "featured": true }},
          { "term": { "starred": true }}
        ],
        "must_not": { "term": { "deleted": false }}
      }
    }
  }
}
```

query dsl: filter caching

```
{
  "filtered": {
    "query": { "match": { "title": "full text search" }},
    "filter": {
      "bool": {
        "must": { "range": { "created": { "gte": "now - 1d / d" }}}},
        "should": [
          { "term": { "featured": true }},
          { "term": { "starred": true }}
        ],
        "must_not": { "term": { "deleted": false }}
      }
    }
  }
}
```


analytics (aggregations dsl)



Types of Aggregations

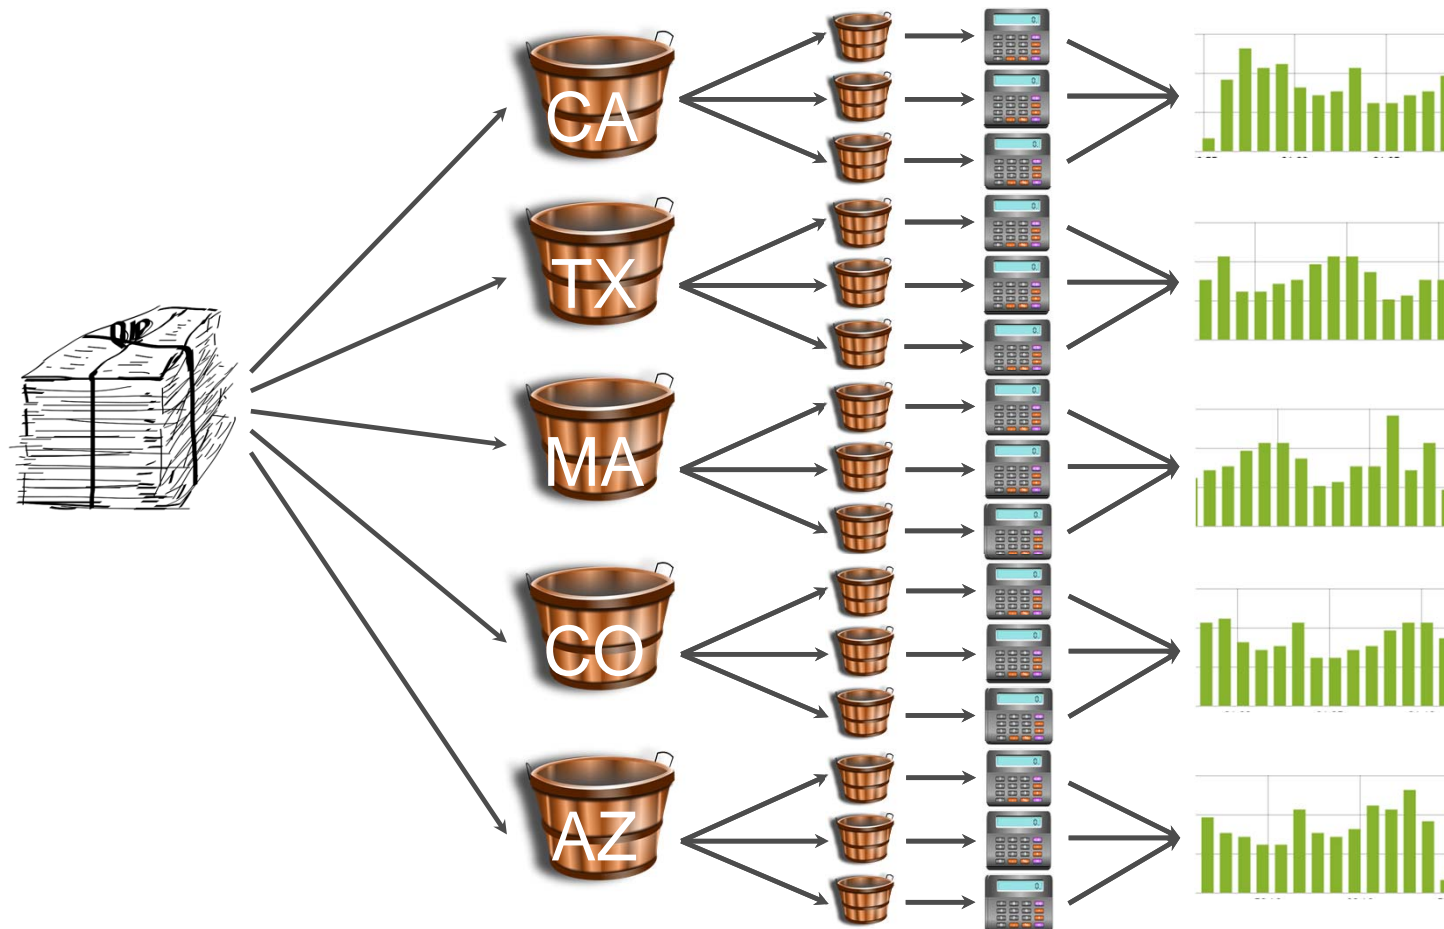
Buckets

- Terms
- Date Histogram
- Filter
- Range
- Nested
- Children
-

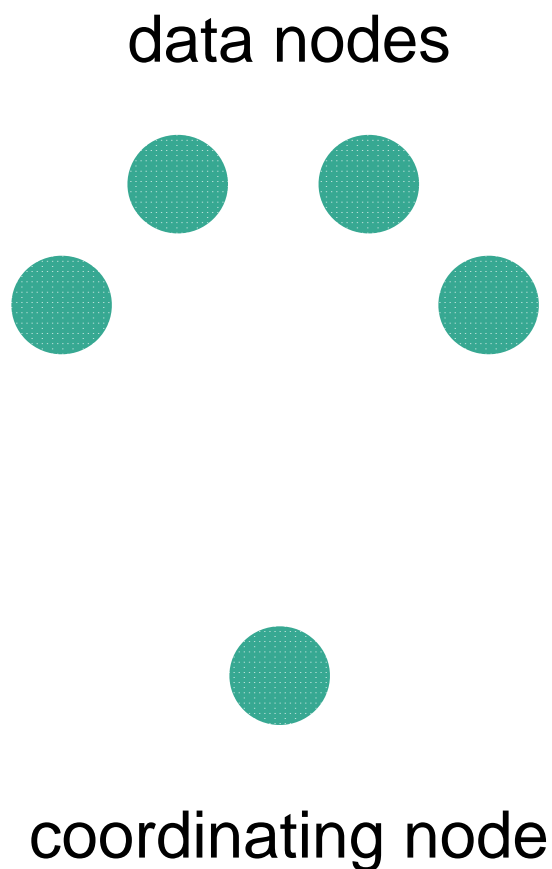
Metrics

- Stats
- Percentile
- Cardinality
- Top hits
- Scripted
- Max | Min | Avg
-

aggs = buckets + calculated metric

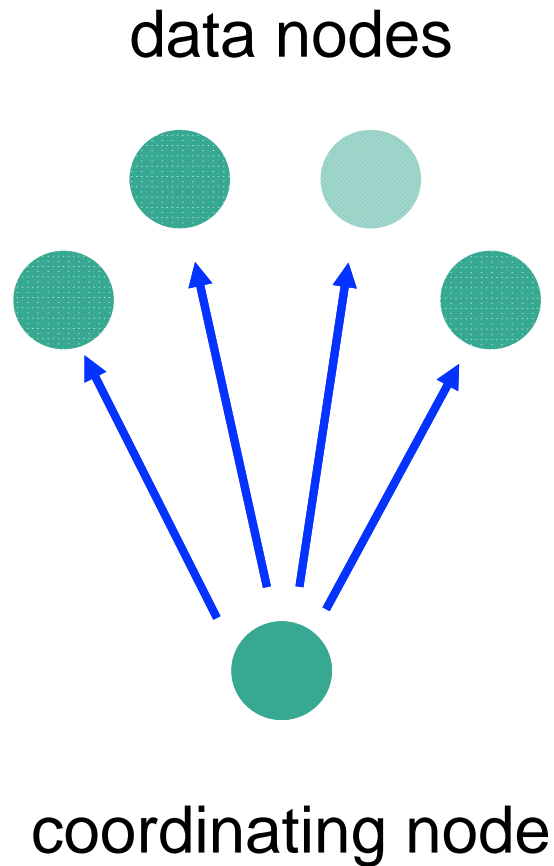


How do aggs work?



- 'inline' with search query
- execute in isolation on each shard
- 4 phases
 - parse
 - collect
 - combine
 - reduce

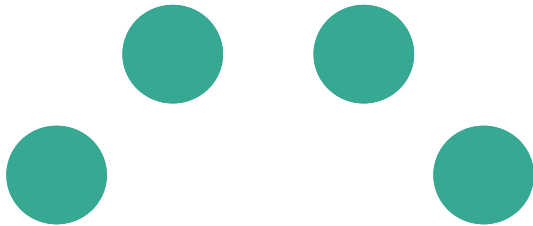
Phase 1 : Parse



- Coordinating node splits the request into shard request
- shards parse aggregation and initialize data structures

Phase 2 + 3: Collect & Combine

data nodes

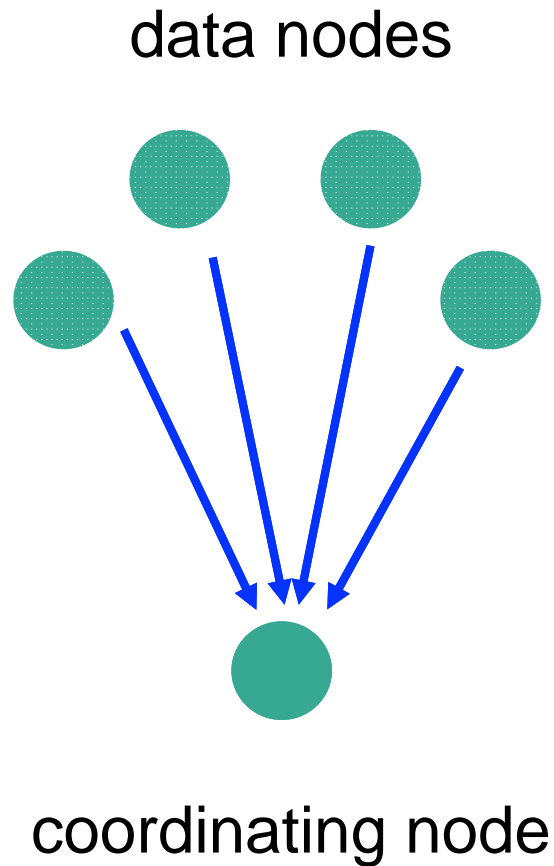


- shards process all matching documents
- once done, they combine the aggregated data into an aggregation



coordinating node

Phase 4: Reduce



- shards send their aggregation to the coordinating node
- coordinating node reduces them into a single aggregation

Designed for speed and scale

- Single network round-trip
- Single pass through the data on shards
- Aggregates are computed in-memory
- Trades accuracy for speed in some use cases
- Aggregations can be composed
- Near real-time response times

Aggregation DSL Example

Request

```
..
  "aggs": {
    "by_date": {
      "date_histogram": {
        "field": "timestamp",
        "interval": "day"
      },
      "aggs": {
        "max_temperature": {
          "max": {
            "field": "temperature"
          }
        }
      }
    }
  }
  ...
```

Response

```
..
  "aggregation": {
    "by_date": {
      "buckets": [
        {
          "key": "2015-01-01T00:00:00.000Z",
          "doc_count": 24,
          "max_temperature": {
            "value": 23
          }
        }
      ]
    }
  }
  ...
```

Kibana Deep-Dive

Pere Urbon-Bayes
Software Engineer



DEMO

Real-World Use Cases

Alan Hardy
Solutions Architect



Use Case: Large-Scale Log Analytics for Infrastructure and Application Monitoring

- Netflix Platform Engineering team deployed a centralized ELK stack as part of their real-time query service
- Billions of events ingested into a system called Suro using Kafka, representing infrastructure events, application logs, API calls
- Hadoop used for batch analytics, ES+Kibana for real-time interactive analytics



Scale & Environment

- Multiple clusters with varying sizes
- Indexing billions events daily
- Elastic{ON} 2015 video available talking about details

Verizon Business

Use Case: Large-Scale Log Analytics for Security

- The ELK stack is powering security analytics for Verizon's Data Breach Investigation Report (DBIR)
- Previous solution was based on Hadoop and while functional was not optimized to deliver real-time, interactive insights
- Indexed data from Hadoop to enable real-time interactive analytics



Scale & Environment

- One cluster with 128+ nodes
- 500+ billion documents
- 10B+ documents indexed daily
- Elastic{ON} 2015 video available talking about details

Q & A

