NETWAYS®

ELK – Stack

Elasticsearch, Logstash, Kibana

Munich – 19.10.2015

NETWAYS®

INTRODUCTION

Bernd Erk

- CEO at NETWAYS GmbH
- Co-Founder Icinga
- **@gethash**
- **info@netways.de**

NETWAYS GmbH

- Open Source Service Provider
- Located in Nuremberg
- About 45 employees right now

- Technical areas
  - Open Source Systems Management
  - Open Source Datacenter
- Custom Open Source solutions

NETWAYS Products

NETWAYS®

INTRODUCTION LOGS & EVENTS

Logs

Logs -> Flow of **unstructured** data

Oct 4 16:57:24 web sshd[25828]: Received disconnect from 10.10.0.31: 11: disconnected by user

Consists of timestamp and message

Events

Event -> Flow off **structured** data

Event {

Time: Oct 4 16:57:24

Process: sshd

State: Received disconnect from 10.10.0.31

Client: 10.10.0.31

}

consists of detailed attributes

Log & Eventmanagement

Logs > Event > Analyse (Correlation) > Action

Tools

- Nagios & Icinga Addons
  - Check_logfiles
  - EventDB
- Check_MK Event Console
- Logmanagement-Tools
  - ELK-Stack
  - Graylog
  - Fluentd

ELK Stack

NETWAYS®

ARCHITECTURE & INSTALLATION

Overview

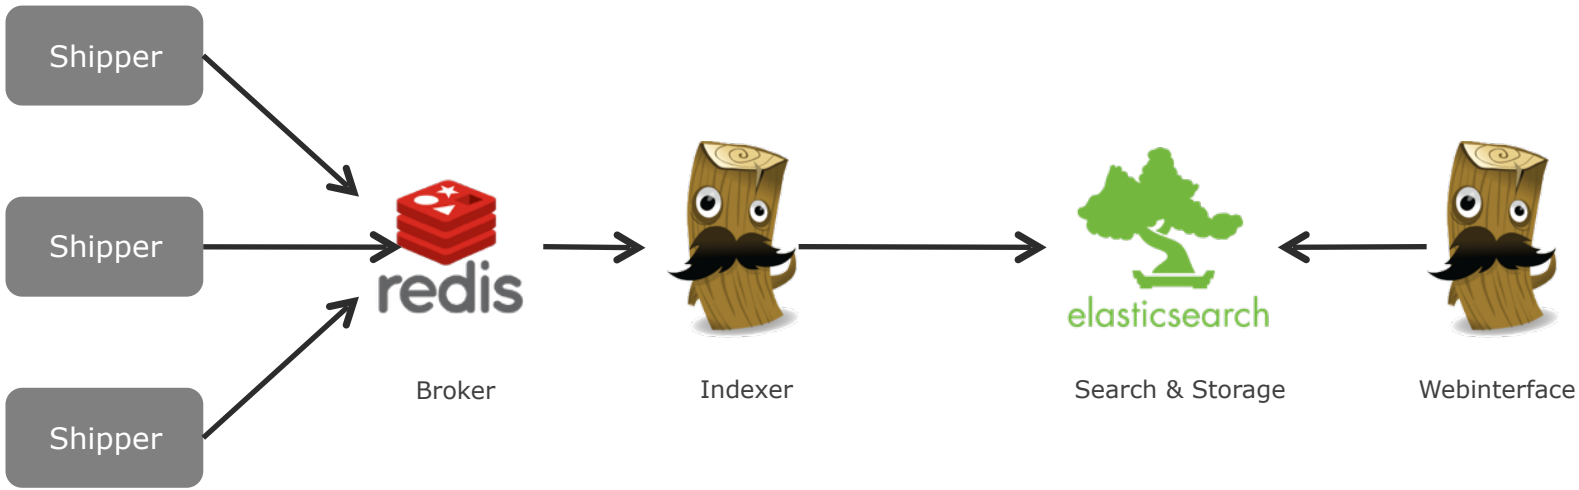
Shipper
Shipper
Shipper
redis
elasticsearch
Broker
Indexer
Search & Storage
Webinterface

Logstash

- Logmanagement based on JRuby
- Configurable “Pipe”
- Flexible Plugin-Architecture for
  - Input
  - Filter
  - Output
- Single File Deployment

Logstash - IO

Outputs

- amqp
- boundary
- circonus
- cloudwatch
- datadog
- datadog_metrics
- elasticsearch
- elasticsearch_http
- elasticsearch_river
- email
- exec
- file
- ganglia
- gelf
- gemfire
- google_cloud_storage
- graphite
- graphtastic
- hipchat
- relp
- s3
- snmptrap
- sqlite
- sqs
- stdin
- stomp
- syslog
- tcp
- twitter
- udp
- unix
- varnishlog
- websocket
- wmi
- xmpp
- zenoss
- zeromq

Inputs

- amqp
- drupal_dblog
- elasticsearch
- eventlog
- exec
- file
- ganglia
- gelf
- gemfire
- generator
- graphite
- heroku
- imap
- irc
- log4j
- lumberjack
- pipe
- rabbitmq
- redis
- http
- irc
- jira
- juggernaut
- librato
- loggly
- lumberjack
- metriccatcher
- mongodb
- nagios
- nagios_nsca
- null
- opentsdb
- pagerduty
- pipe
- rabbitmq
- redis
- riak
- riemann
- s3
- sns
- sqs
- statsd
- stdout
- stomp
- syslog
- tcp
- udp
- websocket
- xmpp
- zabbix
- zeromq

Logstash - Installation

- Download - www.elastic.co/downloads/logstash
- bin/logstash agent -f <config-file>

Redis

- NoSQL in memory based on C
- Support for various “Datatypes”
  - Strings / Hashes / Lists
  - Sets and Sorted Sets
- Support for various replication scenarious
- Very high performance

```
$ ./redis-benchmark -r 1000000 -n 2000000 -t get,set,lpush,lpop -q
SET: 122556.53 requests per second
GET: 123601.76 requests per second
LPUSH: 136752.14 requests per second
LPOP: 132424.03 requests per second
```

Redis - Installation

- Download - http://redis.io/download
- make
- make test
- make install
- /usr/local/bin/redis-server

Elasticsearch

- Schema free RESTful server based on Java
- Based on Lucene Core
- “Comparable” with Apache Solr
- Distributed Architecture using
  - Shards
  - Replicas
  - Gateways
- Realtime search base for Kibana

Elasticsearch – Installation

- Download – www.elastic.co/downloads/elasticsearch
- Unpack the archive
- Run bin/elasticsearch

NETWAYS®

INPUT OFF LOGS

Overview - Logshipping

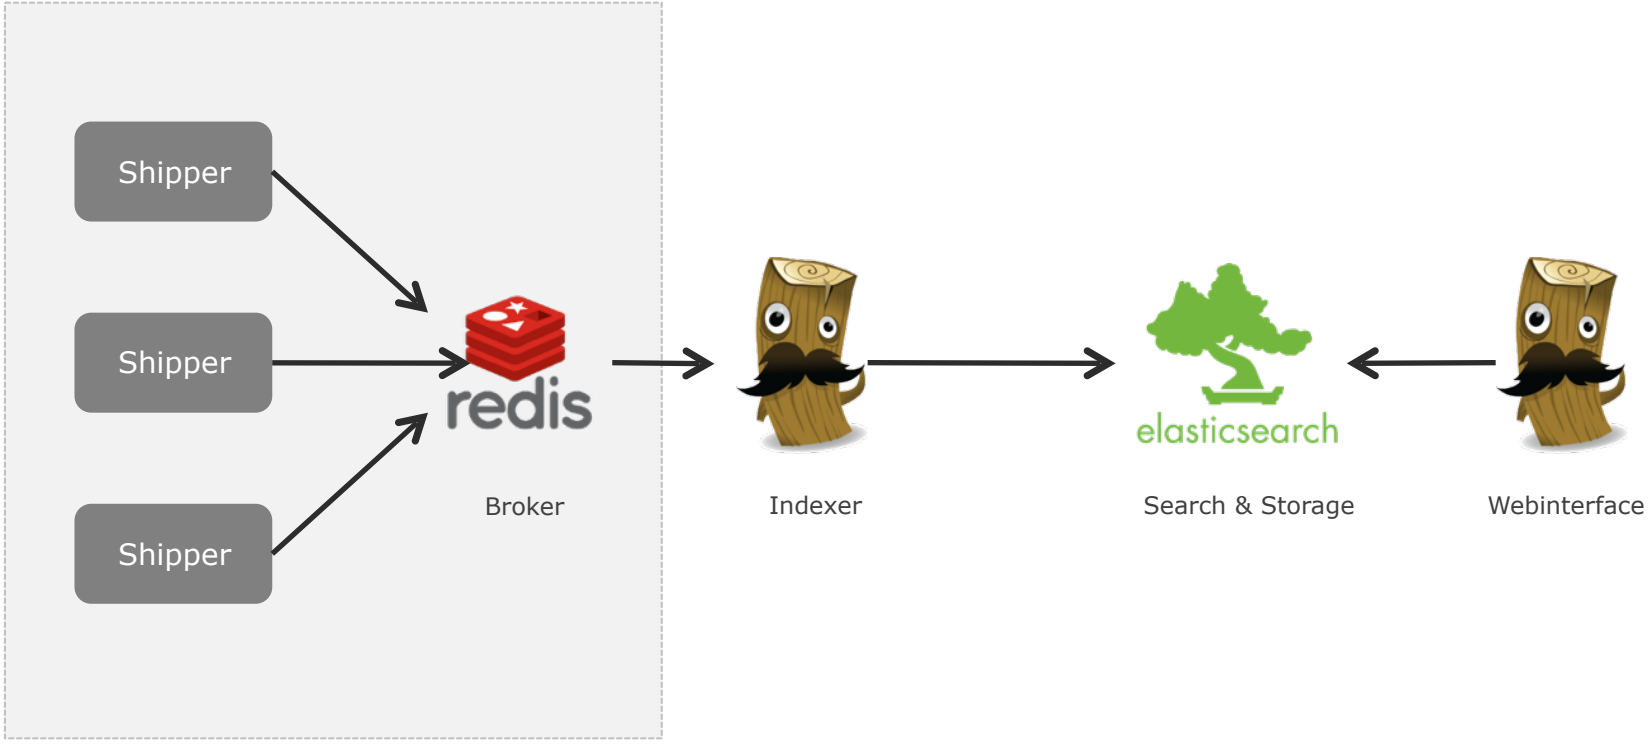

Logstash - Shipper

- Shipment of logs to Logstash
  - Logstash
  - Logstash Forwarder
  - Syslog
  - Log4J
  - Gelf
  - File-Read
  - Many more

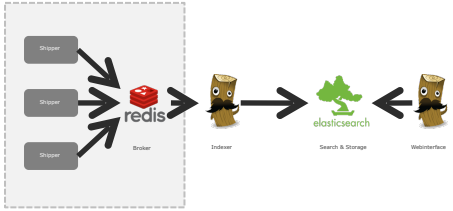

Logstash – Shipper - Configuration

- Configuration

```
input {
  file {
    path => "/root/demodata/access.log.1”
    type => "apache-access"
  }
}
output {
  stdout {
      debug => true
  }
  redis {
    host => "127.0.0.1"
    data_type => "list"
    key => "logstash.apache"
  }
}
```

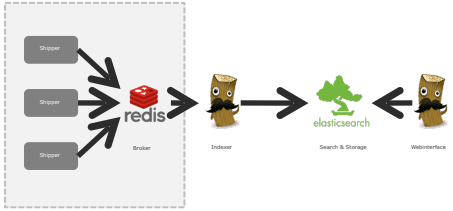

- bin/logstash agent –f logstash_shipper.conf

NETWAYS®

EVENT INDEXING

Overview - Indexing

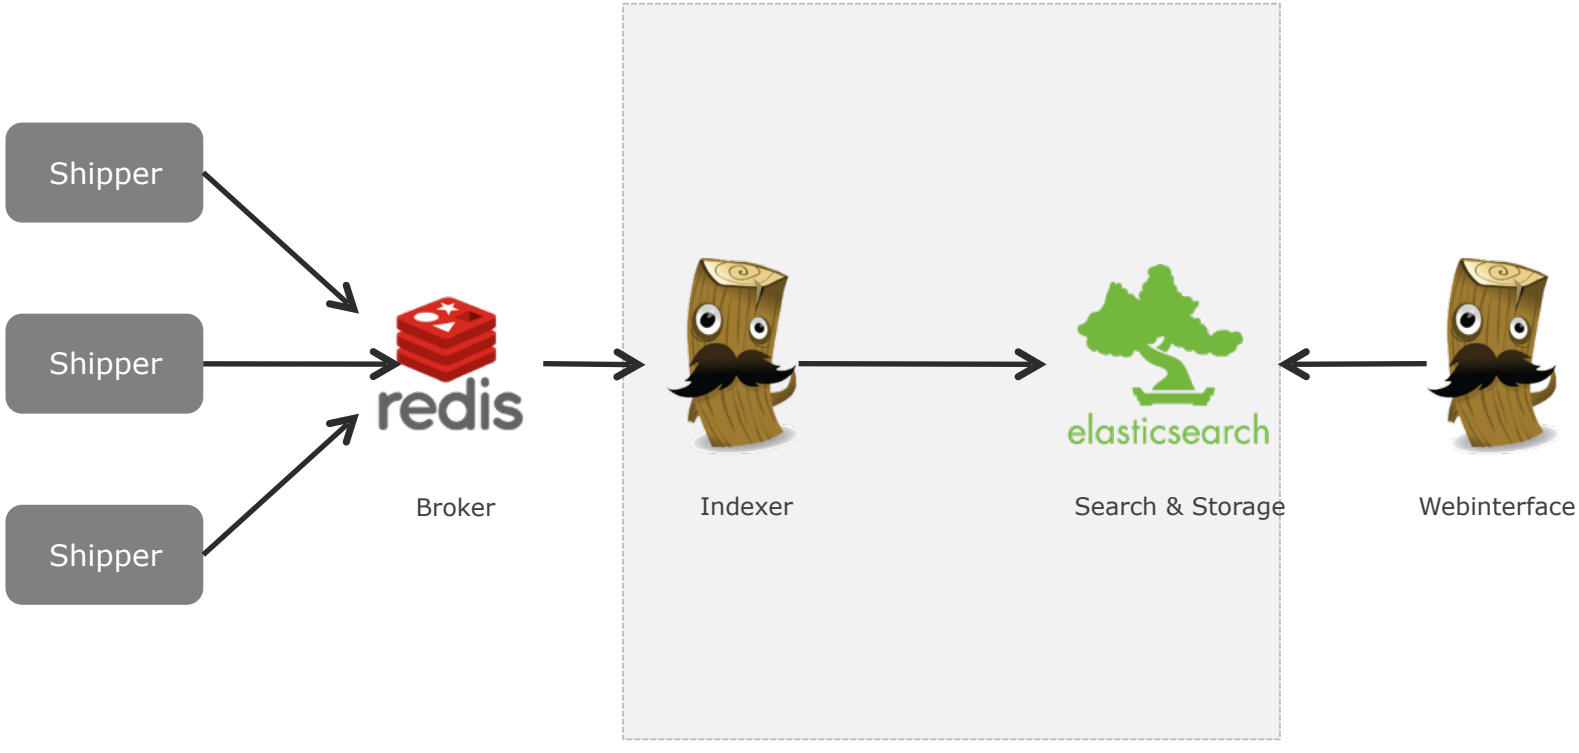

Logstash - Indexer

- Configuration

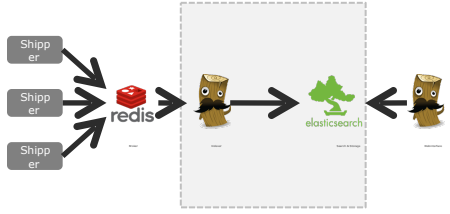


```
input {
  redis {
    host => "127.0.0.1"
    type => "redis-input"
    # these settings should match the output of the agent
    data_type => "list"
    key => "logstash.apache”
  }
}
output {
  stdout {
     debug => true
  }
  elasticsearch_http {
    host => "127.0.0.1"
  }
}
```

Bring your stuff in order

- We need more then a timestamp and message
- We need structured and queryable information
- We need **grok**

Grok - Example

55.3.244.1 GET /index.html 15824 0.043

%{IP:client} %{WORD:method} %{URIPATHPARAM:request} %{NUMBER:bytes} %{NUMBER:duration}

client: 55.3.244.1

method: GET

request: /index.html

bytes: 15824

duration: 0.043

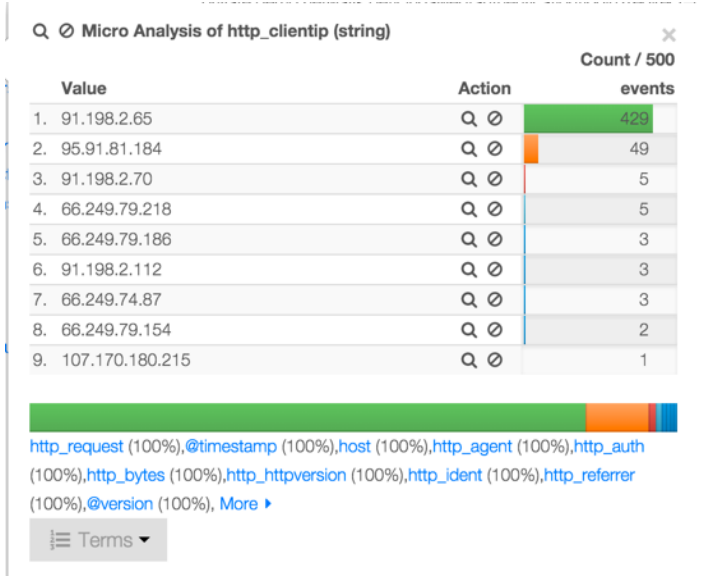

Demo

http://grokconstructor.appspot.com/do/match#result

Logstash – Indexer- Apache

- Configuration for Apache-Logs

```
input {
  redis {
    host => "127.0.0.1"
    type => "apache-access”
    data_type => "list"
    key => "logstash.apache”
    format => "json_event"
  }
}
filter {
  if [type] == "apache-access" {
    grok {
      match => [ "message", "%{COMBINEDAPACHELOG}" ]
    }
  }
}
output {
  elasticsearch_http {
     host => "127.0.0.1”
  }
}
```

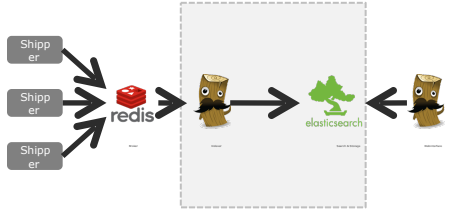

Logstash – Indexer - GEOIP

- Configuration for Geo-Data

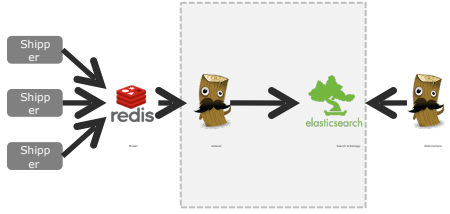


```
input {
  redis {
    host => "127.0.0.1"
    type => "apache-access”
    data_type => "list"
    key => "logstash.apache”
  }
}
filter {
  grok {
    type => "apache-access"
    pattern => "%{COMBINEDAPACHELOG}"
  }
  geoip {
    source => "clientip"
    add_tag => ["geotag"]
  }
}
output {
  elasticsearch_http {host => "127.0.0.1”}
}
```

NETWAYS®

INTERFACES & API

Overview – Interfaces

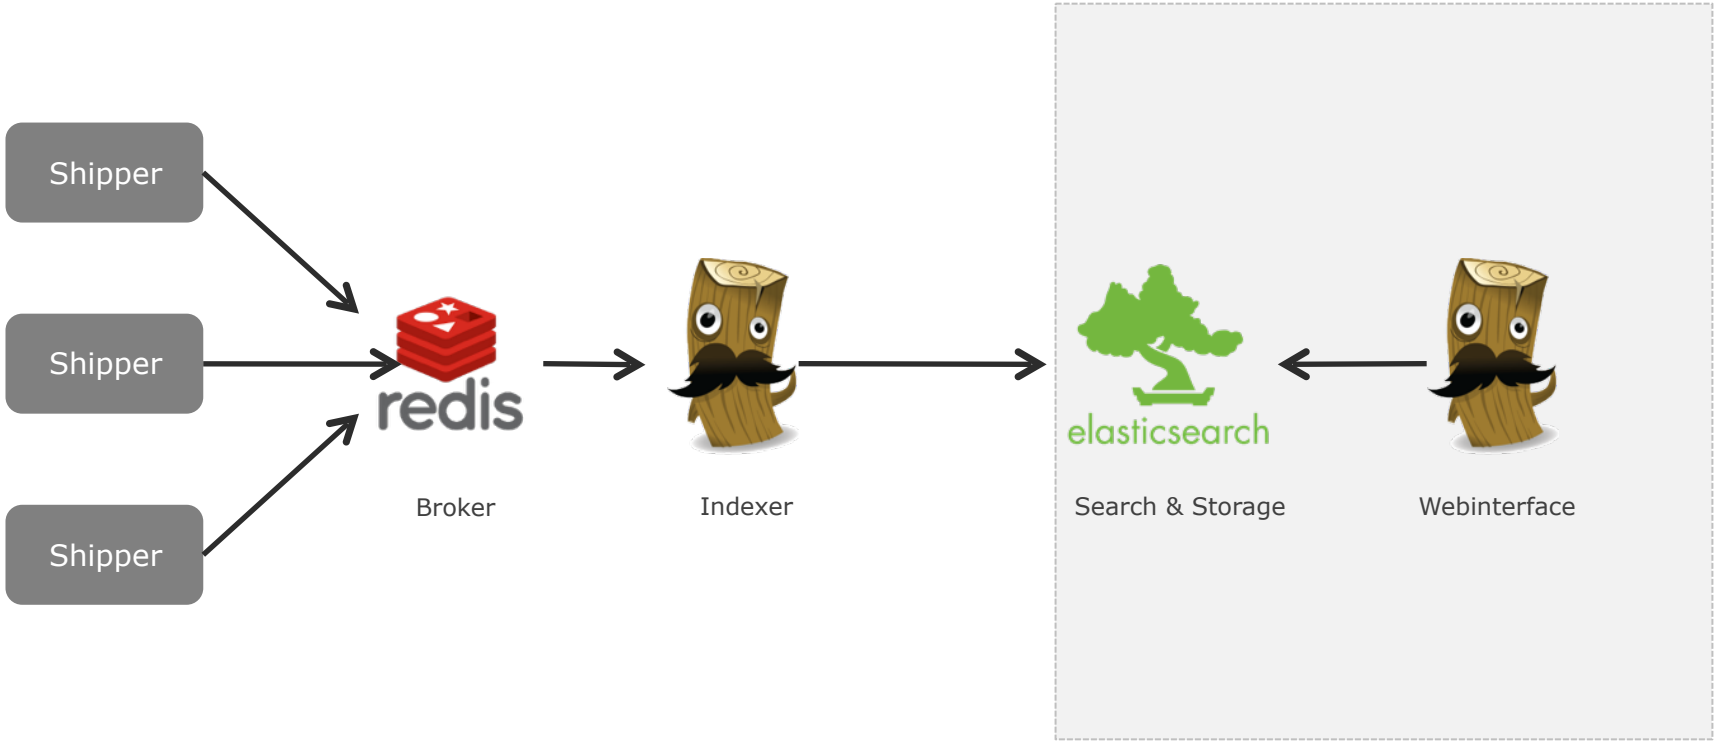

Kibana

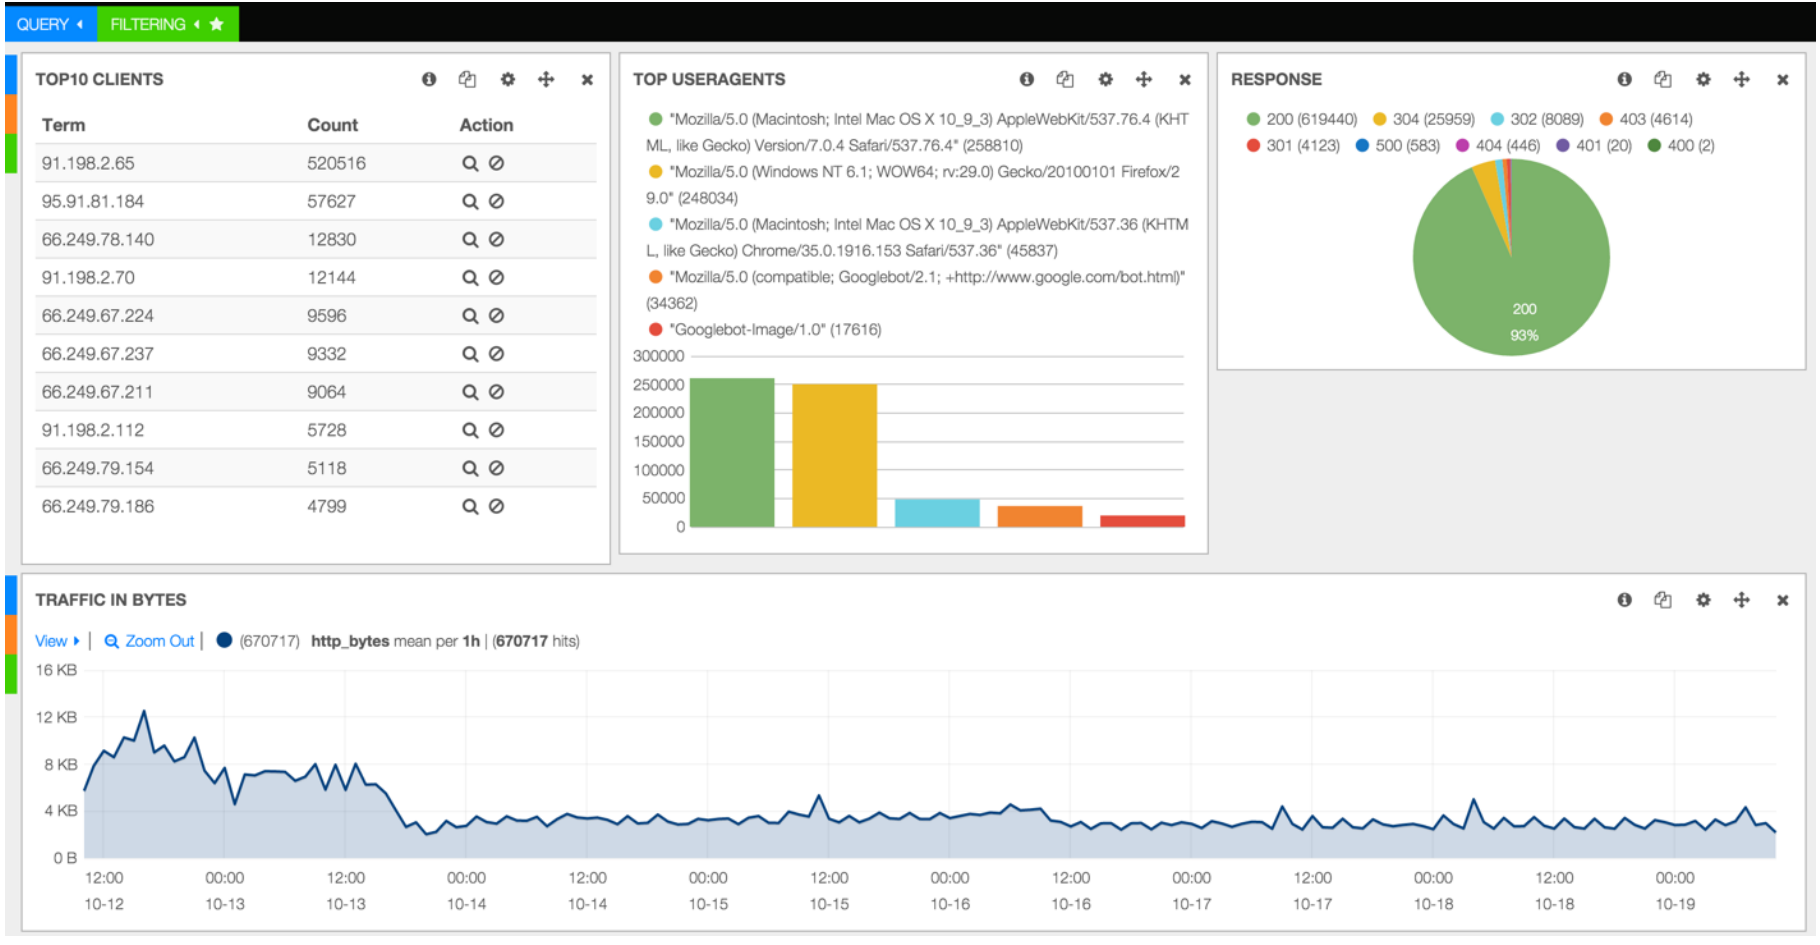
QUERY
FILTERING
TOP10 CLIENTS
Term
Count
Action
91.198.2.65
520516
95.91.81.184
57627
66.249.78.140
12830
91.198.2.70
12144
66.249.67.224
9596
66.249.67.237
9332
66.249.67.211
9064
91.198.2.112
5728
66.249.79.154
5118
66.249.79.186
4799
TOP USERAGENTS
"Mozilla/5.0 (Macintosh; Intel Mac OS X 10_9_3) AppleWebKit/537.76.4 (KHTML, like Gecko) Version/7.0.4 Safari/537.76.4" (258810)
"Mozilla/5.0 (Windows NT 6.1; WOW64; rv:29.0) Gecko/20100101 Firefox/29.0" (248034)
"Mozilla/5.0 (Macintosh; Intel Mac OS X 10_9_3) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/35.0.1916.153 Safari/537.36" (45837)
"Mozilla/5.0 (compatible; Googlebot/2.1; +http://www.google.com/bot.html)" (34362)
"Googlebot-Image/1.0" (17616)
300000
250000
200000
150000
100000
50000
0
RESPONSE
200 (619440)
304 (25959)
302 (8089)
403 (4614)
301 (4123)
500 (583)
404 (446)
401 (20)
400 (2)
200
93%
TRAFFIC IN BYTES
View
Zoom Out
(670717) http_bytes mean per 1h | (670717 hits)
16 KB
12 KB
8 KB
4 KB
0 B
12:00 10-12
00:00 10-13
12:00 10-13
00:00 10-14
12:00 10-14
00:00 10-15
12:00 10-15
00:00 10-16
12:00 10-16
00:00 10-17
12:00 10-17
00:00 10-18
12:00 10-18
00:00 10-19

DEMO

NETWAYS®

MONITORING ELK-STACK

Monitoring of the ELK-Stack

- Availabillity of services and resources
  - Shipment
  - Caching and Indexing
  - Storage
- Realtime monitor for Elasticsearch

NODE PING

- http://berk-logstash.demo.netways.de/es/


```
{
    "status" : 200,
    "name" : "Richard Fisk",
    "cluster_name" : "elasticsearch",
        "version" : {
        "number" : "1.4.5",
        "build_hash" : "2aaf797f2a571dcb779a3b61180afe8390ab61f9",
        "build_timestamp" : "2015-04-27T08:06:06Z",    "build_snapshot" :
        false,
        "lucene_version" : "4.10.4”
    },
    "tagline" : "You Know, for Search”
}
```

ElasticHQ

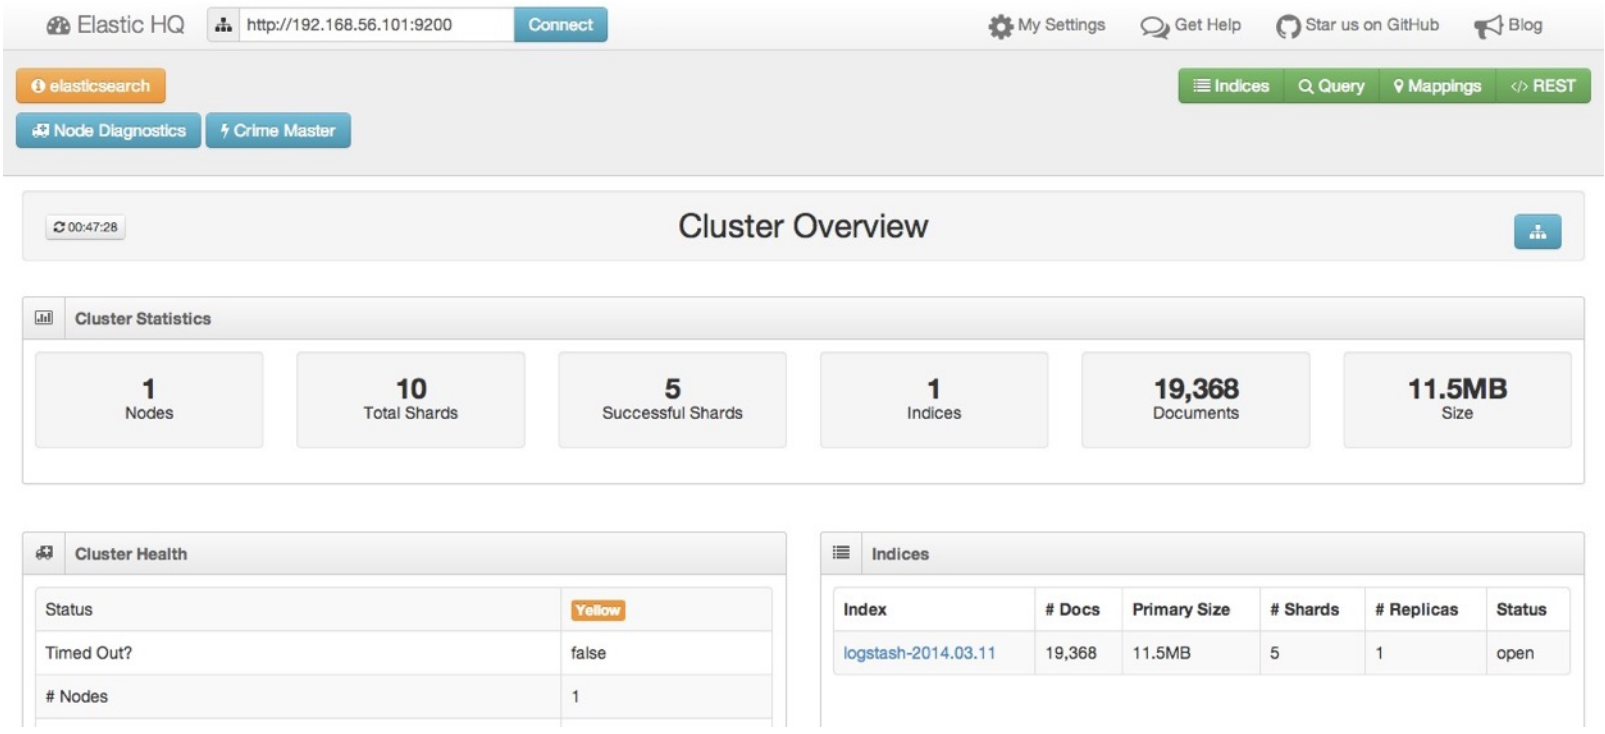
Elastic HQ
http://192.168.56.101:9200
Connect
My Settings
Get Help
Star us on GitHub
Blog
elasticsearch
Indices
Query
Mappings
REST
Node Diagnostics
Crime Master
00:47:28
Cluster Overview
Cluster Statistics
1
Nodes
10
Total Shards
5
Successful Shards
1
Indices
19,368
Documents
11.5MB
Size
Cluster Health
Status
Yellow
Timed Out?
false
Nodes
1
Indices
Index
Docs
Primary Size
Shards
Replicas
Status
logstash-2014.03.11
19,368
11.5MB
5
1
open

DEMO

NETWAYS®

INTEGRATION NAGIOS, ICINGA AND CHECK_MK

Realtime Loganalysis

- Analyse various in sources in realtime
- Check for patters and states
  - Facilitites
  - Regex
  - Programs
- Submission as passive events

Overview Logstash and Nagios

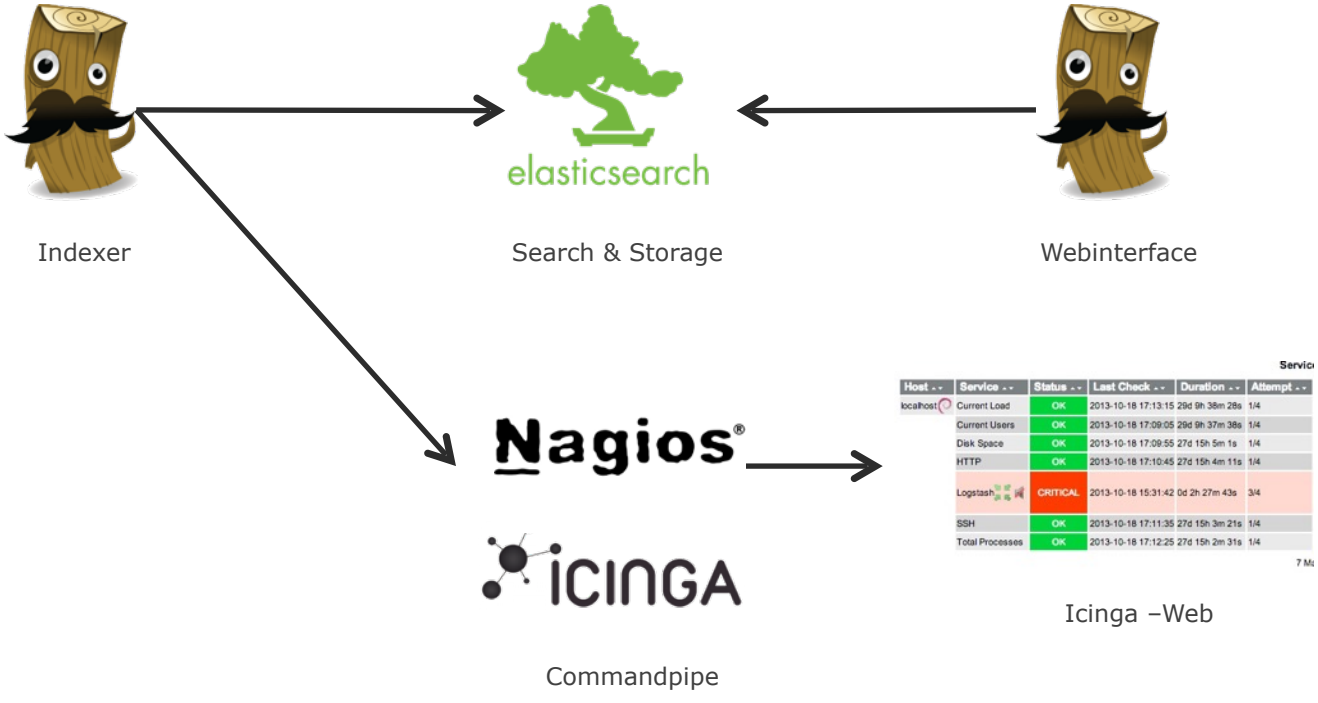

Indexer

Search & Storage

Webinterface

Commandpipe

Icinga –Web

Logstash - Nagios

- Configuration for Nagios-Alert

```
input {
  …
}


filter {
  if [type] == "syslog" {
    grok {match => [ "message", "%{SYSLOGBASE}" ] }
    grep {
      match => [ "message", "Error" ]
      drop => false
      add_tag => "nagios-update"
      add_field => [
      # "nagios_host", "%{@source_host}",
      "nagios_host", "localhost",
      "nagios_service", "Logstash",
      "nagios_level", "2”]
    }}}


output {
  elasticsearch {host => "127.0.0.1”}
  nagios {
    commandfile => "/var/lib/icinga/rw/icinga.cmd"
  }}
```

Overview Check_MK

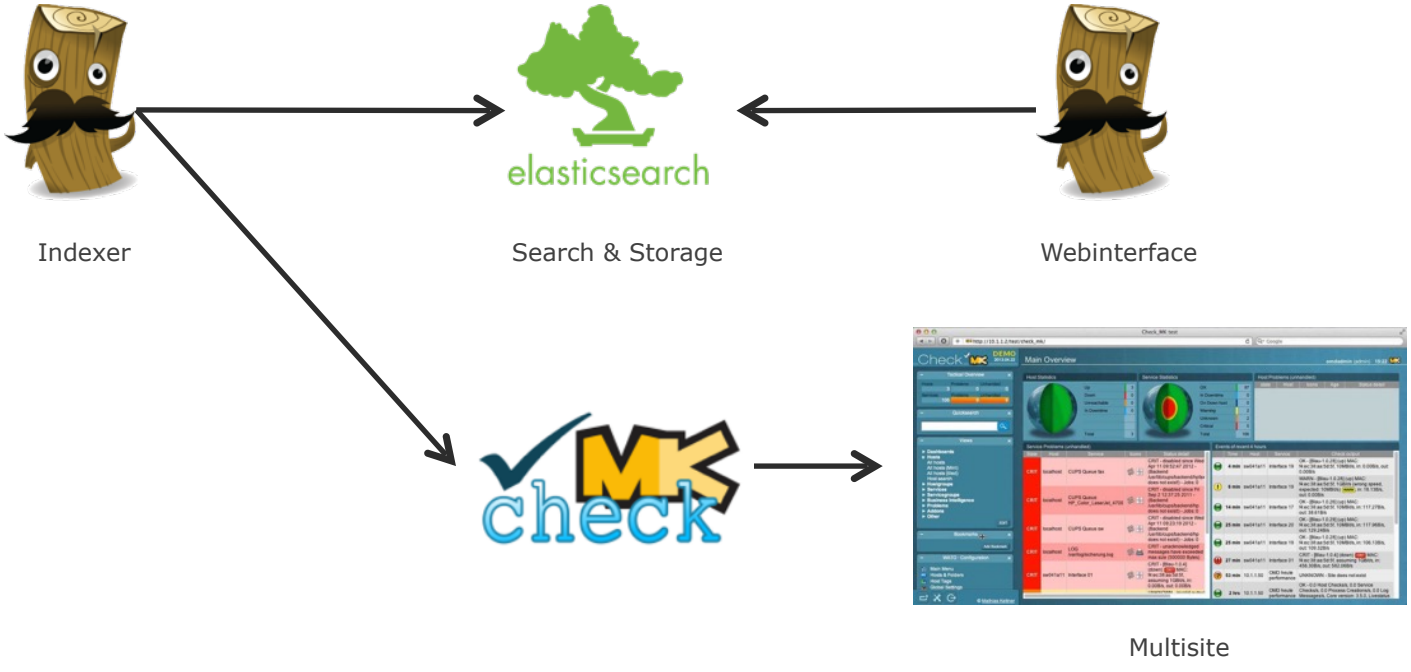

Using the integrated Syslog-Server

- Enable the integrated Syslog server
  - omd config set MKEVENTD_SYSLOG on
- Since version 1.2.3i2 TCP is also available

Logstash -> Syslog

- Configuration for Syslog output

```
input {
  …
}

filter {
  if [type] == "syslog" {
    grok {match => [ "message", "%{SYSLOGBASE}" ] }
    grep {
      match => [ "message", "Error" ]
      drop => false
      add_tag => ”checkmk”
}}}

output {
  elasticsearch {host => "127.0.0.1”}
  syslog{
    facility => “local0”
    host => 192.168.1.1
    port => tcp
    severity => “critical”
  }}
```

CONCLUSION

Conclusion ELK-Stack

- Support for a huge number of APIs and programs
- Scalable storage backend with Elasticsearch
- Flexible Query-Interface with Kibana
- Highly integrable in all popular stacks
- Collect **everything** and analyse **later**!

NETWAYS®

Q&A

THANK YOU

www.netways.de

blog.netways.de

www.netways.org


netways


netways


netways

NETWAYS®