

Introduction to ELK stack

– 巨量資料處理、搜尋、及分析工具介紹 –

計資中心網路組 邵喻美

madeline@ntu.edu.tw

Topics

- Why big data tool for network traffic and log analysis
- What is ELK stack, and why choose it
- ELK stack intro
- ELK use cases
- Implementation of ELK on network & account anomaly detection

Network operation and security management issues

- Lots of users
 - Faculty & staff & students → more than 40000 users on campus
- Lots of systems
 - Routers, firewalls, servers....
- Lots of logs
 - Netflow, syslogs, access logs, service logs, audit logs....
- Nobody cares until something go wrong....

What kind of data we're dealing with...

```
u_ex160714.log:128: 2016-07-14 00:00:02 W3SVC1 EXCHCAS01 172.16.6.163 GET /owa/forms/premium/StartPage.aspx &Initial+Bud
14T00:00:16.642Z,000000000069E4CE,4,172.16.6.163:995,207.46.8.221:51209,b98502058,0,4,8280,uidl,, "R=ok;Budget=""Conn:0,HangingConn:0,AD:$null/$null/CAS:$null/$null/0%,
14T00:00:17.283Z,000000000069E4CF,0,172.16.6.163:995,122.116.67.113:28801,, -2147483648,0,51,OpenSession,,
14T00:00:17.361Z,000000000069E4CE,5,172.16.6.163:995,207.46.8.221:51209,b98502058,0,4,9483,list,, "R=ok;Rows=911;TotalSize=298858039;Budget=""Conn:0,HangingConn:0,AD:$
14T00:00:17.486Z,000000000069E4CF,1,172.16.6.163:995,122.116.67.113:28801,,0,4,49,capa,,R=ok
14T00:00:17.642Z,000000000069E4CF,2,172.16.6.163:995,122.116.67.113:28801,d96943014,124,31,338,auth,NTLM,"R=ok;RpcC=11;RpcL=78;LdapC=2;LdapL=16;Msg=User:鄧智生:4d7073
14T00:00:17.658Z,000000000069E4CF,3,172.16.6.163:995,122.116.67.113:28801,d96943014,0,4,9,stat,, "R=ok;Rows=0;TotalSize=0;Budget=""Conn:0,HangingConn:0,AD:$null/$null/
14T00:00:17.689Z,000000000069E4CF,4,172.16.6.163:995,122.116.67.113:28801,d96943014,0,4,61,quit,, "R=ok;Budget=""Conn:0,HangingConn:0,AD:$null/$null/1%,CAS:$null/$null
14T00:00:19.048Z,000000000069E4D0,0,172.16.6.163:995,209.85.216.168:30398,, -2147483648,0,51,OpenSession,,
14T00:00:19.486Z,000000000069E4D0,1,172.16.6.163:995,209.85.216.168:30398,,0,14,5,user,b03902108,R=ok
14T00:00:19.658Z,000000000069E4CD,7,172.16.6.163:995,114.136.127.71:41166,b02406038,0,4,61,quit,, "R=ok;Budget=""Conn:0,HangingConn:0,AD:$null/$null/0%,CAS:$null/$null
14T00:00:19.892Z,000000000069E4D0,2,172.16.6.163:995,209.85.216.168:30398,b03902108,203,10,34,pass,****,"R=ok;RpcC=15;RpcL=125;LdapC=3;LdapL=31;Msg=User:盧承德:3cd64
14T00:00:20.095Z,000000000069E4D0,3,172.16.6.163:995,209.85.216.168:30398,b03902108,0,4,49,capa,, "R=ok;RpcL=-1;LdapL=-1;Budget=""Conn:0,HangingConn:0,AD:$null/$null/0
14T00:00:20.298Z,000000000069E4D0,4,172.16.6.163:995,209.85.216.168:30398,b03902108,0,4,12,list,, "R=ok;Rows=0;TotalSize=0;Budget=""Conn:0,HangingConn:0,AD:$null/$null
14T00:00:20.501Z,000000000069E4D0,5,172.16.6.163:995,209.85.216.168:30398,b03902108,0,4,8,uidl,, "R=ok;Budget=""Conn:0,HangingConn:0,AD:$null/$null/0%,CAS:$null/$null/
14T00:00:20.704Z,000000000069E4D0,6,172.16.6.163:995,209.85.216.168:30398,b03902108,0,4,61,quit,, "R=ok;Budget=""Conn:0,HangingConn:0,AD:$null/$null/0%,CAS:$null/$null
14T00:00:21.564Z,000000000069E4D1,0,172.16.6.163:995,209.85.213.27:13985,, -2147483648,0,51,OpenSession,,
14T00:00:21.657Z,000000000069E4D2,0,172.16.6.163:995,207.46.8.221:41211,, -2147483648,0,51,OpenSession,,
14T00:00:22.017Z,000000000069E4D1,1,172.16.6.163:995,209.85.213.27:13985,,0,14,5,user,r00a21067,R=ok
14T00:00:22.220Z,000000000069E4D2,1,172.16.6.163:995,207.46.8.221:41211,,0,25,5,user,b97401037@ntu.edu.tw,R=ok
SmJoKzL2TPMB44UbnInBLDngOdZPMBO4ed93cNFggCwbfBU/LG7Br0HA==" https://mail.ntu.edu.tw/owa/auth/logon.aspx?replaceCurrent=
1&url=https%3a%2f%2fmail.ntu.edu.tw%2fowa%2f mail.ntu.edu.tw 200 0 0 27354 931 1671
```

How we “traditional” system managers treat logs

- Set up one or more log servers for receiving logs from servers/routers/appliances
- Unix commands -- grep + awk + sed + sort + uniq + perl + shell

```
sed '1d' all-$day.log > all-$day-1.log
sed 's/\KHTML\,/KHTML/' all-$day-1.log > all-$day-2.log
awk 'BEGIN { FS=","; OFS="," } { if (($1 !~ /service/) && ($4 !~ /exinstaller/) && ($4 !~ /extest/) && ($6 !~ /fe80/) && ($6 !~ /172.16./) && ($6 !~ /172.28./)) print $1,$2,$3,tolower($4),$5,$6,$7 }' all-$day-2.log > all-$day.log
```

- compute stats and send out report/alert

```
/usr/bin/sort -t "," -k 4 all-$day-geo.out > all-$day-geo-sorted.out
$dir/computeAccount-uniq.pl | /usr/bin/awk '{FS="+"}{print $1}' | /usr/bin/sort | uniq -c | /usr/bin/sort | /usr/bin/awk -v d=$date '{if ($1>1) {print d,$2,$1}}' > stats-$day.out
```

Plain text reports or stats trends webpage

Amount of data....

- Router
 - Netflow – 43GB daily
- Wifi
 - NAT log – 4.8TB daily
 - Auth log
- WAF/Firewall
- Server access logs & events
- Mail server log ~18GB daily
 - POP3 – avg. 7GB daily
 - SMTP – avg. 1.75GB daily
 - Exchange – avg. 140MB daily
 - OWA – avg. 8.4GB daily
 - MessageTrackingLog – avg. 100MB daily

What is ELK, and why choose it


Splunk vs. ELK on Google Trend



One of the leaders in security information and event management (SIEM) market

The ELK Stack is now downloaded 500,000 times every month, making it the world's most popular log management platform

How do Netflix, Facebook, Microsoft, LinkedIn, and Cisco monitor their logs? With ELK.



Why ELK?

- Rapid on-premise (or cloud) installation and easy to deploy
- Scales vertically and horizontally
- Easy and various APIs to use
 - Ease of writing queries, a lot easier than writing a MapReduce job
- Availability of libraries for most programming/scripting languages
 - Elastic offers a host of language clients for Elasticsearch, including Ruby, Python, PHP, Perl, .NET, Java, and Javascript, and more
- Tools availability
- It's free (open source), and it's quick



Logstash is a log pipeline tool that accepts inputs from various sources, executes different transformations, and exports the data to various targets

→ collects and parses logs

Elasticsearch is a NoSQL database that is based on the Lucene search engine

→ indexes and stores the information

Kibana is a visualization layer that works on top of Elasticsearch

→ presents the data in visualizations that provide actionable insights

ELK modules

Open Source —

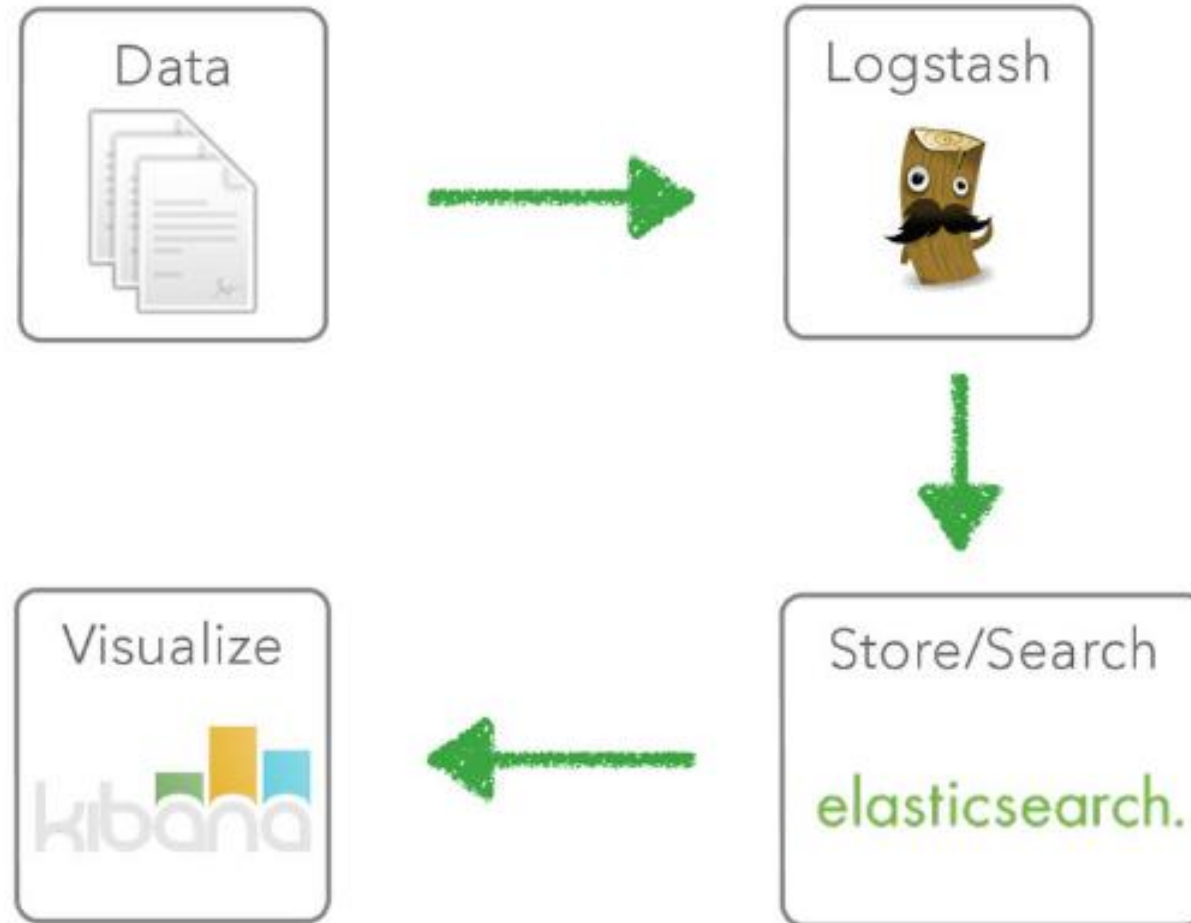
- ElasticSearch
- Logstash
- Kibana
- Beats
 - data shippers – collect, parse & ship

Extension plugins —

- [Alerting \(Watcher\)](#)
 - Proactively monitoring and alerting based on elasticsearch queries or conditions
- [Security \(Shield\)](#)
 - Protect and provide security to elastic stack
- [Monitoring \(Marvel\)](#)
 - Monitor and diagnose health and performance of elastics cluster
- [Graph](#)
 - discover and explore the relationships live in data by adding relevance to your exploration

let's look into ELK stack

The ELK stack



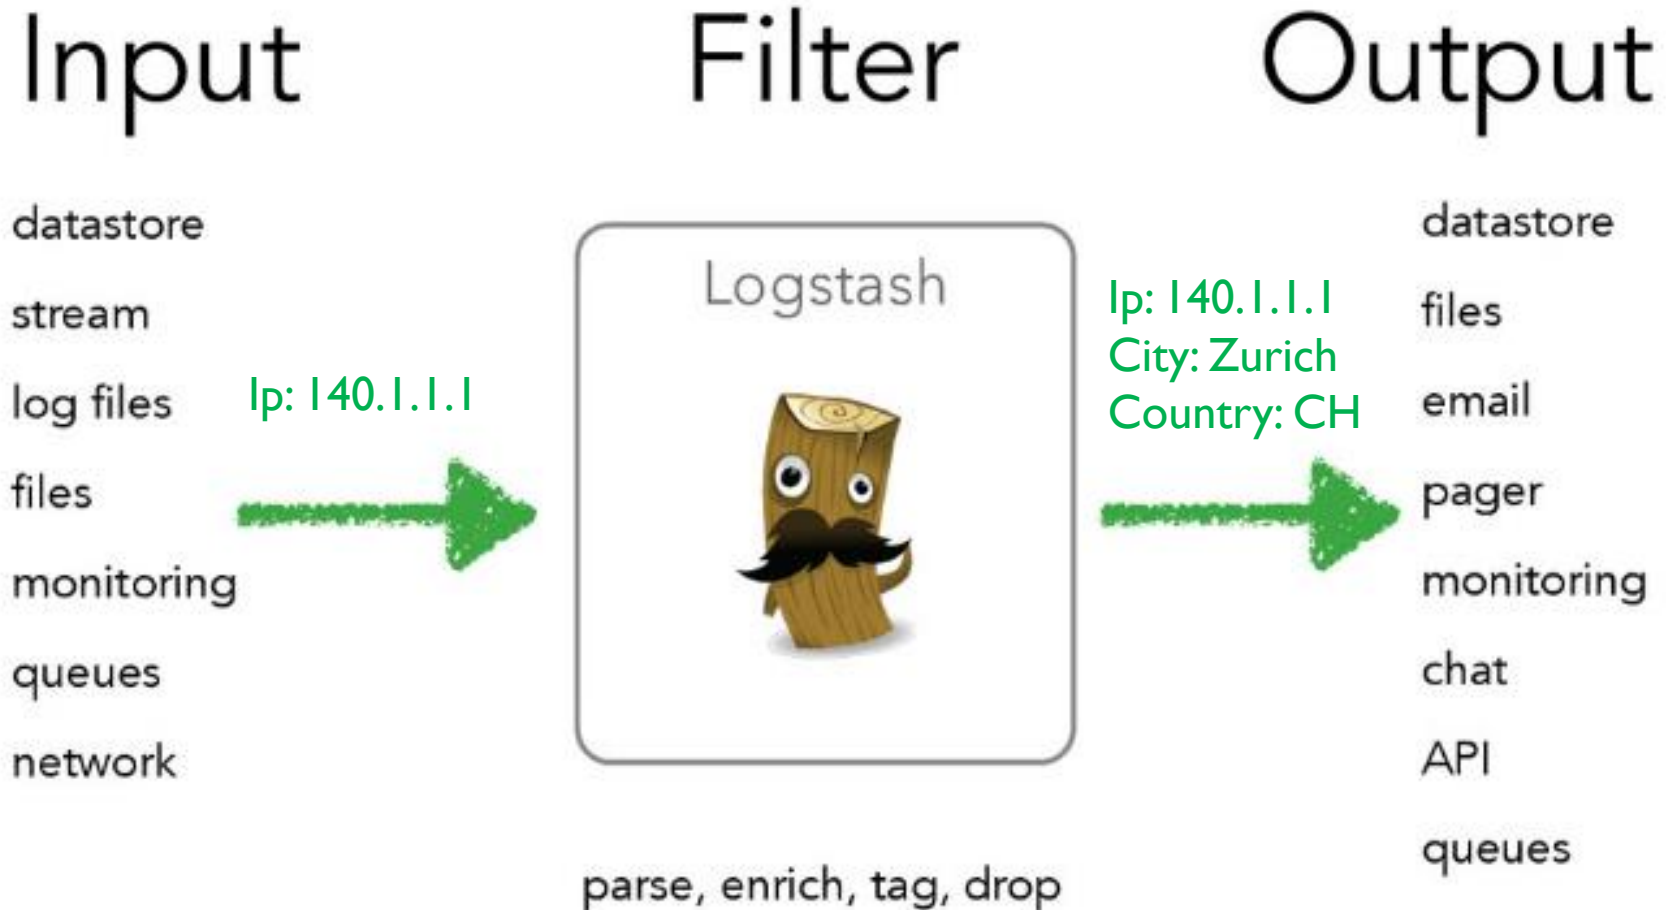
Elasticsearch-Logstash-Kibana

Logstash

- Managing events and logs
- Collect data
- Parse data
- Enrich data
- Store data
- Open Source: Apache License 2.0



Logstash architecture



How logstash works

```
input {  
  file {  
    path => "/tmp/access_log"  
    start_position => "beginning"  
  }  
}
```

```
filter {  
  grok {  
    match => { "message" => "%{COMBINEDAPACHELOG}" }  
  }  
}
```

```
output {  
  elasticsearch {  
  }  
}
```

Logstash Input plugins

- Stdin – Reads events from standard input
- File – Streams events from files (similar to “tail -0F”)
- Syslog – Reads syslog messages as events
- Eventlog – Pulls events from the Windows Event Log
- Imap – read mail from an IMAP server
- Rss – captures the output of command line tools as an event
- Snmptrap – creates events based on SNMP trap messages
- Twitter – Reads events from the Twitter Streaming API
- Irc – reads events from an IRC server
- Exec – Captures the output of a shell command as an event
- Elasticsearch – Reads query results from an Elasticsearch cluster
-

Logstash Filter plugins

- grok – parses unstructured event data into fields
- Mutate – performs mutations on fields
- Geoip – adds geographical information about an IP address
- Date – parse dates from fields to use as the Logstash timestamp for an event
- Cidr – checks IP addresses against a list of network blocks
- Drop – drops all events
- ...

Logstash Output plugins

- Stdout – prints events to the standard output
- Csv – write events to disk in a delimited format
- Email – sends email to a specified address when output is received
- Elasticsearch – stores logs in Elasticsearch
- Exec – runs a command for a matching event
- File – writes events to files on disk
- mongoDB – writes events to MongoDB
- Redmine – creates tickets using the Redmine API
-

Elasticsearch-Logstash-Kibana

ElasticSearch




Schema-flexible

- Built on top of [Apache Lucene™](#), a full-text search-engine library
 - A Schema-free, REST & JSON based distributed search engine with real-time analytics
 - Capable of scaling to hundreds of servers and petabytes of structured and unstructured data
 - Open Source: Apache License 2.0
- Real scalability comes from horizontal scale*
- **Wikipedia** uses Elasticsearch to provide full-text search with highlighted search snippets, and *search-as-you-type* and *did-you-mean* suggestions
 - **The Guardian** uses Elasticsearch to combine visitor logs with social-network data to provide real-time feedback to its editors about the public's response to new articles
 - **Stack Overflow** combines full-text search with geolocation queries and uses *more-like-this* to find related questions and answers
 - **GitHub** uses Elasticsearch to query 130 billion lines of code

Elasticsearch vs. Relational DB

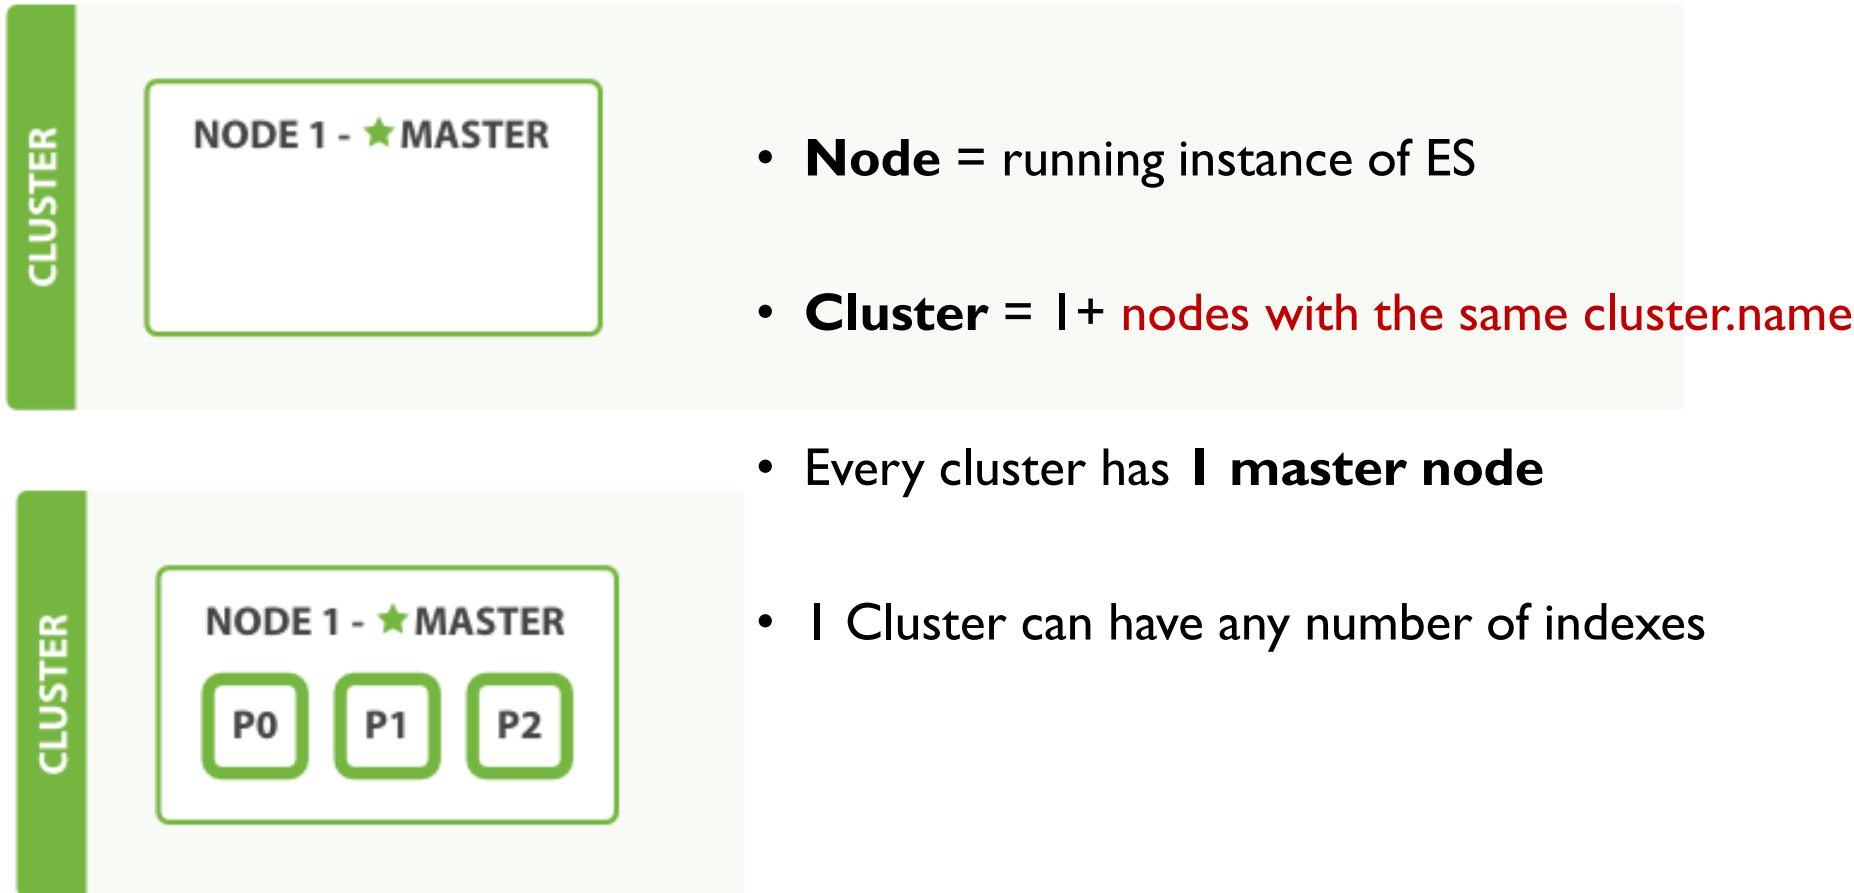
ElasticSearch	Relational DB
Index	Database
Type	Table
Document	Row
Field	Column
Shard	Partition
Mapping	Schema
- (everything is indexed)	Index
Query DSL (<i>domain specific language</i>)	SQL



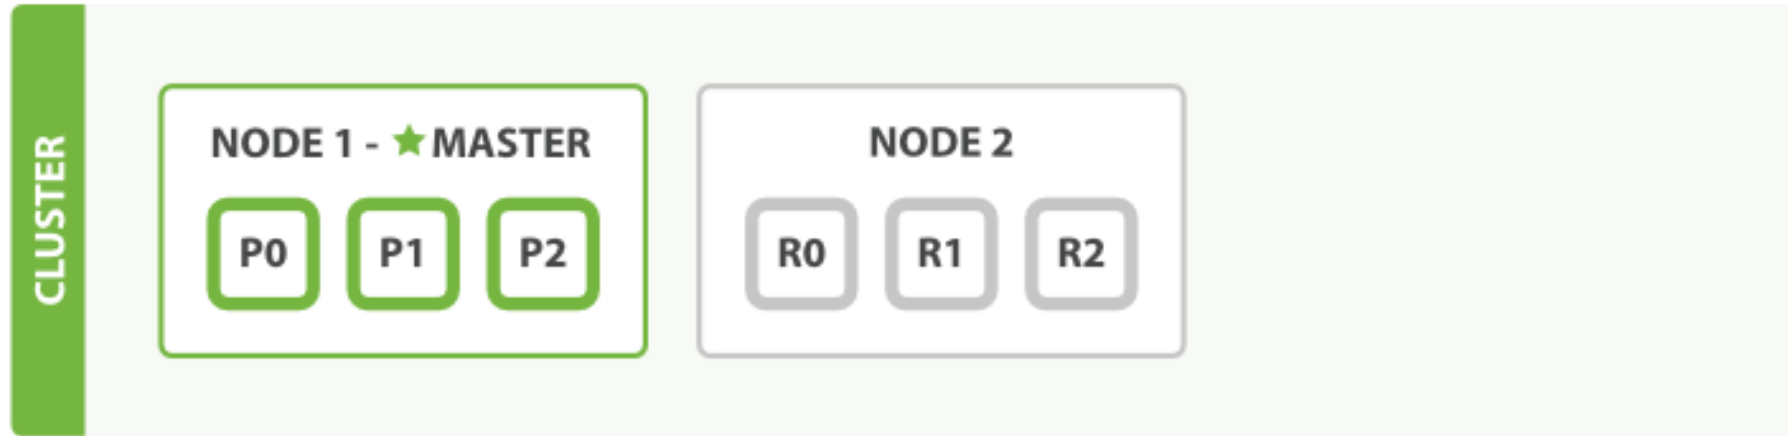
What is a shard

- a shard is a single instance of Lucene, and is a complete search engine in its own right
- Documents are stored and indexed in shards → shards are allocated to nodes in your cluster
- As your cluster grows or shrinks, Elasticsearch will automatically migrate shards between nodes so that the cluster remains balanced
- A shard can be either a *primary* shard or a *replica* shard
 - Each document in your index belongs to a single primary shard
 - A replica shard is just a copy of a primary shard

ElasticSearch clustering – single node cluster

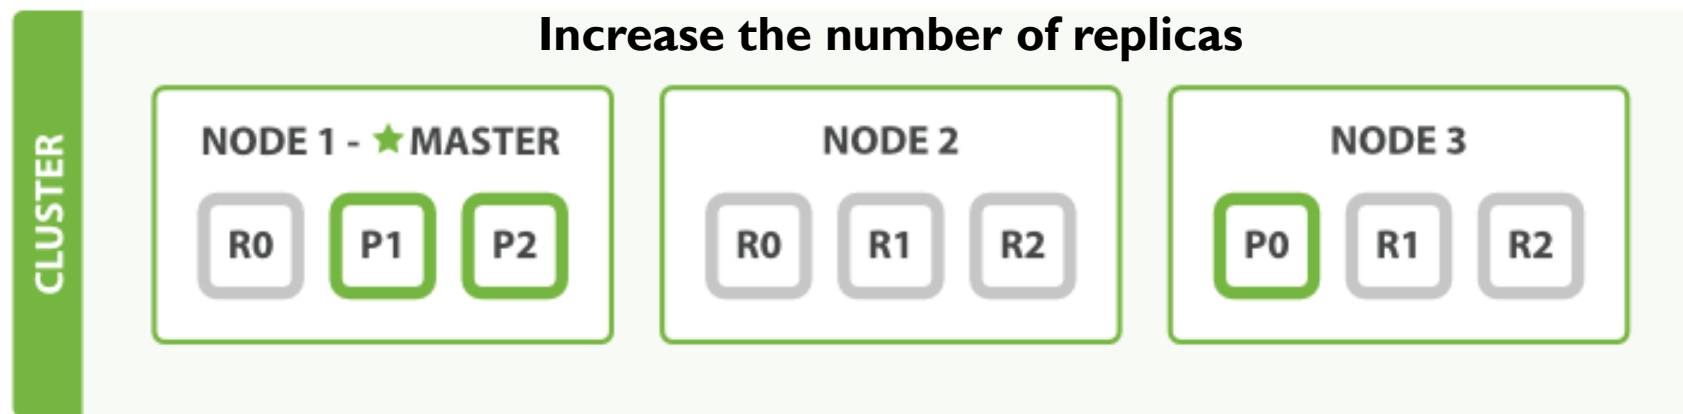


ElasticSearch clustering – adding a second node



- A cluster consists of one or more nodes with the same cluster.name
 - All primary and replica shards are allocated
 - Each index has one primary (P) and one replica (R) shard
 - Clients talk to any node in the cluster

ElasticSearch clustering – adding a third node



- More primary shards:
 - faster indexing
 - more scale
- More replicas:
 - faster searching
 - more failover

Talking to Elasticsearch

HTTP method or verb: GET, POST, PUT, HEAD, or DELETE

- RESTful API with JSON over HTTP
 - Over port 9200
 - Access via web client, or command line by `curl` command

```
curl -X<VERB> '<PROTOCOL>://<HOST>:<PORT>/<PATH>?<QUERY_STRING>' -d '<BODY>'
```

- JSON (JavaScript Object Notation) ← the standard format used by NoSQL

```
curl -XGET 'http://localhost:9200/_count?pretty' -d '{
  "query": {
    "match_all": {}
  }
}'
```

- Elasticsearch clients
 - Java API, Java REST client, JavaScript API, PHP API, Python API, Perl API...

Indexing a document

Index Type Id

```
curl -XPUT localhost:9200/test/product/1 -d '{"category": "electronics", "price": 129.99, "id": 1, "name": "ipod"}'
---
{"ok":true,"_index":"test","_type":"product","_id":"1","_version":1}
```

Document

```
curl -XPOST localhost:9200/test/product -d '{"category": "electronics", "price": 129.99, "name":"ipod"}'
---
{"ok":true,"_index":"test","_type":"product","_id":"9wrADN4eS8uXm3gNpDvEJw","_version":1}
```

- Store a document in an index so that it can be retrieved and queried
- Like the INSERT keyword in SQL

Retrieving documents

```
curl -XGET 'localhost:9200/test/product/1?pretty'
---
{
  "_index" : "test",
  "_type" : "product",
  "_id" : "1",
  "_version" : 2,
  "_exists" : true, "_source" : {"category": "electronics", "price":
129.99, "name": "ipod"}
}
```

- Using GET method to retrieve document
- We can retrieve a specific document if we happen to know its id

Performing Queries

- Using the `q=<query>` form performs a full-text search by parsing the query string value
- Query with **query DSL**, which is specified using a JSON request body

```
curl -XGET 'localhost:9200/test/product/_search?
q="ipod"&format=yaml'
---
took: 104
timed_out: false
_shards:
  total: 1
  successful: 1
  failed: 0
hits:
  total: 1
  max_score: 0.15342641
  hits:
    - _index: "test"
      _type: "product"
      _id: "1"
      _score: 0.15342641
      _source:
        category: "electronics"
        price: 129.99
        name: "ipod"
```

```
GET /megacorp/employee/_search
{
  "query" : {
    "match" : {
      "last_name" : "Smith"
    }
  }
}
```

Query DSL – Combining Filters

```
SELECT product
FROM   products
WHERE  (price = 20 OR productID = "XHDK-A-1293-#fJ3")
      AND (price != 30)
```



Bool Filter

```
{
  "bool" : {
    "must" : [],
    "should" : [],
    "must_not" : [],
    "filter": []
  }
}
```

```
GET /my_store/products/_search
{
  "query" : {
    "constant_score" : { ❶
      "filter" : {
        "bool" : {
          "should" : [
            { "term" : {"price" : 20}}, ❷
            { "term" : {"productID" : "XHDK-A-1293-#fJ3"}} ❸
          ],
          "must_not" : {
            "term" : {"price" : 30} ❹
          }
        }
      }
    }
  }
}
```


Query DSL – Nesting Boolean Queries

```
SELECT document
FROM   products
WHERE  productID      = "KDKE-B-9947-#kL5"
      OR ( productID = "JODL-X-1937-#pV7"
          AND price    = 30 )
```



```
GET /my_store/products/_search
{
  "query" : {
    "constant_score" : {
      "filter" : {
        "bool" : {
          "should" : [
            { "term" : {"productID" : "KDKE-B-9947-#kL5"}}, ❶
            { "bool" : { ❷
              "must" : [
                { "term" : {"productID" : "JODL-X-1937-#pV7"}}, ❸
                { "term" : {"price" : 30}} ❹
              ]
            }
          ]
        }
      }
    }
  }
}
```

ELK use cases

Use Cases




Social



Location



User-Activity



Machine
(Log files)



Documents

Search



WIKIPEDIA
The Free Encyclopedia



rightmove



Logging



BNP PARIBAS

TOMTOM

NETFLIX

Security
Analytics



mozilla



Analytics



theguardian

Eventbrite

User cases

“ Elasticsearch, Logstash, and Kibana allow for real-time

verizon

Use Case	Logging, Analytics
Products	Elasticsearch, Logstash




“ With the ELK stack, we log more than 30K messages and 100K documents four times every day from the Mars Rover to optimize our space missions.”

Dan Isla, Data Scientist

Use Case	Search, Logging, Analytics
Products	Elasticsearch, Logstash, Kibana

Use Case	Search, L
Products	Elasticse

Cisco Talos Security Intelligence and Research Group: Hunting for Hackers




Cisco Talos use ELK to analyze...

- Sandbox data cluster
 - Dynamic malware analysis reports
 - Search for related pattern, malewares
 - ES stats
 - 10 nodes
 - 3 TB
 - 100k reports/day
 - ~8 months of data
- Honeypot cluster
 - Collect attackers' attempt
 - { Account, password } pair
 - Executed commands
 - url of download files
 - Suspicious command center for report back

```
"_index": "logstash-telnet-sqs-2016.02.10",
"_type": "telnet-sqs",
"_source": {
  "Event.Type": "ConnectionLost",
  "@timestamp": "2016-02-10T23:51:10.000Z",
  "Event.Session": "1272f0ccd05111e5bb400242ac110001",
  "Net.IP.Src": "117.158.195.59",
  "User.Name": "admin",
  "User.Pass": "1234",
  "Net.Port.Src": 23,
  "@version": "1",
  "type": "telnet-sqs",
```

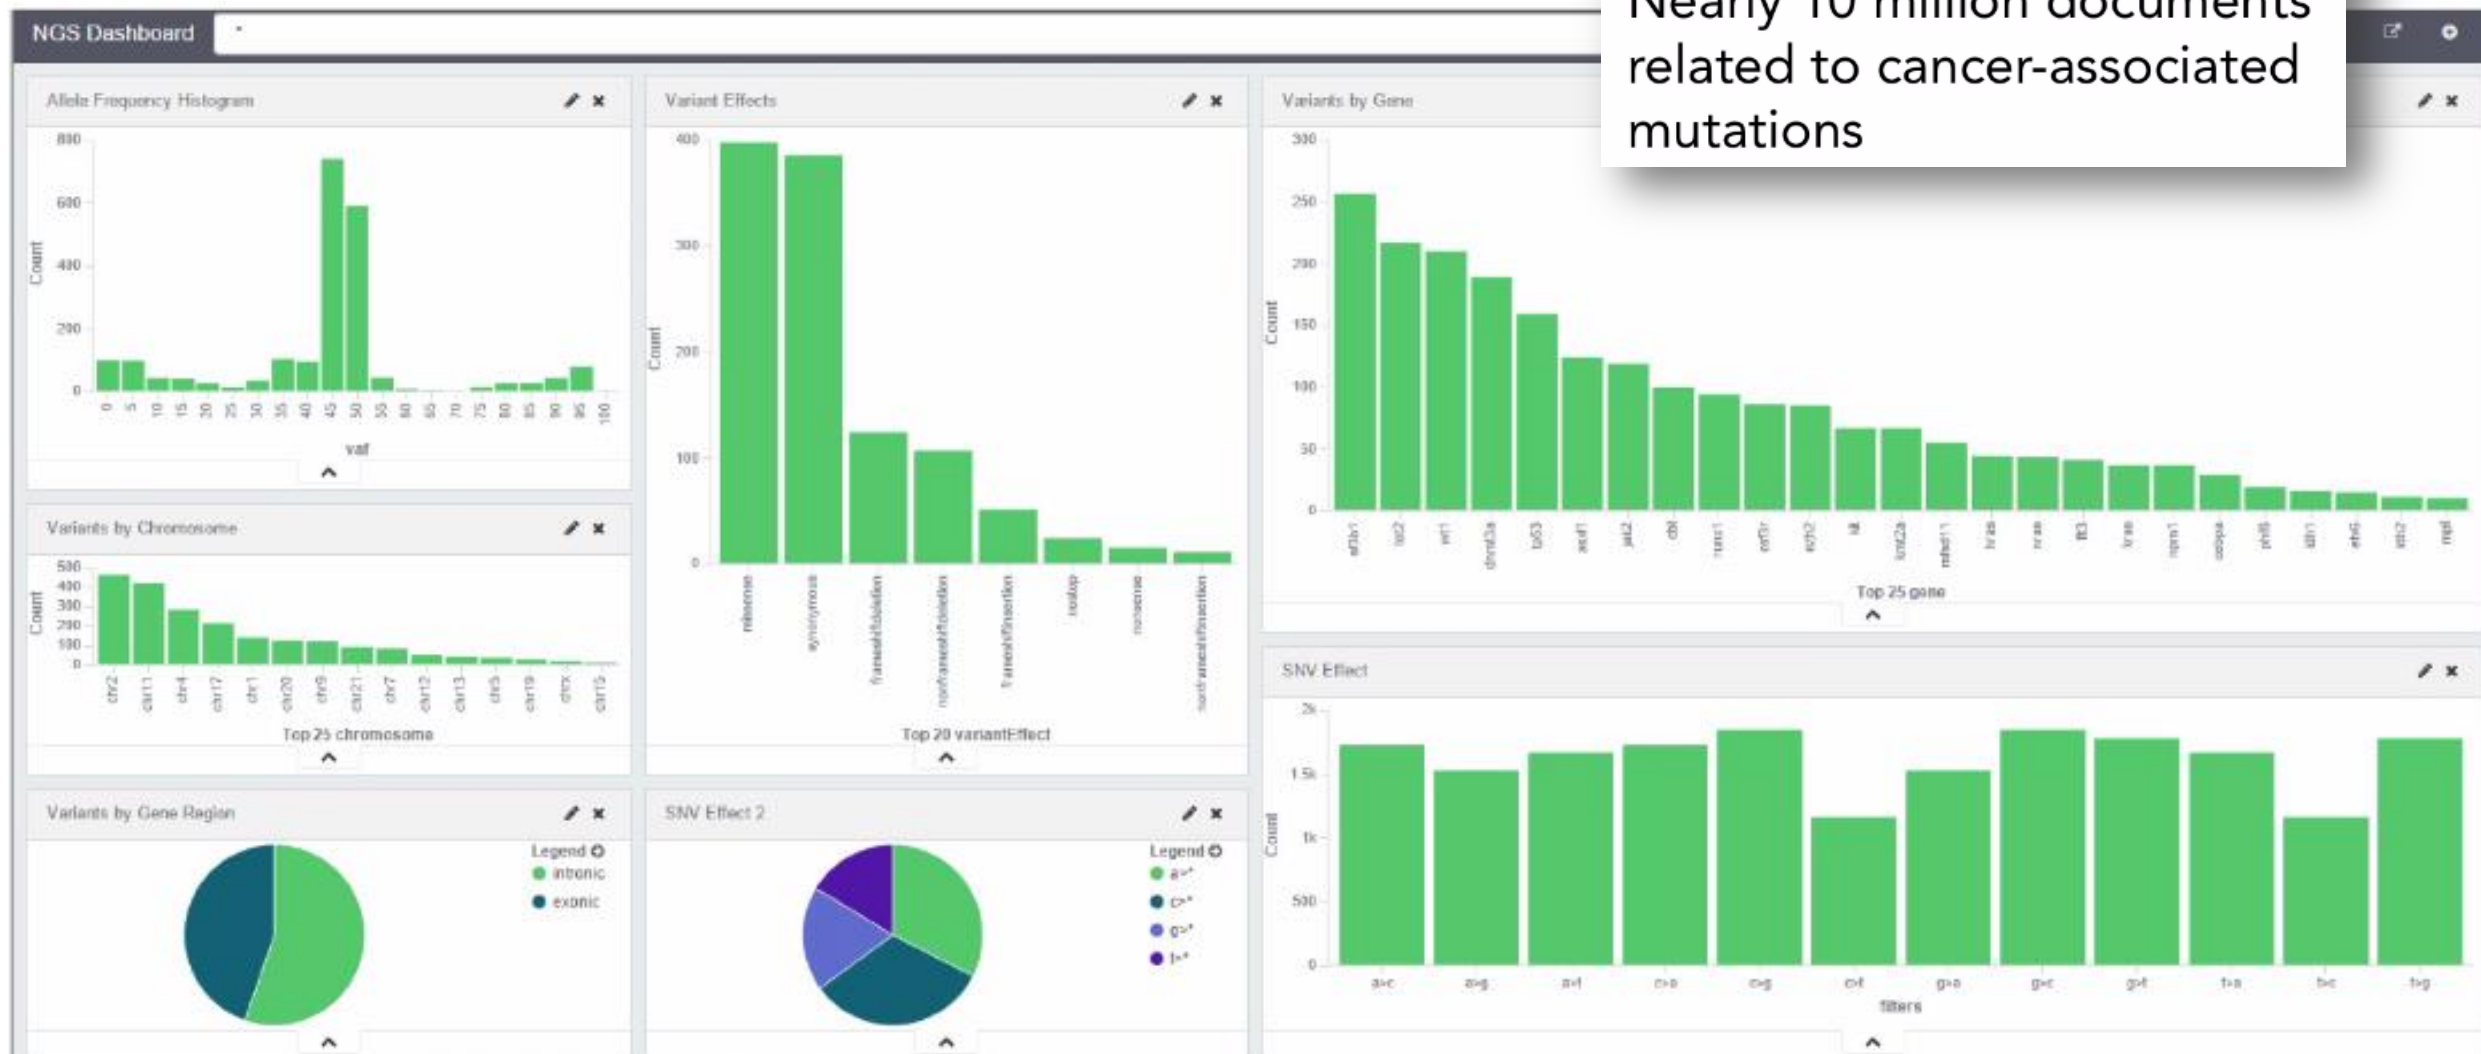

Yale's {elastic}SEARCH – The Search for Cancer's Causes and Cures

Sequencing and Interpretation Pipeline

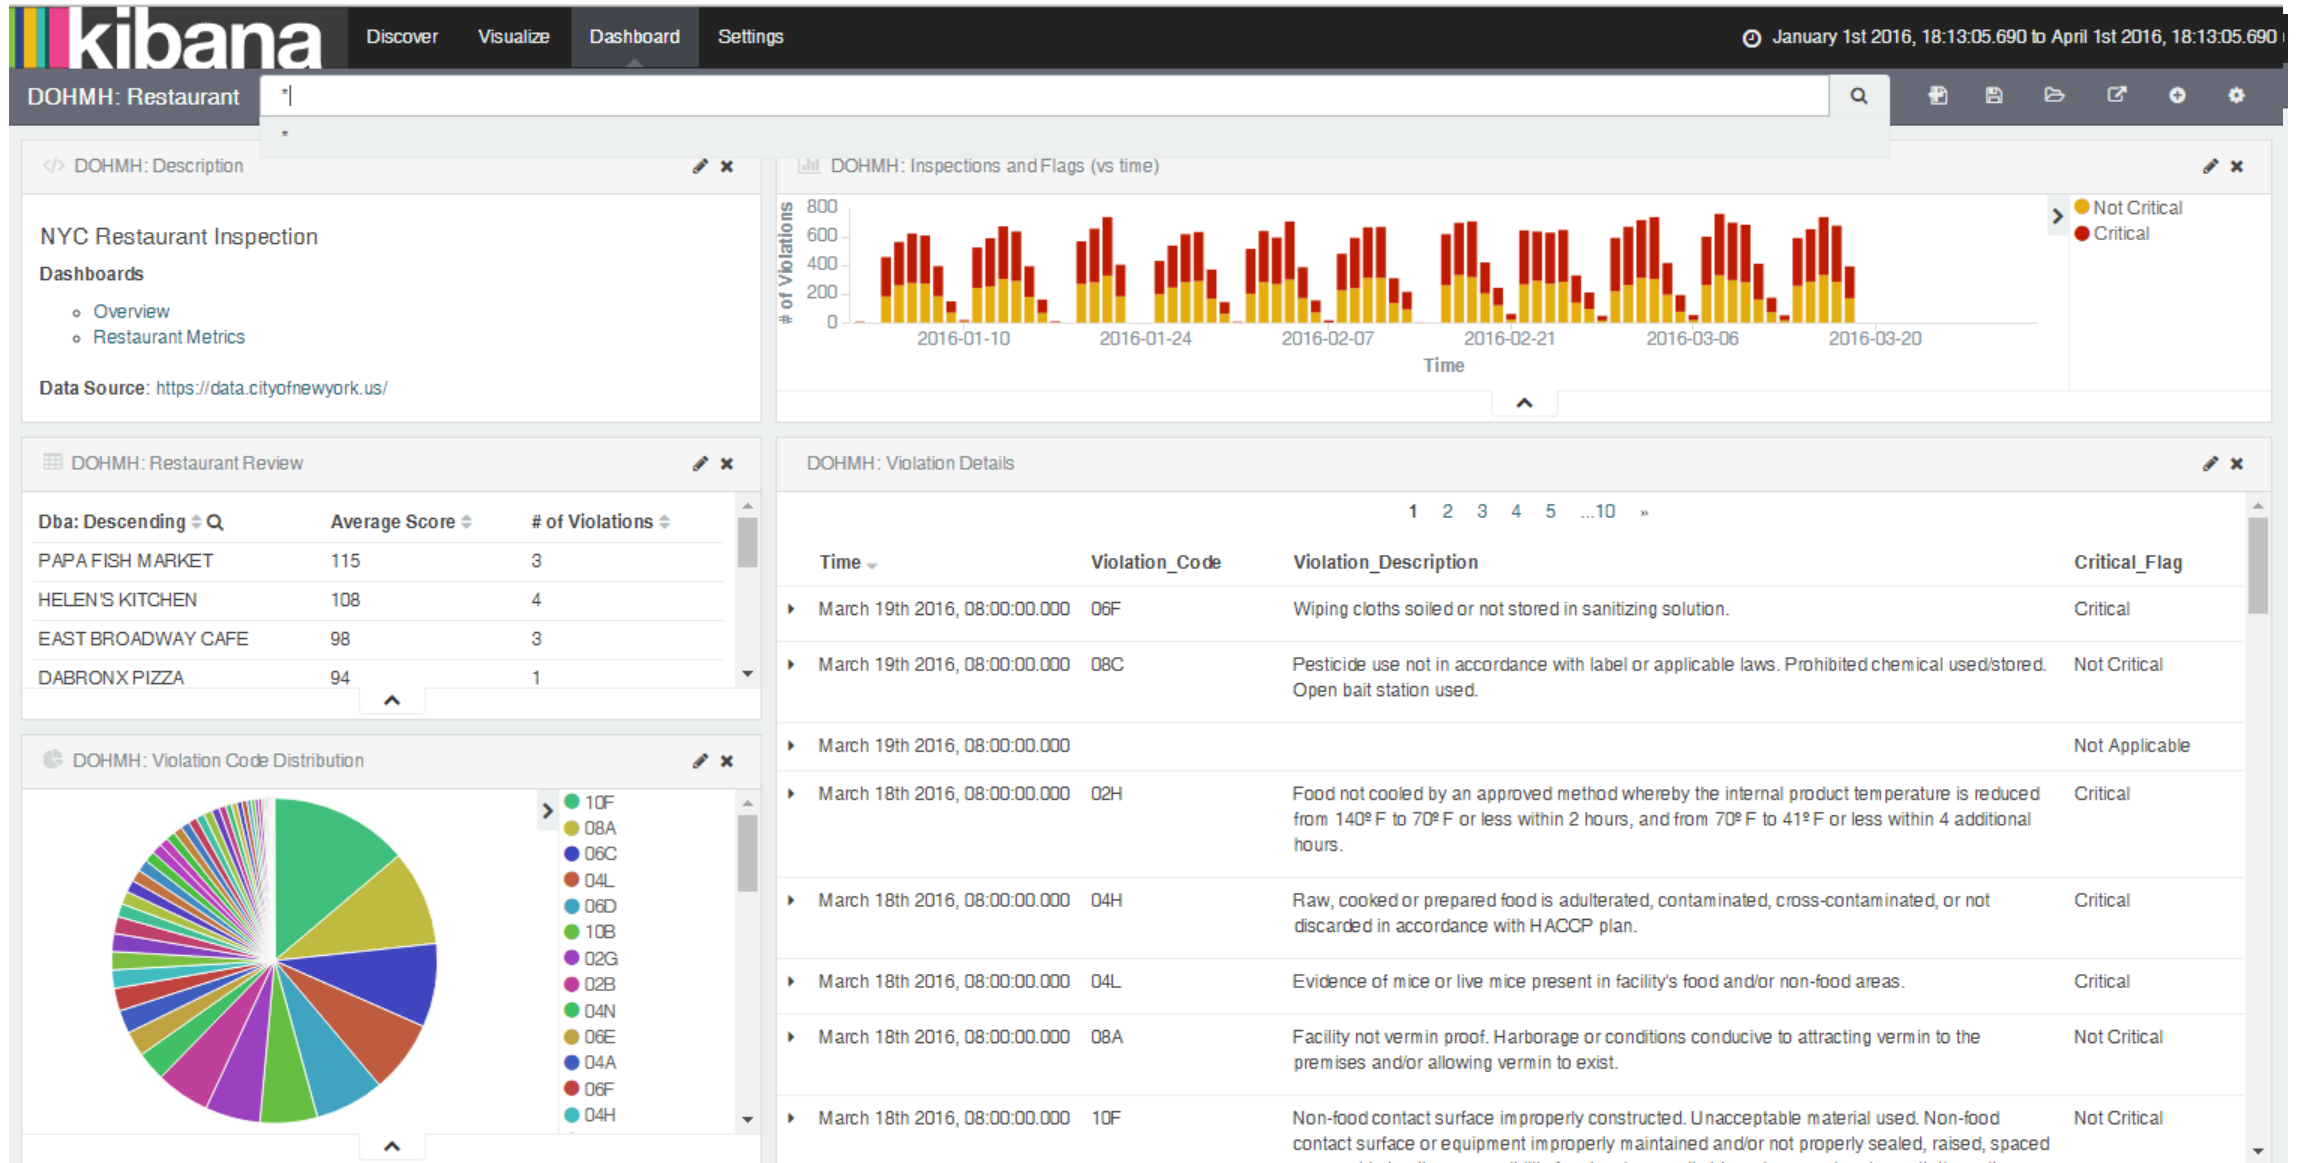


- With Next generation sequencing technology, the lab can process 8 million patients specimens yearly
- How to interpret this amount of data → what software can be used

Over 60 million variant annotations
Nearly 10 million documents related to cancer-associated mutations



NYC restaurants inspection @ELK



Implementation of ELK on Campus network & account anomaly detection

Network anomaly detection – case 1

- 從SMTP異常登入偵測疑似被盜用的帳號
 - 系統管理者通常這麼做.....

撰寫程式：從日誌檔統計帳號登入來源國家 → 產生統計檔 → 管理者判斷
→ grep, awk, sort, uniq, shell script, perl.....

```
2016-07-07T17:54:17.853Z,EXCHCAS03\Client EXCHCAS03,08D3757739790262,28,172.16.6.170:587,93.153.205.115:22343,*,SMTPSubmit,SMTPAccept
2016-07-07T17:54:17.853Z,EXCHCAS03\Client EXCHCAS03,08D3757739790262,29,172.16.6.170:587,93.153.205.115:22343,*,NTUCC\b99901017,auth
2016-07-07T17:54:17.853Z,EXCHCAS03\Client EXCHCAS03,08D3757739790262,30,172.16.6.170:587,93.153.205.115:22343,>,235 2.7.0 Authentication
2016-07-07T17:54:17.853Z,EXCHCAS03\Client EXCHCAS03,08D3757739790262,31,172.16.6.170:587,93.153.205.115:22343,*,b99901017,auth
```



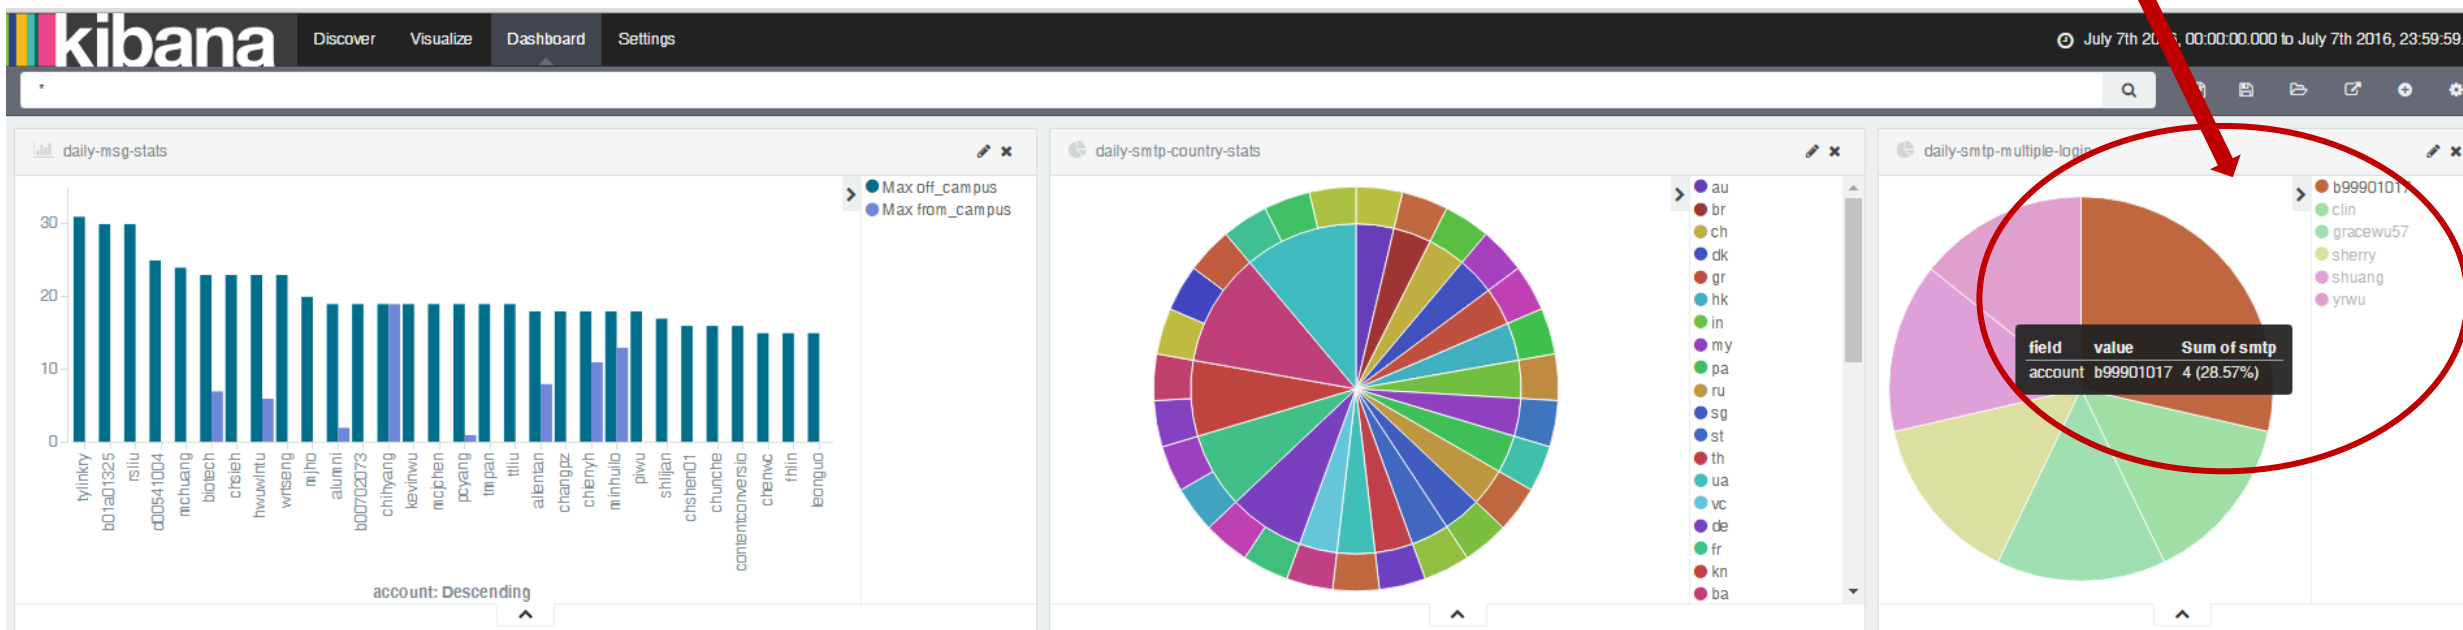
b99901048	-	140.112.108.229	=>	(KHARKIV,KHARKIVS'KA OBLAST',UA)
b99901017	-	159.224.53.134	=>	(Edeia,Goias,BZ)
b99901017	-	201.148.127.74	=>	(MOUNTAIN HOME,ARKANSAS,US)
b99901017	-	216.134.234.250	=>	(SAINT PETERSBURG,SAINT PETERSBURG CITY,RU)
b99901052	-	111.82.122.127	=>	(TAIPEI TAIWAN,TW)

Network anomaly detection – case I

- 透過ELK....

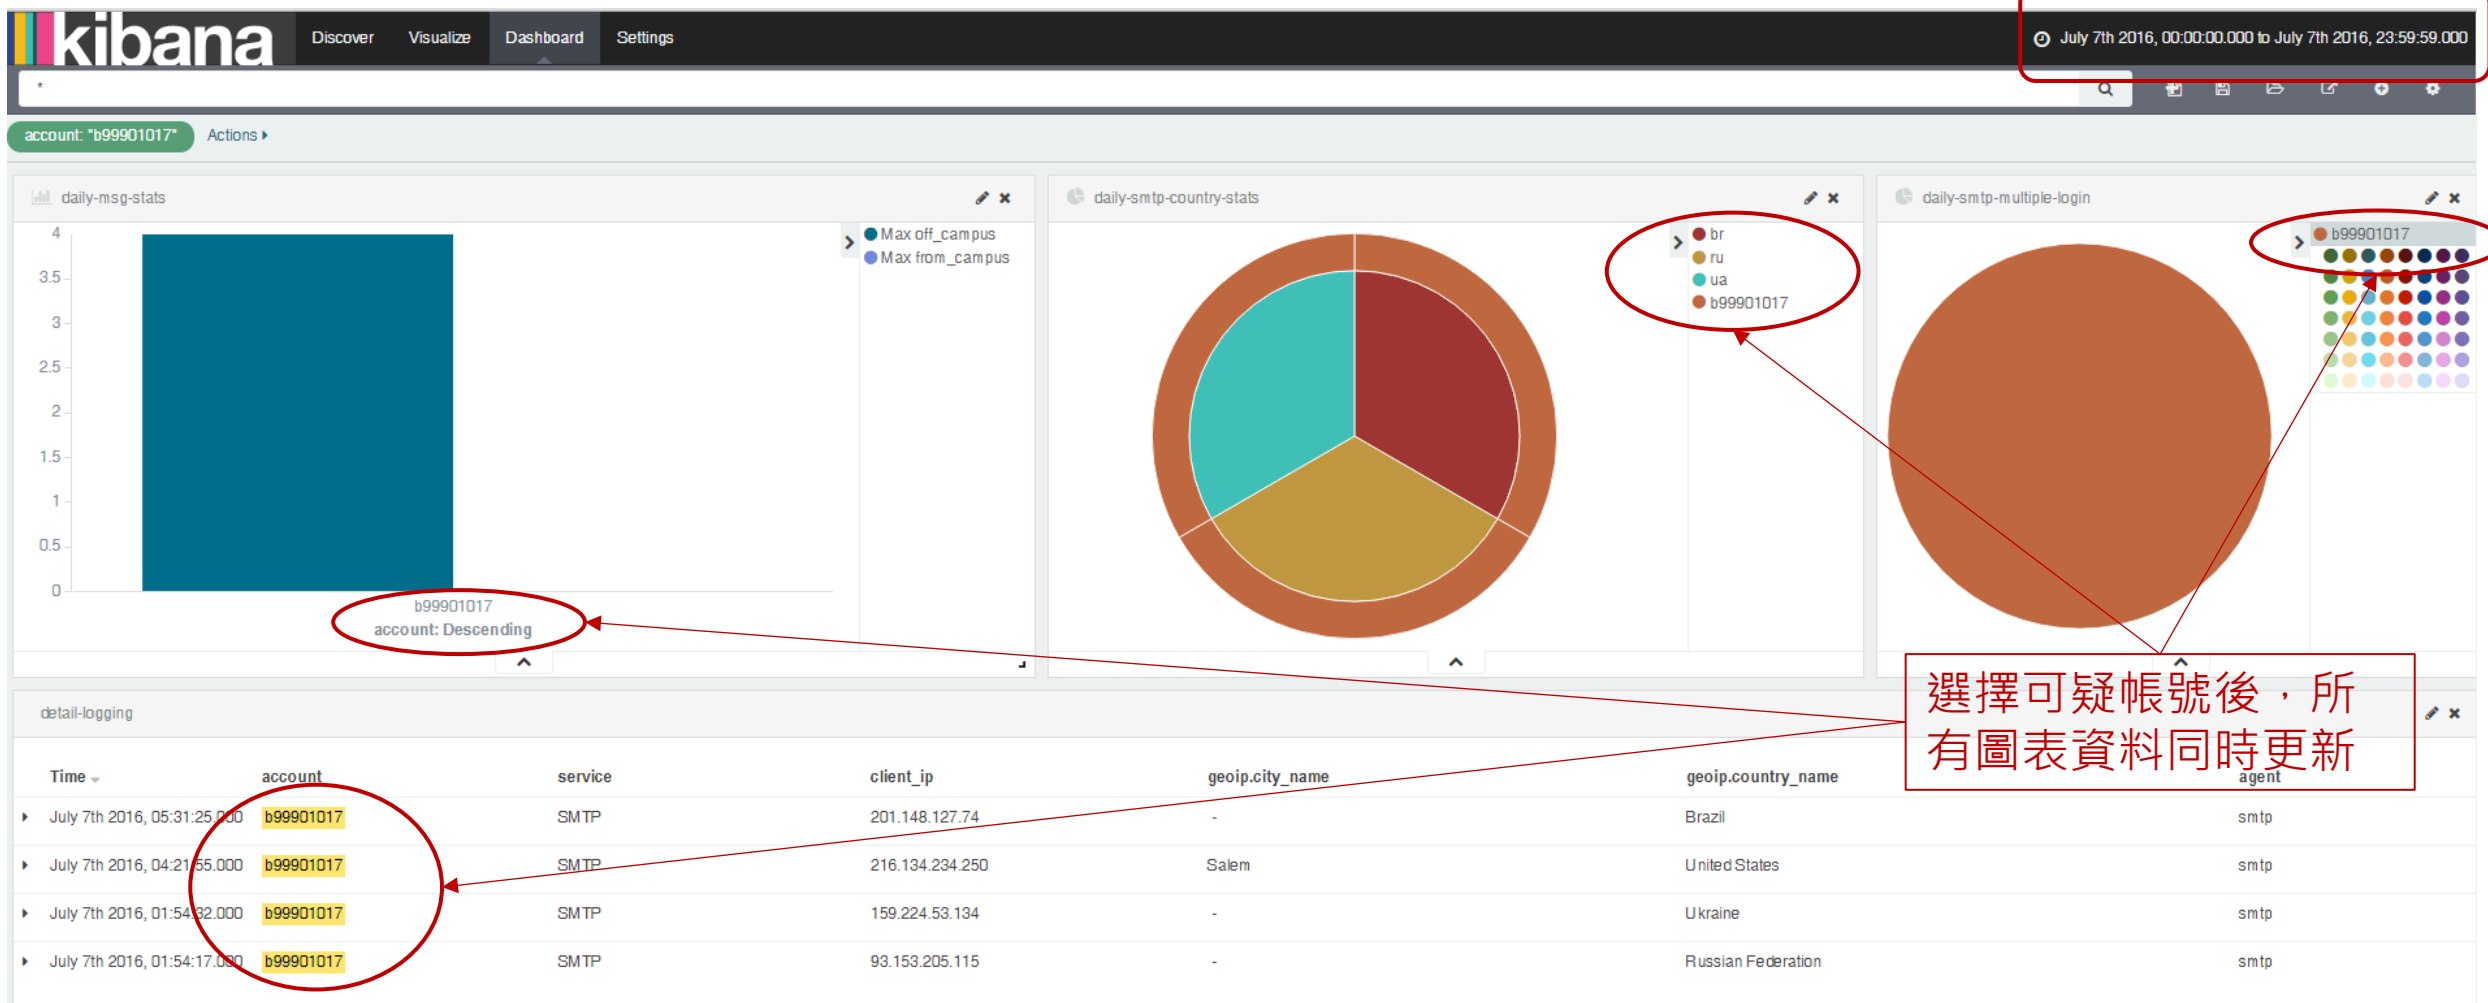
將log匯入ELK之後，透過Kibana介面呈現統計數據，容易辨識&確認異常情況

從4個國家登入SMTP



Network anomaly detection – case 1

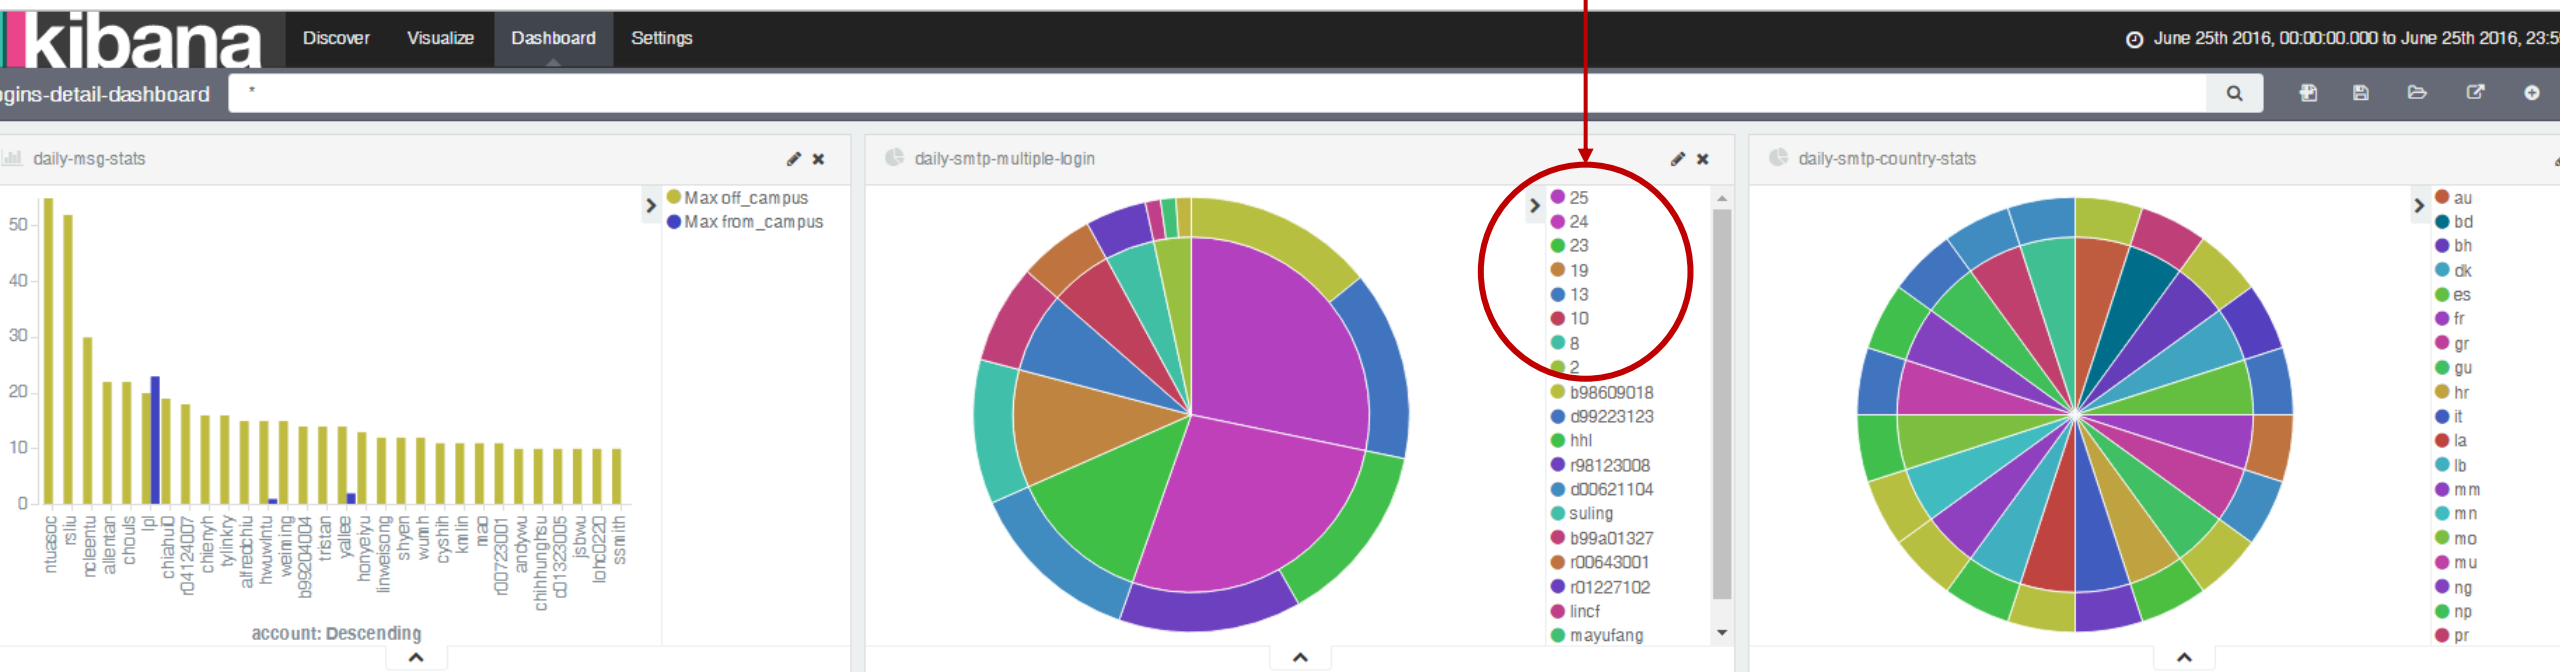
調整資料區間可比對該帳號的使用習慣



選擇可疑帳號後，所有圖表資料同時更新

Network anomaly detection – case 2


單日帳號透過SMTP登入之國家數排名



detail-logging

1 2 3 4 5 ...10 »

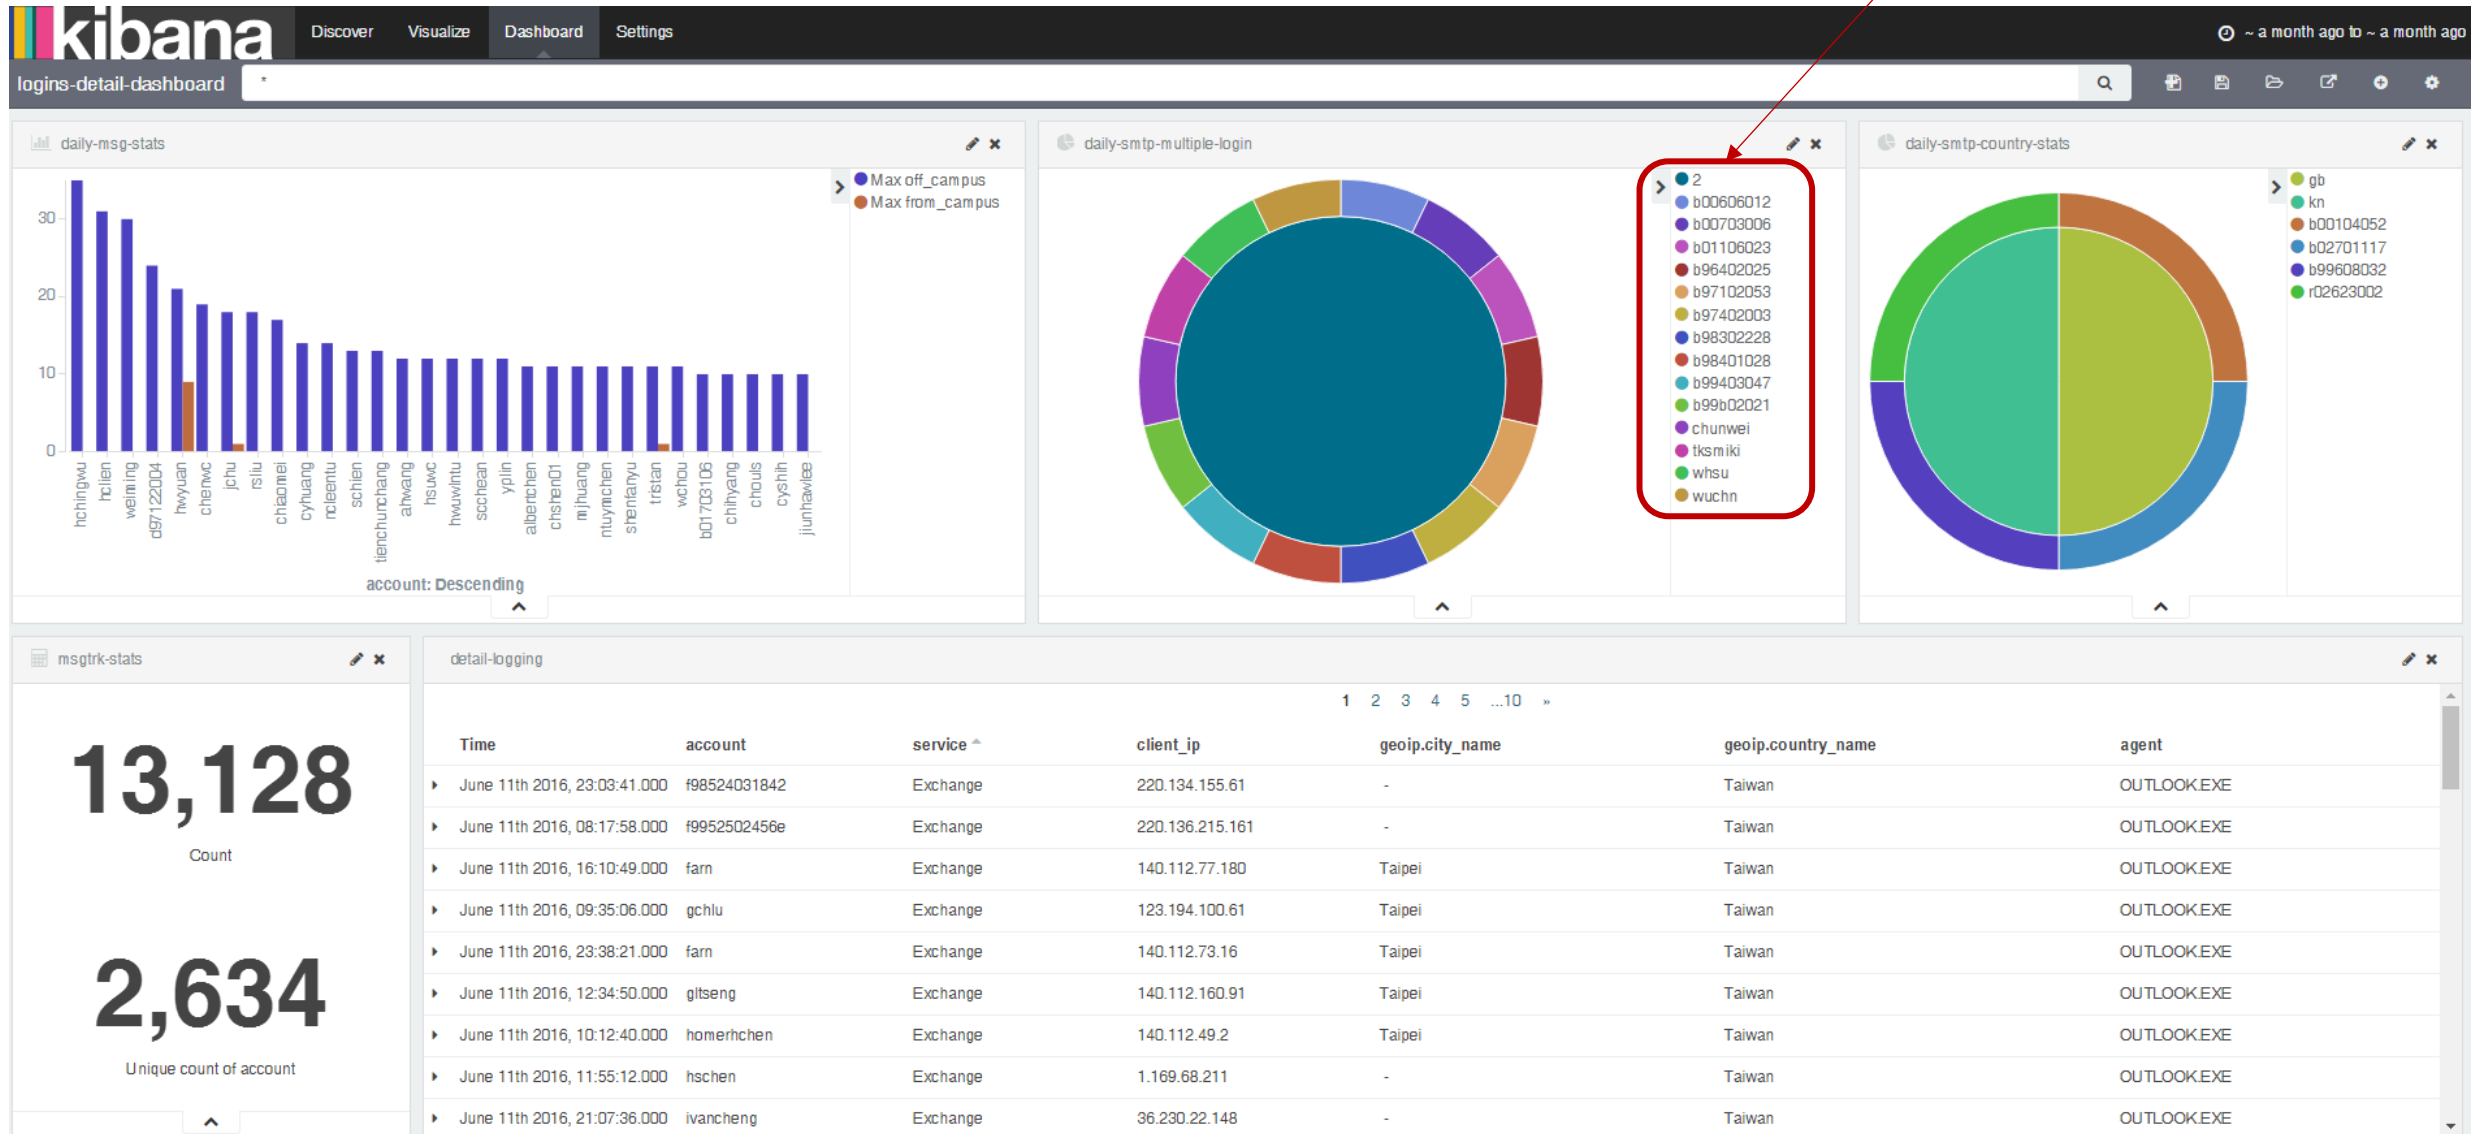
Time	account	service	client_ip	geoip.city_name	geoip.country_name	agent
June 25th 2016, 23:59:59.000	r00325001	POP3	209.85.213.25	Mountain View	United States	pop3
June 25th 2016, 23:59:58.000	r01922030	POP3	209.85.218.30	Mountain View	United States	pop3
June 25th 2016, 23:59:57.000	b97801056	POP3	209.85.216.139	Mountain View	United States	pop3
June 25th 2016, 23:59:56.000	jauchen	Exchange	42.72.10.174	Taipei	Taiwan	OUTLOOK.EXE

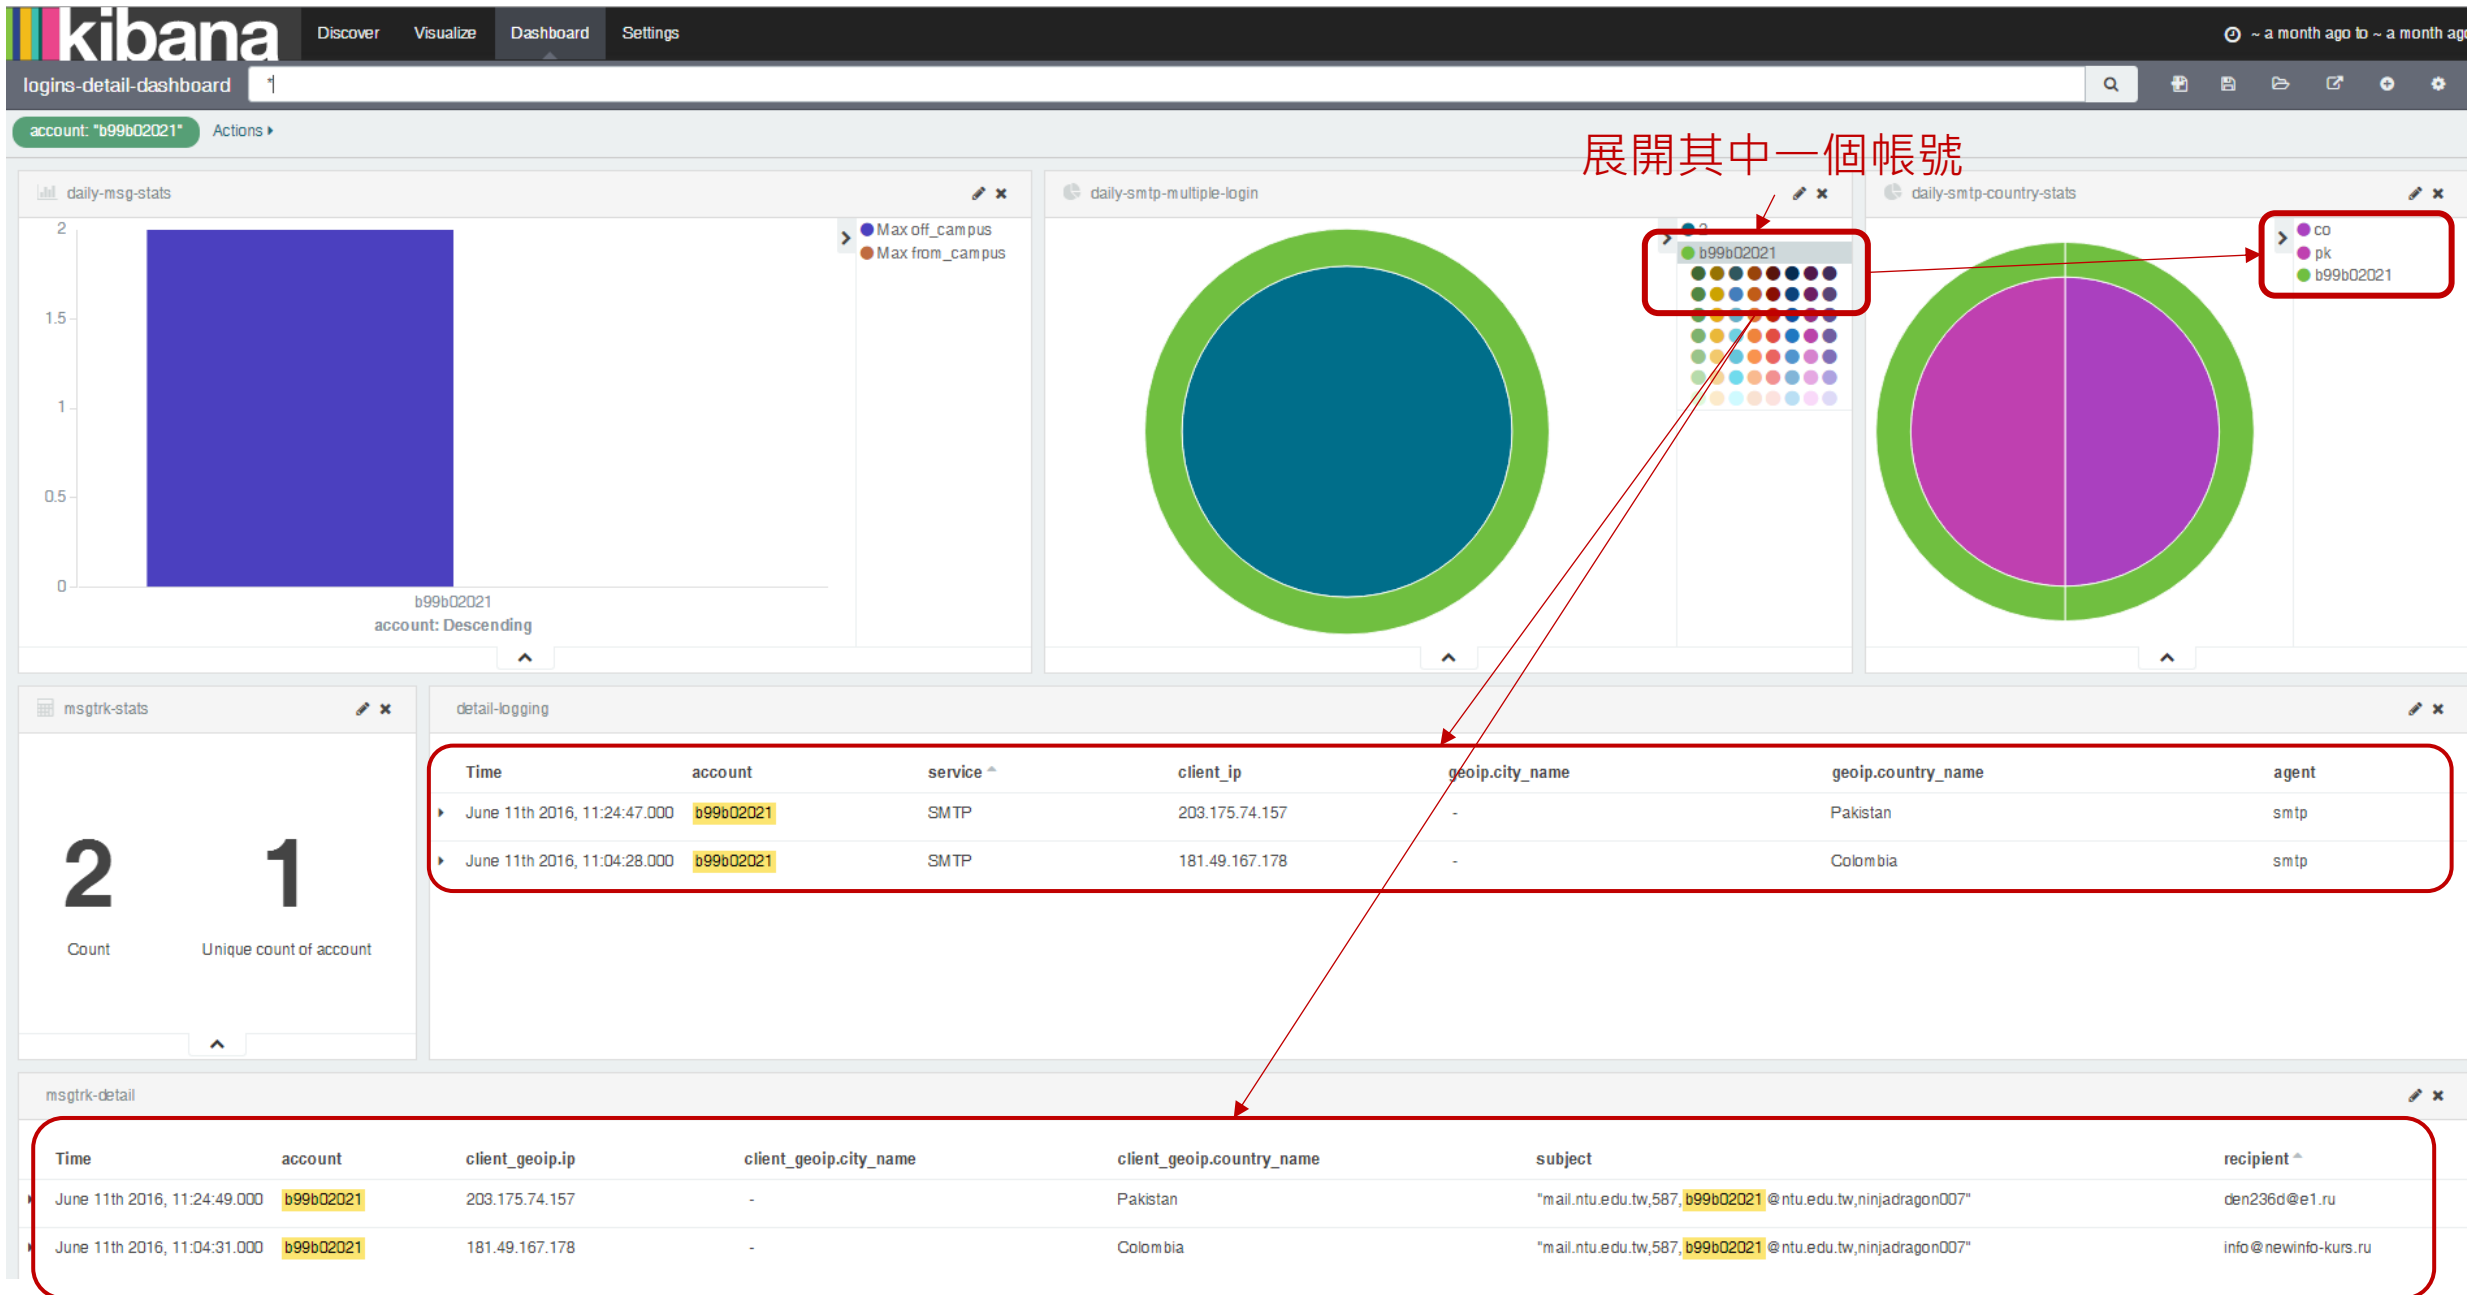


Time	account	service	client_ip	geoip.city_name	geoip.country_name	agent
June 25th 2016, 19:29:08.000	b98609018	SMTP	117.248.21.109	Bangalore	India	smtp
June 25th 2016, 19:26:52.000	b98609018	SMTP	192.117.182.186	-	Israel	smtp
June 25th 2016, 19:26:13.000	b98609018	SMTP	113.20.116.18	Hanoi	Vietnam	smtp
June 25th 2016, 19:25:53.000	b98609018	SMTP	195.174.56.77	Çorlu	Turkey	smtp
June 25th 2016, 19:25:05.000	b98609018	SMTP	114.79.36.246	-	Indonesia	smtp
June 25th 2016, 19:24:50.000	b98609018	SMTP	111.119.175.158	-	Pakistan	smtp
June 25th 2016, 19:24:49.000	b98609018	SMTP	126.78.184.107	Kazo	Japan	smtp
June 25th 2016, 19:23:48.000	b98609018	SMTP	111.91.88.13	Thane	India	smtp
June 25th 2016, 19:23:47.000	b98609018	SMTP	190.93.52.170	San Luis	Argentina	smtp
June 25th 2016, 19:23:01.000	b98609018	SMTP	123.237.224.187	Rajkot	India	smtp

Network anomaly detection – case 3

單日內從多國登入SMTP





利用ELK協助偵測異常帳號行為

Month	偵測到被盜帳號數量
105.01	17
105.02	11
105.03	14
105.04	6
105.05	10
105.06	44
105.07	71


- 105.05開始嘗試利用ELK
 - 匯入及分析帳號登入記錄
 - Mail protocols
 - VPN
 - Outbound Message logs
 - 研究log性質以擷取更有意義的特性
 - 反覆修正改進Kibana呈現統計數據的方式 → 以快速發現異常情況

Network anomaly detection


- netflow

某連線學校反應網路連線異常，
觀察區網骨幹流量圖可明顯發現
兩個異常凸起點

北區區網總流量分析 北區區網周流量分析



每日 圖表 (5 分鐘 平均)




該校與區網之頻寬為1 Gb，而在兩
個攻擊時間點發生時，1Gb頻寬幾
乎完全被攻擊頻寬佔滿，導致該校
對外連線近乎中斷


Network anomaly detection - netflow

* 6:TCP 17:UDP 1:ICMP

- Top Protocol Packets在運作正常之網路環境中，協定分佈比例最高應為Protocol: 6 TCP
- 該校之協定分佈為Protocol: 17 UDP，封包佔據了超過 98% 之網路流量
- 由此顯示該校正遭受**UDP Flooding 攻擊**




正常協定分佈比例




異常協定分佈比例

Network anomaly detection - netflow

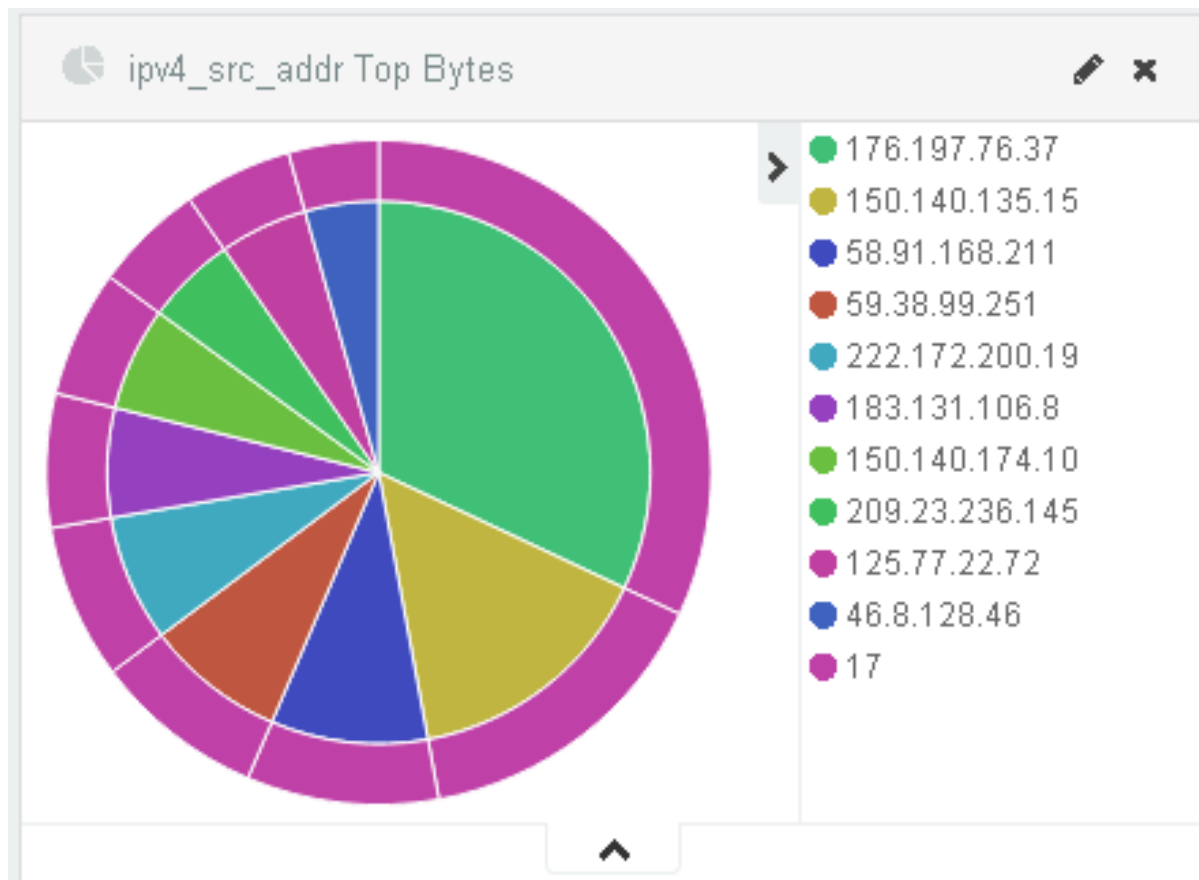


- 觀察 I4_src_port Top Bytes顯示Source Port 123 佔了絕大比例
- Port 123為網路校時NTP Server所使用之Port
- 由此判斷**此攻擊來源為NTP攻擊**



- 接著觀察 ipv4_dst_addr Top Bytes顯示目的IP 120.97.6.56佔據該校最大流量
- 由此可知**此次受攻擊目的IP位址**

Network anomaly detection - netflow



- 分析攻擊來源：
 - 由ipv4_src_addr Top Packets統計圖可觀察多個來源IP
 - 由此可知**此攻擊型態為 Distributed Denial of Service (DDoS)**

What we can do next...

- Anomaly network traffic detection
 - Detect from analyzing across netflow、 FW, and server logs
- Suspicious account usage
 - Account → IP address → Netflow
- Suspicious email discovery
 - Establish “sender address”-“sender mail server” relationship from daily mail log → discover suspicious sender-forgery email spoofing
- And there're many more ...