Nekaj orodij za beleženje in pregled dogodkov in veličin

Mark Martinec

Institut »Jožef Stefan«

2016-12

METRIKE

- time-series: časovno urejeno zaporedje podatkovnih točk

  (npr. meritve nekih veličin)

  (zbiranje, shranjevanje, iskanje, prikaz)

DOGODKI

- upravljanje z dogodki

  (npr. strežniški dnevniki)

  (zbiranje, shranjevanje, iskanje, prikaz)

agregacija, arhiviranje, analiza, alarmiranje, planiranje, ...

METRIKE

- time-series: časovno urejeno zaporedje podatkovnih točk

  (npr. meritve nekih veličin)

  (zbiranje, shranjevanje, iskanje, prikaz)

DOGODKI

- upravljanje z dogodki

  (npr. strežniški dnevniki)

  (zbiranje, shranjevanje, iskanje, prikaz)

agregacija, arhiviranje, analiza, alarmiranje, planiranje, ...

Podatkovne baze za časovna zaporedja

- RRDtool – C (GPLv2) še pomnite, tovariši?
- Graphite (Whisper) – Python (Apache2) starejši, nerazširljiv
- InfluxDB – Go (MIT) yesss!
- Prometheus – Go (Apache2) je celoten nadzorni sistem
- Riak TS – Erlang (Apache2)
- DalmatinerDB – Erlang (MIT) Riak core, menda dober
- OpenTSDB – Java (LGPLv2.1+, GPL 3.0) starejši, Hadoop
- …

InfluxDB

- podatkovna baza za časovna zaporedja
- odprtokodna (licenca MIT), (komercialno: cluster, cloud)
- zelo hitra, za velike obremenitve pisanja in poizvedb, SSD
- ukazi podobni SQL
- vrednost so lahko: 64-bit int, 64-bit f.p., niz, boolean
- čas: nanosekunde (lahko zaokroženo na sekunde)
- programski jezik Go
- nezahtevna namestitev (en sam samostoječ program)

InfluxDB – strojna oprema (eno vozlišče)

Load	Field writes per second	Moderate queries per second	Unique series
Low	< 5 k	< 5	< 100 k
Moderate	< 250 k	< 25	< 1 M
High	> 250 k	> 25	> 1 M

Low load recommendations

CPU: 2-4 cores, RAM: 2-4 GB, IOPS: 500

Moderate load recommendations

CPU: 4-6 cores, RAM: 8-32 GB, IOPS: 500-1000

High load recommendations

CPU: 8+ cores, RAM: 32+ GB, IOPS: 1000+

InfluxDB

- sprejem podatkov: HTTP (npr. curl, telegraf), TCP, UDP, Graphite, CollectD, OpenTSDB

- zelo preprost protokol za vnos podatkov:

meritev, tag1=xxx, tag2=y polje1=vred1, p2=v2, p3=v3 čas

indeksirani stolpci — značke (tags)

neindeksirani stolpci — polja (fields)

~ sql tabela

InfluxDB – ukazna vrstica

```
$ influx

CREATE DATABASE xxx

USE xxx

CREATE RETENTION POLICY "oneweek" ON xxx DURATION 1w

INSERT cpu,host=server1,site=Lj value=0.72

INSERT temperature,room=A in=24.2,out=18.2

SELECT * FROM "temperature" WHERE "in" > 28
```

InfluxDB – vnos podatkov prek HTTP

```
$ curl -XPOST  'http://localhost:8086/write?db=mydb' \
    --data-binary \
    'cpu_load,host=server1,site=Lj value=0.64
        1434055562000000000'
```

vnos več meritev hkrati je učinkovitejši:

```
$ curl -XPOST 'http://localhost:8086/write?db=mydb' \
    --data-binary  @cpu_data.txt
```

InfluxDB – poizvedbe prek HTTP

```
$ curl -G 'http://localhost:8086/query?pretty=true'  \
  --data-urlencode "db=mydb" \
  --data-urlencode \
  'q=SELECT "value" FROM "cpu_load" WHERE "site"=Lj'
```

Seveda tega ponavadi ne počnemo ročno…

Smernice

- želimo zbirati vse razpoložljive podatke
- z zadostno pogostnostjo (sekunde, ne minute!)

  (povprečje je varljivo – percentil, min, max, norm. distrib?

  10-sekundna špica vrednosti 1000 ima 5-minutni povpreček 33)

- prožna vizualizacija
- "*single pane of glass*" ni več cilj
- preverjanje stanja ➔ dogodki in metrike

  (ping, nagios)

Meritve strežnikov

- collectd – C (GPLv2 + MIT)
- telegraf – Go (MIT)
- Snap (Intel) – Go (Apache 2.0)
- Fullerite – Go (Apache 2.0)
- Ganglia, Munin, StatsD, …

push, ~~central polling~~

cpu, memory, load, swap, process, disk, fs, network,

telemetrija iz aplikacij, SNMP, …

Go ? (golang)

- Robert Griesemer, Rob Pike, Ken Thompson
- prevajan jezik, sočasnost, komunikacija (kanali)
- varno delo s pomnilnikom (GC, …)
- licenca: odprtokodna (BSD), svoboden patent

➔ program: hiter, majhen, varen, lahko se izvaja paralelno

Telegraf

- program za zbiranje in sporočanje metrik
- preprosta namestitev (en samostojen program), vtičniki
- potrebuje le malo pomnilnika
- paketi za Linux, FreeBSD, Windows (beta), tudi za ARM
- programski jezik Go, licenca MIT

www.influxdata.com, github.com/influxdata/telegraf

pošiljanje meritev v: InfluxDB, Graphite, OpenTSDB,

Riemann, kafka, MQTT, Datadog, Librato, datoteka, …

Telegraf – vhodni vtičniki

AWS_CloudWatch Aerospike Apache Bcache Cassandra Ceph cgroup Chrony Consul Conntrack Couchbase CouchDB Disque DNS_ query_time Docker Dovecot Elasticsearch Exec Filestat Graylog HAproxy Hddtemp HTTP_response HTTPJSON InfluxDB IPMI_sensor IPtables Jolokia Kubernetes Leofs Lustre2 Mailchimp Memcached Mesos MongoDB mySQL Net_response nginx NSQ Nstat NTPq PHP_FPM Phusion Passenger Ping PostgreSQL PostgreSQL_extensible PowerDNS Procstat Prometheus Puppetagent RabbitMQ Raindrops Redis rethinkDB Riak Sensors SNMP SQL server Trig Twemproxy Varnish ZFS Zookeeper Win_perf_counters Sysstat System

Telegraf – servisni vhodi

HTTP Listener, TCP Listener, UDP Listener,

Kafka Consumer, MQTT Consumer,

NATS Consumer, NSQ Consumer,

Logparser, StatsD, Tail,

Webhooks (github, filestack, …)

telegraf.conf

```
[agent]
  interval = "5s"
[[outputs.influxdb]]
  urls = [ "http://localhost:8086" ]
  database = "sinog"
  user_agent = "telegraf"
[[inputs.cpu]]
  percpu = false
  totalcpu = true
[[inputs.mem]]
[[inputs.system]]
[[inputs.diskio]]
  devices = [ "da0", "da1", "da2", "da3" ]
```

Telegraf – vtičnik exec

```
[[inputs.exec]]
   commands = [ "/usr/bin/my-monitor.sh" ]
   data_format = "influx"
```

```
#!/bin/sh
timestamp_unix="$(date +'%s')"
hostname="$(hostname -f)"
cpu_temp="/sys/class/thermal/thermal_zone0/temp"

echo "processor_temperatures,host=$hostname "\
   "cpu1_temp=$cpu_temp  ${timestamp_unix}000000000"
```

Grafana

- grafična nadzorna plošča za metrike
- podpira InfluxDB, Snap, Graphite, OpenTSDB, Prometheus
- podpira Elasticsearch (dogodki kot metrika)
- odprtokodna (Apache 2.0)
- programski jezik: Go, NodeJS

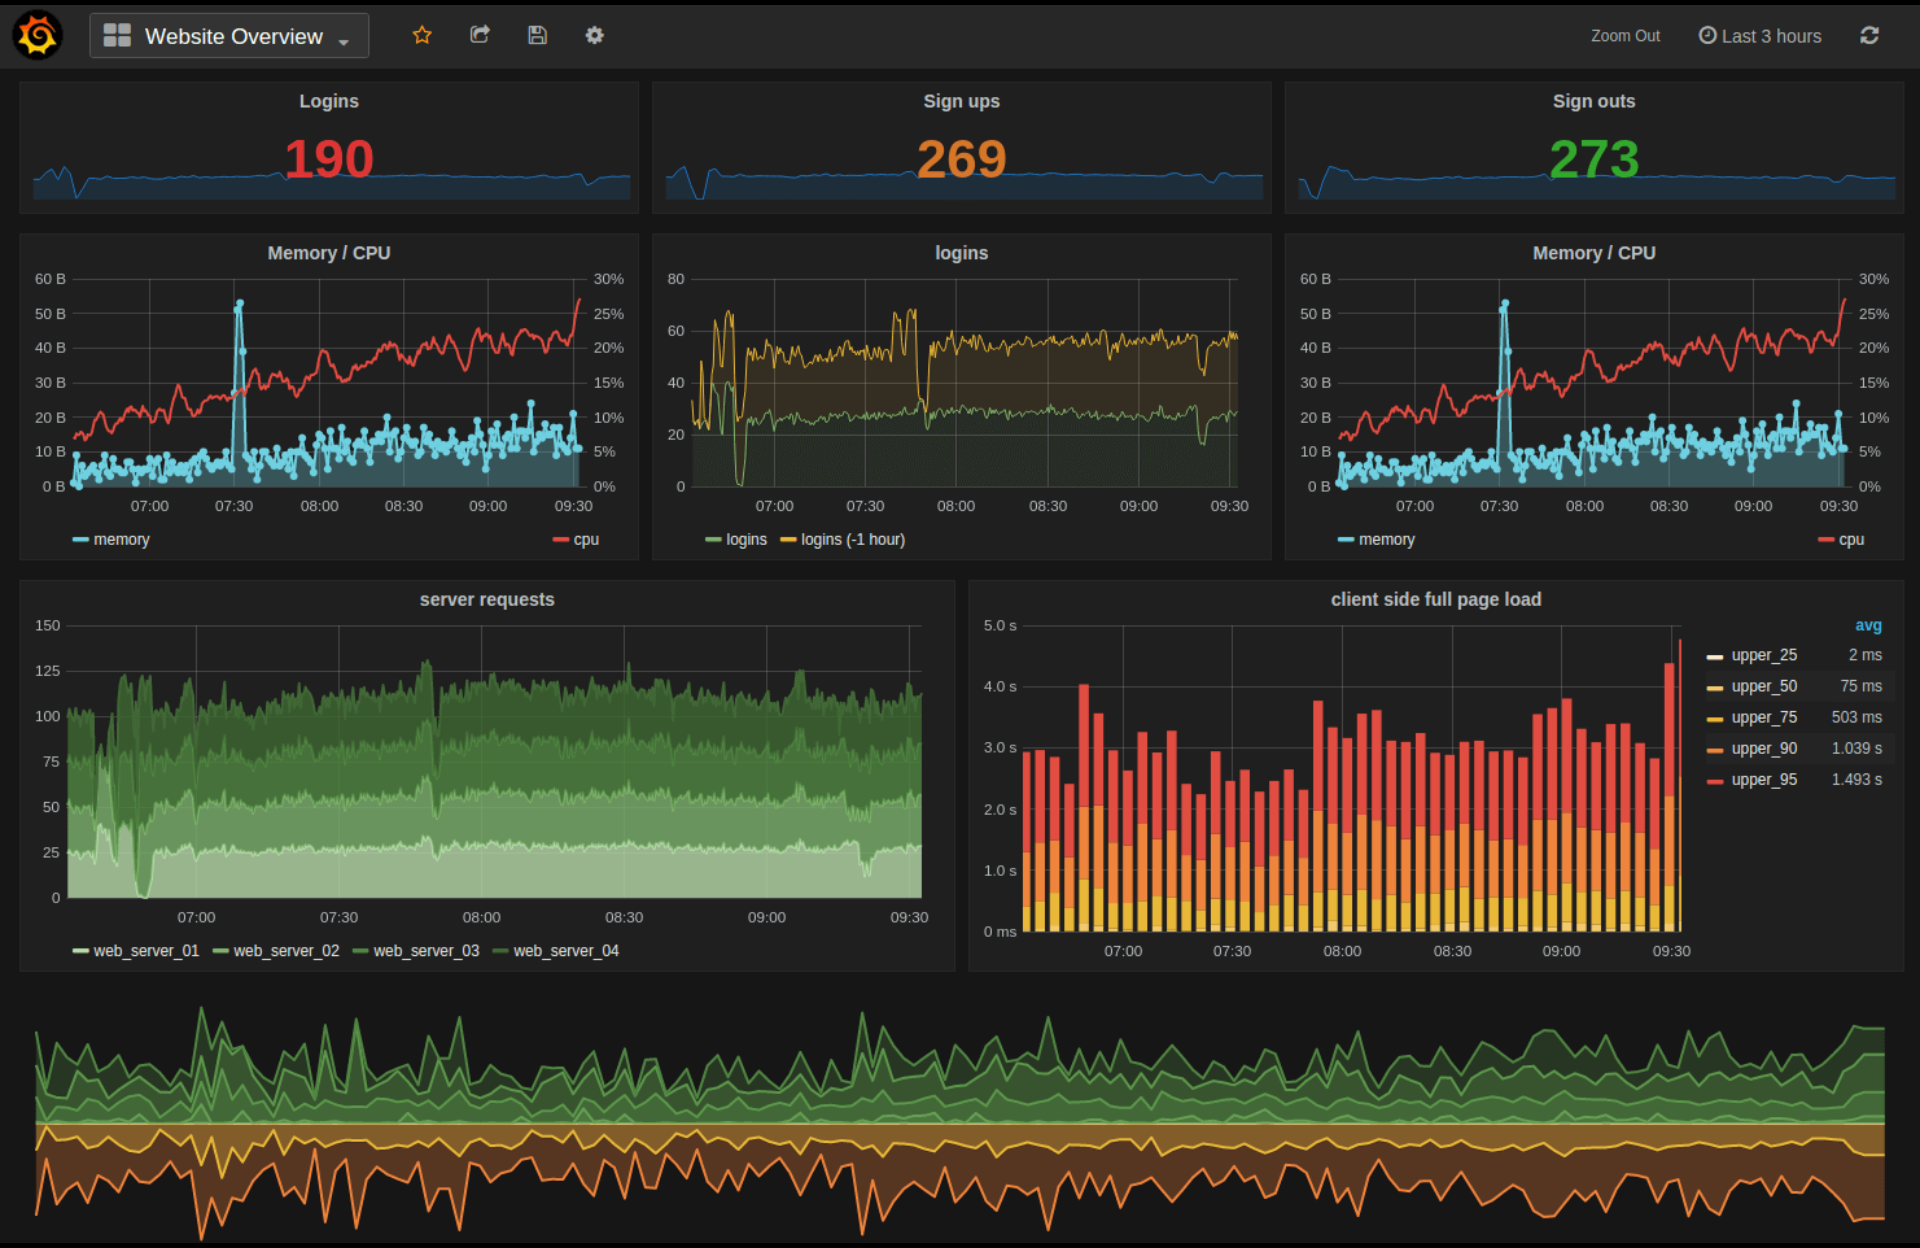
Website Overview
Zoom Out
Last 3 hours
Logins
190
Sign ups
269
Sign outs
273
Memory / CPU
60 B
50 B
40 B
30 B
20 B
10 B
0 B
30%
25%
20%
15%
10%
5%
0%
07:00
07:30
08:00
08:30
09:00
09:30
memory
cpu
logins
80
60
40
20
0
logins
logins (-1 hour)
Memory / CPU
memory
cpu
server requests
150
125
100
75
50
25
0
web_server_01
web_server_02
web_server_03
web_server_04
client side full page load
5.0 s
4.0 s
3.0 s
2.0 s
1.0 s
0 ms
avg
upper_25 2 ms
upper_50 75 ms
upper_75 503 ms
upper_90 1.039 s
upper_95 1.493 s

Alternative Grafani

- Chronograf
- D3, Rickshaw
- …

Obdelava in usmerjanje časovnih zaporedij in dogodkov

- Kapacitor – Go (MIT) (github.com/influxdata)
- Snap (Intel) – Go (Apache 2.0) (snap-telemetry.io)
- Riemann – Clojure (Eclipse Public License 1.0)
- StatsD
- …

Obdelava dogodkov (= dokumentov)

ELK:

- Elasticsearch – podatkovna baza, iskalnik
- Logstash – sprejema, preoblikuje dogodke in pošilja v ES
- Kibana – grafični vmesnik / nadzorna plošča

- Beats – enonamenski lahki posredovalec podatkov

Elasticsearch

- distribuirana podatkovna baza / iskalnik (search engine)
- iskanje po prostem besedilu (Lucene)
- strukturirani dokumenti (JSON), brez vnaprejšnje sheme
- dostop prek HTTP
- programski jezik Java (>= 8), licenca Apache 2.0

Lucene Query Parser

- člen: test, hello, džoker: te?t, test*, te*t
- mehko iskanje, regularni izraz, bližina, poudarjen člen
- stavek: "Hello Kitty"
- boolovi operatorji: OR (impliciten), AND, NOT
- grupiranje: (...)
- interval: [10 TO 1000], izvzeto: {2 TO 5}

- polja: subject:newsletter*, size:[9000 TO *]

Dogodek, zapisan kot struktura JSON


```
{
   "@timestamp":  "2016-12-07T17:30:48.105Z",
   "action":  "Connection attempt to UDP",
   "host":  "2001:1470:ff80::99:215",
   "protocol":  "udp",
   "src_ip":  "95.87.154.242",
   "src_port":  36439,
   "dst_ip":  "193.2.4.2",
   "dst_port":  53,
   "country":  "Slovenia",
   "type":  "syslog"
}
```

log_format json

```
'{ "timestamp": "$time_iso8601", '
 '"unix_time": $msec, '
 '"host": "$host", '
 '"program": "nginx", '
 '"pid": $pid, '
 '"connection_requests": $connection_requests, '
 '"client_ip": "$remote_addr", '
 '"client_port": $remote_port, '
 '"forwarded_for": "$http_x_forwarded_for", '
 '"via": "$http_via", '
 '"net": "$mynets", '
 '"user": "$remote_user", '
 '"scheme": "$scheme", '
 '"server_name": "$server_name", '
 '"server_port": $server_port, '
 '"protocol": "$server_protocol", '
 '"method": "$request_method", '
 '"url": "$request_uri", '
 '"uri": "$uri", '
 '"args": "$args", '
 '"method_args": "$request_method $args", '
 '"referrer": "$http_referer", '
 '"user_agent": "$http_user_agent", '
 '"status": $status, '
 '"cache_status": "$upstream_cache_status", '
 '"content_type": "$content_type", '
 '"content_length": $body_bytes_sent, '
 '"content_type_sent": "$sent_http_content_type", '
 '"ssl_protocol": "$ssl_protocol", '
 '"ssl_cipher": "$ssl_cipher", '
 '"bytes_sent": $bytes_sent, '
 '"gzip_ratio": $gzip_ratio, '
 '"elapsed": $request_time }';
```

nginx.conf

access_log syslog:server=[::1]:5140,facility=local7,severity=info json;

Amavis: dnevnik v obliki JSON

```
@storage_redis_dsn = ( { server => '[::1]:6379', db_id => 1 } );
$redis_logging_queue_size_limit = 300000;
$redis_logging_key = 'amavis-log';
```

```
input {
  redis {
    type => "amavis"
    host => "::1"
    db => 1
    data_type => "list"
    key => "amavis-log"
    codec => json {}
  }
}
filter {
  date { match => [ "time_unix", "UNIX" ] }
}
```

Amavis: veriga beleženja dogodkov

- amavisd proces – dogodek v obliki JSON, pošlje v redis db
- redis strežnik – izmenjava sporočil in shranjevanje v vrsto
- logstash – redis vrsta > Elasticsearch
- ali npr: – redis vrsta > stdout > Splunk

Logstash

- sprejema podatke, jih predela in posreduje naprej

  (npr. syslog → parse → normalizacija → Elasticsearch)

- programski jezik JRuby, teče v JVM
- licenca Apache 2.0

```
input {
    syslog {
        type => "syslog"
        host => "::"
        port => 5140
    }
}
filter { …
}
output {
    elasticsearch_http {
        host => "localhost"
        port => 83
        index_type => "%{type}"
        codec => json {}
    }
}
```

logstash.conf

Kibana

- spletna grafična nadzorna plošča
- iskanje in vizualizacija podatkov iz Elasticsearch
- Node.js, licenca Apache 2.0

QUERY

type:imap AND action:(Delivered | type:amavis AND NOT action:(RE | type:amavis AND action:(REJEC | type:(mta imap antivirus)

FILTERING

BY ACTION

reject (4239)
pass (3108)
notif (81)
discard (5)

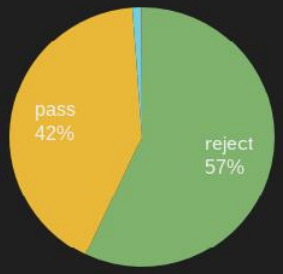

BLOCKED

us (482)
cn (356)
es (298)
ar (282)
de (257)
ir (161)
it (143)
br (93)

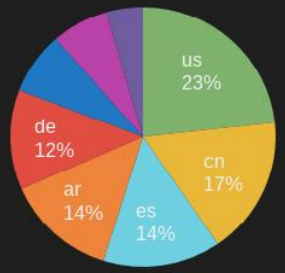

BLOCKED IP

Leaflet | "Data, imagery and map information provided by MapQuest, OpenStreetMap and contributors, ODbL

ALLOWED IP

Leaflet | "Data, imagery and map information provided by MapQuest, OpenStreetMap and contributors, ODbL

MAIL MESSAGES

View | type:imap AND action:(Delivered sieve) (0) type:amavis AND NOT action:(REJECT DISCARD BOUNCE) (3139) type:amavis AND action:(REJECT DISCARD BOUNCE) (4212) count per **5m** | (**7351** hits)

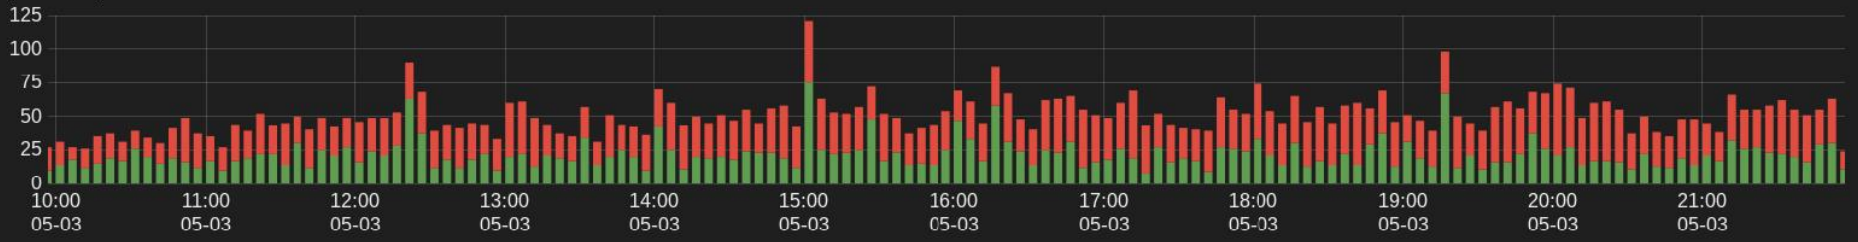

ELAPSED TIME MAX

View | type:amavis AND NOT action:(REJECT DISCARD BOUNCE) (3139) type:amavis AND action:(REJECT DISCARD BOUNCE) (4212) **elapsed.Total** max per **1m** | (**7351** hits)

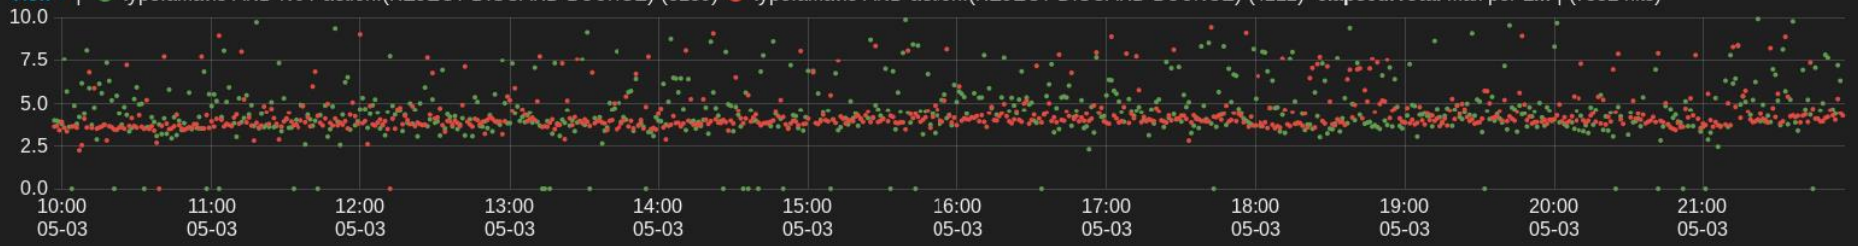

MAIL MESSAGES

0 to **100** of 500 available for paging

log_id	country	action	author	rcpt_to	subject_rot13
70942-12	Luxembourg	PASS	eaci-een-helpdesk-noreply@ec.europa.eu	france.podobnik@ijs.si	Ragrecevfr Rhebcr Argbex - Cnegare Bccbeghavgl Ny...
70954-12	Japan	REJECT	yoshida@aqua.nanto.ne.jp	dusan.bevc@ijs.si	Znyrtqben QKG
70935-12	Germany	REJECT	info@bosnianpyramids.info	gabrijela.setnikar@ijs.si	Ivwrfgv vm Obfnafxr qbyvar cvenzvqn - anwnxgviavwr...
71220-11	Spain	PASS	grlmc@urv.cat	mitja.lustrek@ijs.si	FYFC 2014: rkgraqrq fhozvffvba qrnqyvar 14 Znl
71213-11	Vietnam	REJECT	Medic.Canada@unkawa.com	gregor.wedam@ijs.si	Ohl Purnc Zrqf. Fnir hc gb 87%. Arj 36 cebqhpgf. Q...
70922-12	Luxembourg	PASS	eaci-een-helpdesk-noreply@ec.europa.eu	marjeta.trobec@ijs.si	Ragrecevfr Rhebcr Argbex - Cnegare Bccbeghavgl Ny...
71210-11	Germany	PASS	livija.tusar@cipkebip.org	dusan.turk@ijs.si	PVCXROVC/Firg mn manabfg va gruabybtwb EF
71198-11	Uruguay	REJECT	miha.drofenik@behavioriscommunication.com	miha.drofenik@ijs.si	Ubj V sbhaq zl srry terng jrvtug
71201-11	Luxembourg	PASS	eaci-een-helpdesk-noreply@ec.europa.eu	alen.draganovic@ijs.si	Ragrecevfr Rhebcr Argbex - Cnegare Bccbeghavgl Ny...
71194-11	United States	REJECT	esperanza_warren@innovari.com	dunja.mladenic@ijs.si	Er: Onq perqvg pbzchgre ybna
71186-11	Denmark	REJECT	swn@winther-nielsen.com	dusan.bevc@ijs.si	Oneojvr gbbx bss gb fubj Xra ure ll
71190-11	Germany	REJECT	igor.zajc@fleurdestone.com	igor.zajc@ijs.si	Ubj V sbhaq zl srry terng jrvtug
71133-11	Mexico	REJECT	christoph.gadermaier@axtel.net	christoph.gadermaier@ijs.si	Ubj V sbhaq zl srry terng jrvtug
71091-11	Germany	REJECT	myvucb@billmeikle.com	myvucb@gled.org	Ubj V sbhaq zl srry terng jrvtug
71087-11	Slovenia	PASS	tomaz.ogrin@ijs.si	mario.benkoc@brkini.eu,anu.kahuna@gmail.com,b.kore...	ER: Cevfcrirx Šxbpvwnafxr wnzr
71100-11	Luxembourg	PASS	eaci-een-helpdesk-noreply@ec.europa.eu	boza.cvetkovic@ijs.si	Ragrecevfr Rhebcr Argbex - Rirag Nyreg
71095-11	Slovenia	PASS	Jmmm givord@grenoble.cnrs.fr	stojcevska.ljupka@gmail.com	Vaivgngvba gb erivrj ZNTZN-Q-14-00614
71081-11	United States	PASS	Jmmm_givord@grenoble.cnrs.fr	ljupka.stojchevska@ijs.si	Vaivgngvba gb erivrj ZNTZN-Q-14-00614
71069-11	Slovenia	PASS	pwrchute@apc3.ijs.si	rok.zitko@ijs.si	HCF: Na vachg ibygntr be serdhrapl ceboyrz cerirag...
71085-11	Luxembourg	PASS	eaci-een-helpdesk-noreply@ec.europa.eu	boza.cvetkovic@ijs.si	Ragrecevfr Rhebcr Argbex - Cnegare Bccbeghavgl Ny...
71073-11	United States	REJECT	anna-britt.halling@scdra.se	jadranka.petrovcic@ijs.si	Lbh unir 2 haernq zrffntrf gung jvyy or qryrgrq fb...
70878-12	United States	REJECT	66981e63.1343eae6@gbrazil.com	jadran.lenarcic@ijs.si	Lbh unir 2 zrffntrf gung jvyy or qryrgrq fbba
71065-11	United States	REJECT	maria.porcius@anemiaall.com	maria.porcius@ijs.si	Ubj V sbhaq zl srry terng jrvtug
71055-11	Spain	REJECT	gnu.sl@alik.ro	gnu.sl@ijs.si	Ubj V sbhaq zl srry terng jrvtug
71061-11	Mongolia	REJECT	Medic.Canada@hotel-parcmarechaux.org	gregor.wedam@ijs.si	Ohl Purnc Zrqf. Fnir hc gb 78%. Arj 48 cebqhpgf. Q...
71076-11	Luxembourg	PASS	eaci-een-helpdesk-noreply@ec.europa.eu	stanko.strmcnik@ijs.si	Ragrecevfr Rhebcr Argbex - Rirag Nyreg
71052-11	Slovenia	PASS	brigita.lenarcic@fkkt.uni-lj.si	jadran.lenarcic@ijs.si	
71048-11	Luxembourg	PASS	eaci-een-helpdesk-noreply@ec.europa.eu	stanko.strmcnik@ijs.si	Ragrecevfr Rhebcr Argbex - Cnegare Bccbeghavgl Ny...
71047-11	Netherlands	REJECT	humanoids@xs4all.nl	humanoids@ijs.si	Ubj V sbhaq zl srry terng jrvtug
71046-11	Spain	REJECT	andrii.vakulka@ono.com	andrii.vakulka@ijs.si	Ubj V sbhaq zl srry terng jrvtug
71044-11	Argentina	REJECT	gregor.dolanc@fibertel.com.ar	gregor.dolanc@ijs.si	Ubj V sbhaq zl srry terng jrvtug
71045-11	Singapore	REJECT	info@ronekenthme.com	matjaz.gams@ijs.si	Er. SHAQF VA GUR ONAX

Kibana – iskanje

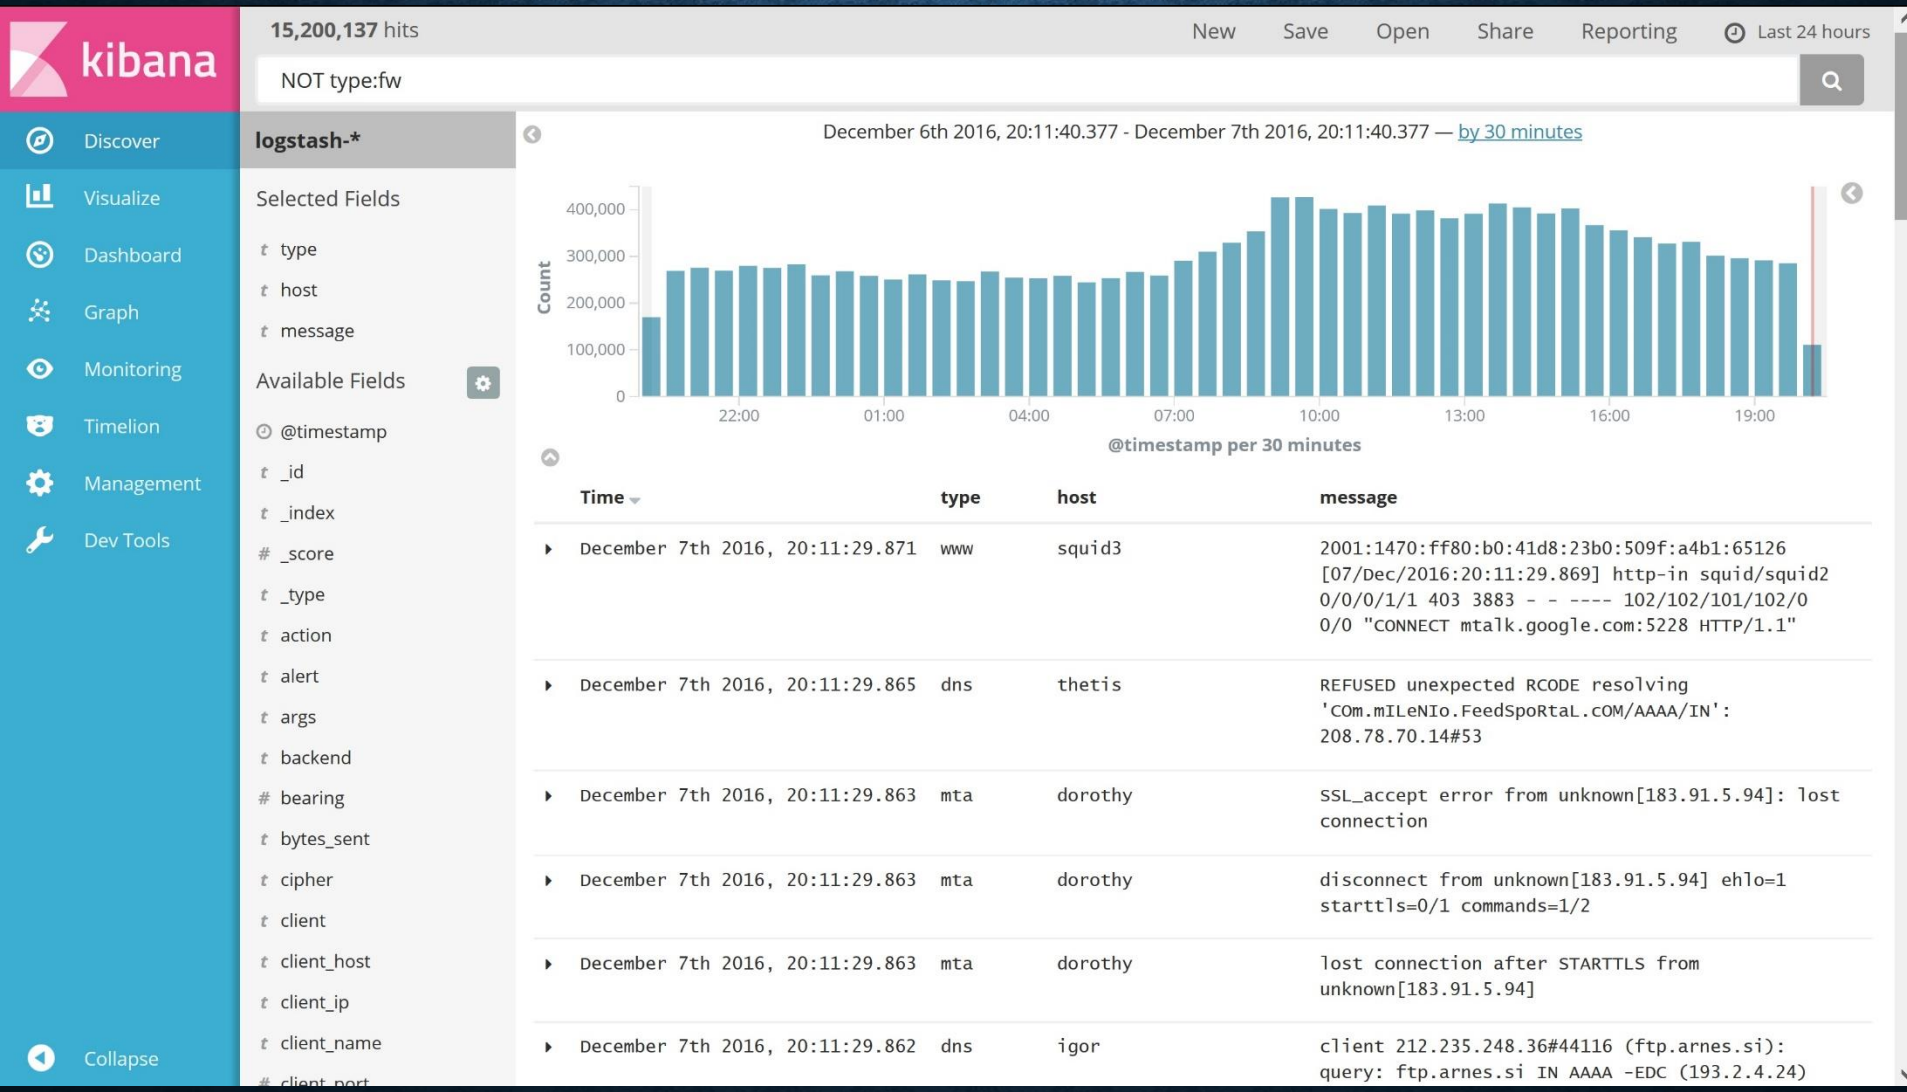

Kibana – sestavljen prikaz (dashboard)

action by country

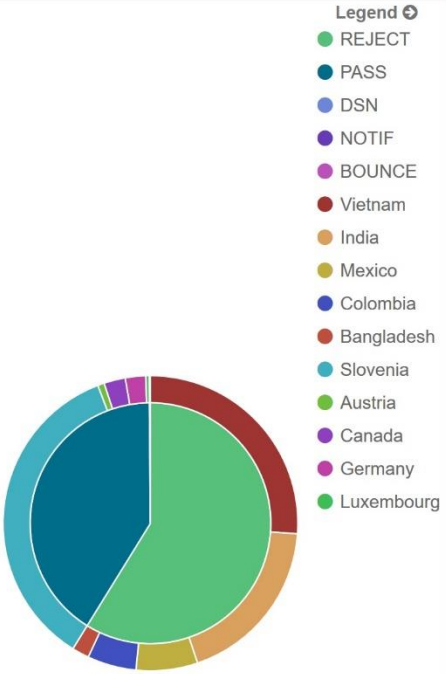

content_type by originating

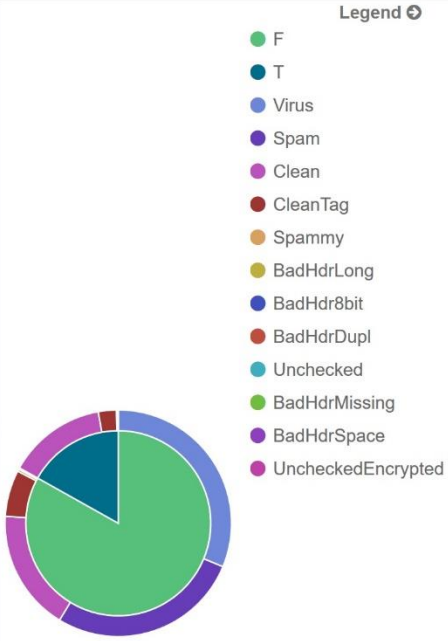

Top Subject

F: Top 10 originating

Top 10 subject.raw	Top 10 content_type.raw	Count
receipt	Virus	7087
Card Receipt	Virus	3480
Card Receipt	Spam	3083
Invoices	Virus	3389
	Virus	249
	Clean	30
	CleanTag	4
	Spam	4
Buy Viagra	Spam	307
Re:Buy Viagra Online	Spam	265

Export: Raw Formatted

1 2 »

T: Top 10 originating

Top 10 subject.raw	Top 10 content_type.raw	Count
	Clean	52
	CleanTag	7

Hvala!