



Elastic Stack Data Ingestion & Visualizations

Aravind Putrevu

Developer | Evangelist

@aravindputrevu | aravindputrevu.in



Agenda

- 1 The Ingest Story
- 2 Logstash: Only for Logs?
- 3 Configuring and Managing Logstash
- 4 Visualizing Data Ingested in Kibana
- 5 Demo

Agenda

- 1 The Ingest Story
- 2 Logstash: Only for Logs?
- 3 Configuring and Managing Logstash
- 4 Visualizing Data Ingested in Kibana
- 5 Demo

Agenda

- 1 The Ingest Story
- 2 Logstash: Only for Logs?
- 3 Configuring and Managing Logstash
- 4 Visualizing Data Ingested in Kibana
- 5 Demo

Agenda

- 1 The Ingest Story
- 2 Logstash: Only for Logs?
- 3 Configuring and Managing Logstash
- 4 Visualizing Data Ingested in Kibana
- 5 Demo


Agenda

- 1 The Ingest Story
- 2 Logstash: Only for Logs?
- 3 Configuring and Managing Logstash
- 4 Visualizing Data Ingested in Kibana
- 5 Demo



Elastic Stack

No enterprise edition
All new versions with 6.2



**SECURITY
ANALYTICS**

**LOG
ANALYTICS**

**METRICS
ANALYTICS**

**BUSINESS
ANALYTICS**

SEARCH

APM



Protect your
data




Be alerted on
changes




Detect anomalies



Monitor your
Elastic Stack



Find links in
your data




Share your
insights



Logstash

Data processing pipeline



**Ingest data of all shapes,
sizes, and sources**


**Parse and dynamically
transform data**


**Transport data to any
output**

**Secure and encrypt data
inputs**


Build your own pipeline

More than 200+ plugins






Popular Data Sources



How it works?



Configuring Logstash

```
# This is a comment. You should use comments to describe
# parts of your configuration.
input {
  ...
}

filter {
  ...
}

output {
  ...
}
```

Configuring Logstash

client ip timestamp

↑ ↑

```
127.0.0.1 -- [11/Dec/2013:00:01:45 -0800] "GET /xampp/status.php HTTP/1.1" 200 3891  
"http://cadenza/xampp/navi.php" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10.9; rv:25.0) Gecko/20100101 Firefox/25.0"
```

↓ ↓

request agent

The diagram illustrates the mapping of a Logstash log entry to its constituent fields. The log entry is shown in two lines: the first line contains the client IP, timestamp, request method and path, status code, and response size; the second line contains the request path and the user agent string. Arrows point from the labels 'client ip', 'timestamp', 'request', and 'agent' to their respective parts in the log entry.

Configuring Logstash

```
input { stdin { } }

filter {
  grok {
    match => { "message" => "%{COMBINEDAPACHELOG}" }
  }
  date {
    match => [ "timestamp" , "dd/MMM/yyyy:HH:mm:ss Z" ]
  }
}

output {
  elasticsearch { hosts => ["localhost:9200"] }
  stdout { codec => rubydebug }
}
```

Configuring Logstash

```
{
  "message" => "127.0.0.1 - - [11/Dec/2013:00:01:45 -0800] \"GI
"@timestamp" => "2013-12-11T08:01:45.000Z",
"@version" => "1",
  "host" => "cadenza",
  "clientip" => "127.0.0.1",
  "ident" => "-",
  "auth" => "-",
  "timestamp" => "11/Dec/2013:00:01:45 -0800",
  "verb" => "GET",
  "request" => "/xampp/status.php",
  "httpversion" => "1.1",
  "response" => "200",
  "bytes" => "3891",
  "referrer" => "\"http://cadenza/xampp/navi.php\"",
  "agent" => "\"Mozilla/5.0 (Macintosh; Intel Mac OS X 10.9;
}
```

Grok

```
2016-09-19T18:19:00 [8.8.8.8:prd] DEBUG this is an example log message
```

```
%{TIMESTAMP_ISO8601:timestamp} \[%{IPV4:ip};%{WORD:environment}\] %{LOGLEVEL:log_level} %{GREEDYDATA:message}
```

```
{
  "timestamp": "2016-09-19T18:19:00",
  "ip": "8.8.8.8",
  "environment": "prd",
  "log_level": "DEBUG",
  "message": "this is an example log message"
}
```

Best Practices for Grok

- Grok may not perform well when a match fails
- Monitor the occurrence of `_grokparsefailures` and then benchmark their cost
- Use anchors such as `^` and `$` to remove ambiguity and aid the regex engine
- Tiered matching increases performance if you don't use anchors, otherwise don't bother. When in doubt, measure!
- Use Monitoring or Metrics API.



```
1 55.3.244.1 DEBUG This is Grok Test!
```

```
1  %{IP:client} %{WORD:loglevel} %{GREEDYDATA:message}
```

Simulate


```
1 {
2   "loglevel": "DEBUG",
3   "client": "55.3.244.1",
4   "message": "This is Grok Test! "
5 }
```



Pipelines in 1, 2, 3

Start Fast, Think Big


- Configure your pipelines, not code them
- Stash your first event in minutes
- Grow iteratively, scale horizontally



Multiple Pipelines


- pipeline.id: my-pipeline_1
path.config: "/etc/path/to/p1.config"
pipeline.workers: 3
- pipeline.id: my-other-pipeline
path.config: "/etc/different/path/p2.cfg"
queue.type: persisted

Centralized Pipeline Management



Logstash Modules

- Elastic Cloud
- Arcsight
- Netflow



*

Uses lucene query syntax




Add a filter +

Netflow: Dashboard Navigation

[Overview](#) | [Conversation Partners](#) | [Traffic Analysis](#) | [Top-N](#) | [Geo Location](#) | [Autonomous Systems](#) | [Flow Exporters](#) | [Raw Flow Records](#)

Netflow: Countries and Cities (flow records)



- United States
- Canada
- Australia
- Netherlands
- China
- Ireland
- Antigua and Barbu...
- New Zealand
- Singapore

Netflow: Destinations and Sources (flow records)




- 172.16.139.250
- 172.16.133.6
- 8.8.8.8
- 172.16.128.169
- 172.16.133.54
- 172.16.128.202
- 172.16.133.132
- 172.16.133.41
- 172.16.133.254

Netflow: Destination and Source Ports (flow records)



- TCP/5440
- TCP/80 (http)
- TCP/443 (https)
- UDP/53 (dns)
- UDP/161 (snmp)
- UDP/1853 (vids-avtp)
- UDP/1900 (ssdp)
- UDP/8200 (trivnet1)
- UDP/137 (netbios-...)

Netflow: Geo Location Heatmap




Data Resiliency

- Persistent Queues
- Dead Letter Queues




Persistent Queue

- Handle bursts of events
- At-least once delivery



Dead Letter Queue



Scaling Logstash

The screenshot displays the Kibana Monitoring dashboard. The left sidebar contains navigation links: Discover, Visualize, Dashboard, Timeline, Machine Learning, APM, Graph, Dev Tools, Monitoring (selected), and Management. The main content area shows three sections: Elasticsearch, Kibana, and Logstash. The Elasticsearch section indicates a health status of 'yellow' and provides an overview of its version (6.2.0), uptime (2 minutes), and jobs (0). It also lists metrics for a single node, including disk available space (17.65%), JVM heap usage (46.22%), and indices (13). The Kibana section shows a 'green' health status and an overview of requests (60) and response time (250 ms). The Logstash section, highlighted with a red rounded rectangle, shows a 'green' health status and an overview of events received (180) and emitted (180). It also lists metrics for a single node, including uptime (2 minutes) and JVM heap usage (28.88%).

Elasticsearch • Health is yellow

Overview

Version	6.2.0
Uptime	2 minutes
Jobs	0

Nodes: 1

Disk Available	17.65%
	41GB / 234GB
JVM Heap	46.22%
	458MB / 990MB

Indices: 13

Documents	26,025
Disk Usage	26MB
Primary Shards	17
Replica Shards	0

Kibana • Health is green

Overview

Requests	60
Max. Response Time	250 ms

Instances: 1

Connections	24
Memory Usage	13.44%
	192MB / 1GB

Logstash

Overview

Events Received	180
Events Emitted	180

Nodes: 1

Uptime	2 minutes
JVM Heap	28.88%
	286MB / 990MB

Pipelines: 1

Nodes: 1 Memory: 226MB / 990MB Events Received: 609 Events Emitted: 609


Events Received Rate (/s)



Events Emitted Rate (/s)




Event Latency (ms)




[Overview](#) [Pipelines](#) [Advanced](#)**127.0.0.1:9600** Events Received: **606** Events Emitted: **606** Config Reloads: **0** Pipeline Workers: **8** Batch Size: **125** Version: **6.0.0** Uptime: **6 minutes**

Events Received Rate (/s)




Events Emitted Rate (/s)




Event Latency (ms)



JVM Heap (MB)




CPU Utilization (%)




System Load




Measure...Measure...Measure...



Scaling Logstash



Scaling Logstash



Logstash Plugins

- Input
- Output
- Filter

Parse/Extract

The Grok Filter

The go-to-tool for parsing fields

```
filter {  
  grok {  
    match => {"message" => "%{TIMESTAMP_8601:ts}%{SPACE}%{GREEDYDATA:message}"}  
  }  
}
```

The Date Filter

Use data strings to set @timestamp

```
filter {  
  date {  
    match => ["timestamp_string", "ISO8601"]  
  }  
}
```

The KV Filter

The easy way to parse data in key/value pairs

```
filter {  
  kv {  
    source => "message"  
    target => "parsed"  
    value_split => ":"  
  }  
}
```

Core operations

The Mutate Filter

The go-to-tool for field manipulation in Logstash

- Convert field types (from strings to integers etc.)
- Add/rename/replace/copy fields
- Upper/lowercase transformation
- Join arrays together (useful for Array=>String operations)
- Merge hashes
- Split fields into Arrays
- Strip whitespace

Core Transformation Filters

- Mutate - Modify / Add Individual Fields
- Split - Divide a single event into multiple events
- Drop - Delete an event

Core Operations Example

```
filter {  
  mutate { lowercase => "account" }  
  if [type] == "batch" {  
    split { field => actions target => action }  
  }  
  if { "action" =~ /special/ } {  
    drop {}  
  }  
}
```

Search Supportive

GeoIP Filter

Enrich IP address information

```
filter {  
  geoip {  
    fields => "my_geoip_field"  
  }  
}
```

User Agent Filter

Enrich browser user agent information

```
filter {  
  useragent {  
    source => "useragent"  
  }  
}
```

Translate Filter

Use local data to map / enrich events

```
filter {  
  translate {  
    dictionary => [ "100", "Continue",  
                   "101", "Switching Protocols",  
                   "merci", "thank you",  
                   "old version", "new version" ]  
  }  
}
```

Elasticsearch Filter

Use Elasticsearch Data to Enrich Events

```
elasticsearch {  
  hosts => ["es-server"]  
  query => "type:start AND operation:%{[opid]}"  
  fields => { "@timestamp" => "started" }  
}
```

JDBC Streaming Filter

Use a database to enrich events

```
filter {  
  jdbc_streaming {  
    jdbc_driver_library => "/path/to/mysql-connector-java-5.1.34-bin.jar"  
    jdbc_driver_class => "com.mysql.jdbc.Driver"  
    jdbc_connection_string => "jdbc:mysql://localhost:3306/mydatabase"  
    jdbc_user => "me"  
    jdbc_password => "secret"  
    statement => "select * from WORLD.COUNTRY WHERE Code = :code"  
    parameters => { "code" => "country_code"  
    target => "country_details"  
  }  
}
```

Add Data to Kibana

Use these solutions to quickly turn your data into pre-built dashboards and monitoring systems.



APM

APM automatically collects in-depth performance metrics and errors from inside your applications.

Add APM



Logging

Ingest logs from popular data sources and easily visualize in preconfigured dashboards.

Add log data



Metrics

Collect metrics from the operating system and services running on your servers.

Add metric data



Security Analytics

Centralize security events for interactive investigation in ready-to-go visualizations.

Add security events

Data already in Elasticsearch? [Set up index patterns](#)

Visualize and Explore Data



APM

Automatically collect in-depth performance metrics and errors from inside your applications.



Dashboard

Display and share a collection of visualizations and saved searches.



Discover

Interactively explore your data by querying and filtering raw documents.



Graph

Surface and analyze relevant relationships in your Elasticsearch data.



Machine Learning

Automatically model the normal behavior of your time series data to detect anomalies.



Timelion

Use an expression language to analyze time series data and visualize the results.



Visualize

Manage and Administer the Elastic Stack



Console

Skip cURL and use this JSON interface to work with your data directly.



Index Patterns

Manage the index patterns that help retrieve your data from Elasticsearch.



Monitoring

Track the real-time health and performance of your Elastic Stack.



Saved Objects

Import, export, and manage your saved searches, visualizations, and dashboards.



Security Settings

Protect your data and easily manage who has access to what with users and roles.



Watcher

Detect changes in your data by creating, managing, and monitoring alerts.

Resources

- <https://www.elastic.co/learn>
- <https://www.elastic.co/blog/category/engineering>
- <https://discuss.elastic.co/>
- <https://fb.com/groups/ElasticIndiaUserGroup>
- <https://elastic.co/community>

Fin!



discuss.elastic.co | aravind@elastic.co | [@aravindputrevu](https://twitter.com/aravindputrevu)