

# INFORMATION SECURITY






# INFORMATION SECURITY





# INFORMATION SECURITY





**it depends...**

**Organisations see the game as chess**, laying out their own assets in a strategic manner trying to outthink the adversary.

**Attackers are playing poker**, trying to bluff and gamble their way in.



Play both  
games!





# CSIRTs, CIRTs, CERTs, SIRTs, OR IRTs



**A Computer Security Incident Response Team (with any of the above acronyms) is a concrete organisational entity that is assigned the responsibility for coordinating and supporting the response to a computer security event or incident.**

The goal of a CSIRT is to minimise and control the damage resulting from incidents, provide effective guidance for response and recovery activities, and work to prevent future incidents from happening.

# KEY REQUIREMENTS



Along with an **effective incident response handbook** (playbook), and a **mandate to protect the organisation**, great teams will also possess:

- **Adequate resources, tooling, and training** availabilities for the team to remain relevant and effective.
- Proper documentation and **understanding of what must be protected**.
- Documented and **reliable relationships** with other groups in the organisation.



# OxCERT



## University of Oxford Computer Emergency Response Team

1994 Founded as a group of  
volunteers across Oxford

1998 FIRST membership

2001 Full time staff as part of  
central Networks Team

2013 Information Security Team

2017 Five full time staff





# THE FOUR CORE QUESTIONS



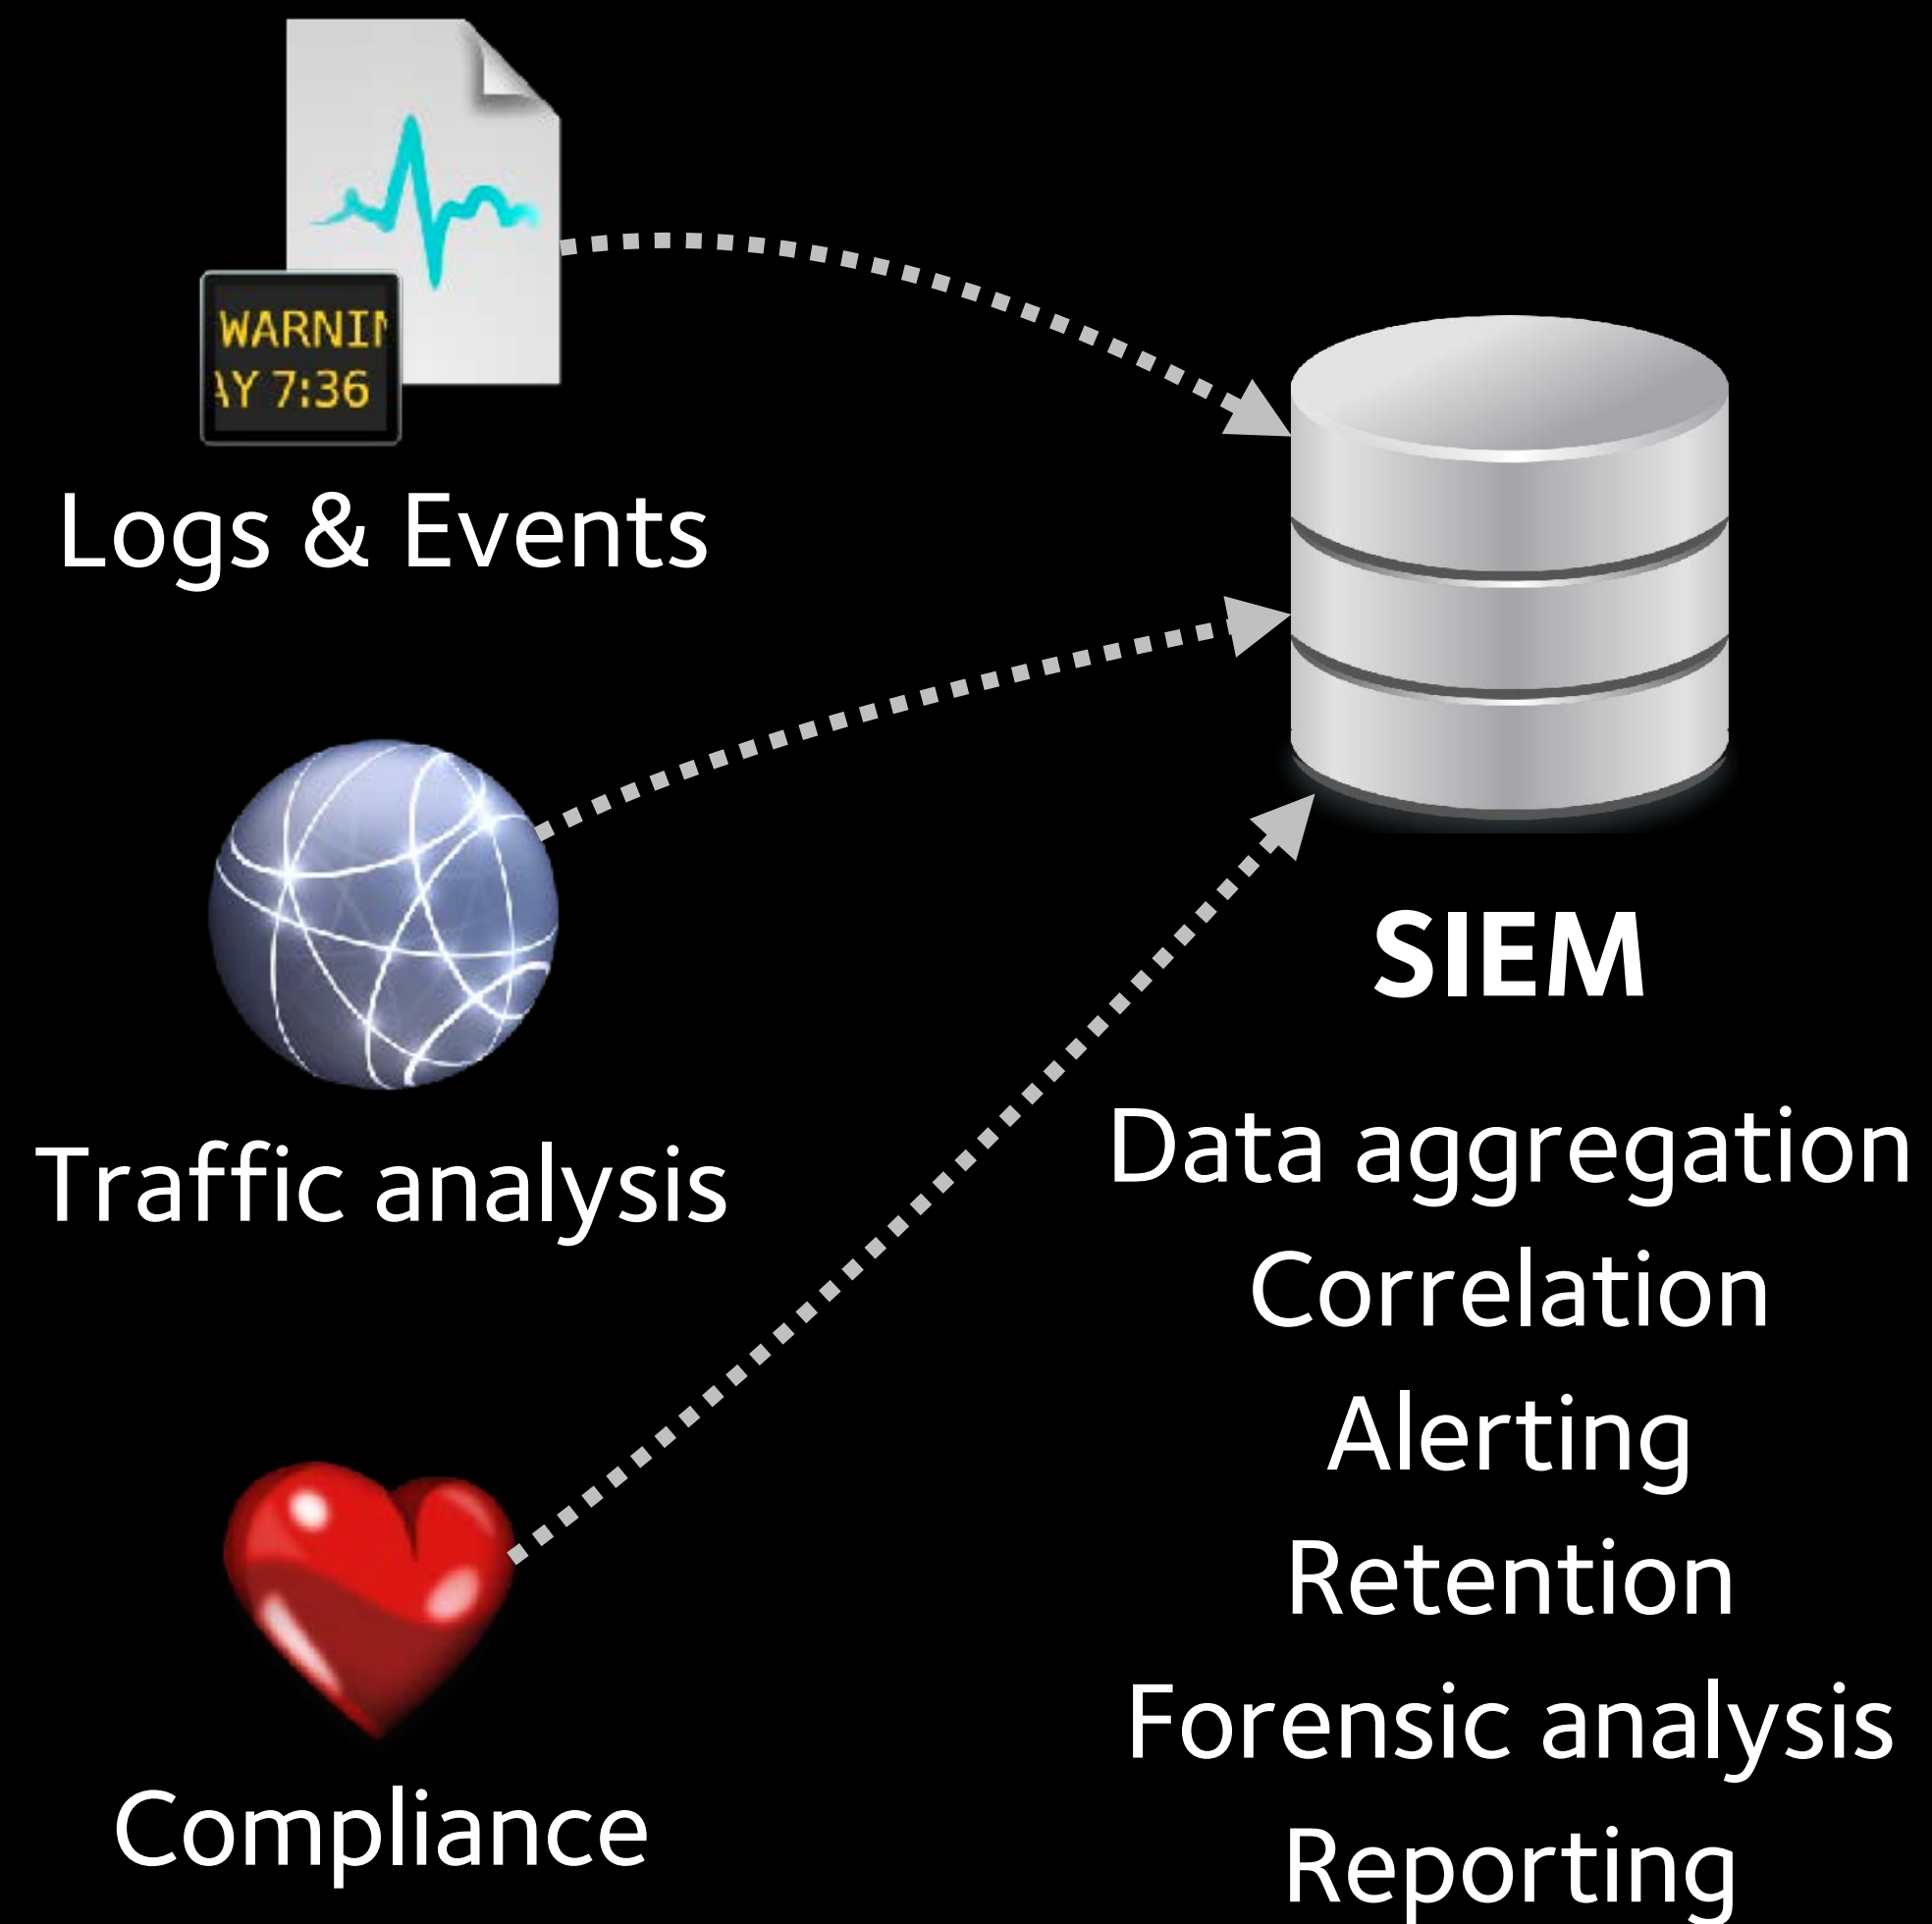
The core foundation to security monitoring and incident response:

1. What are we trying to protect?
2. What are the threats?
3. How do we detect them?
4. How do we respond?

Understanding that there will always be a place for incident prevention, while also recognizing that not every threat can be blocked, ensures a pragmatic approach to detection and response.



# THE TOOL MAKETH THE TEAM





# LOGS & EVENT INGEST



Collection of lightweight data shippers for Elasticsearch  
Filebeat, Winlogbeat, Heartbeat, Metricbeat, Packetbeat



Data processing pipeline with a variety of input formats  
e.g. syslog, log4j, puppet factor, rss, https, rss, snmptrap

- Variety of filter plug-ins
- Powerful post-processing
- Data aggregation and pruning



# NETFLOW – TRAFFIC METADATA



NetFlow also known as Jflow, Netstream, Cflowd, sflow, and IPFIX vary slightly, but the essence is the same: each technology creates a record of connections (a flow) between at least two hosts, including metadata like:

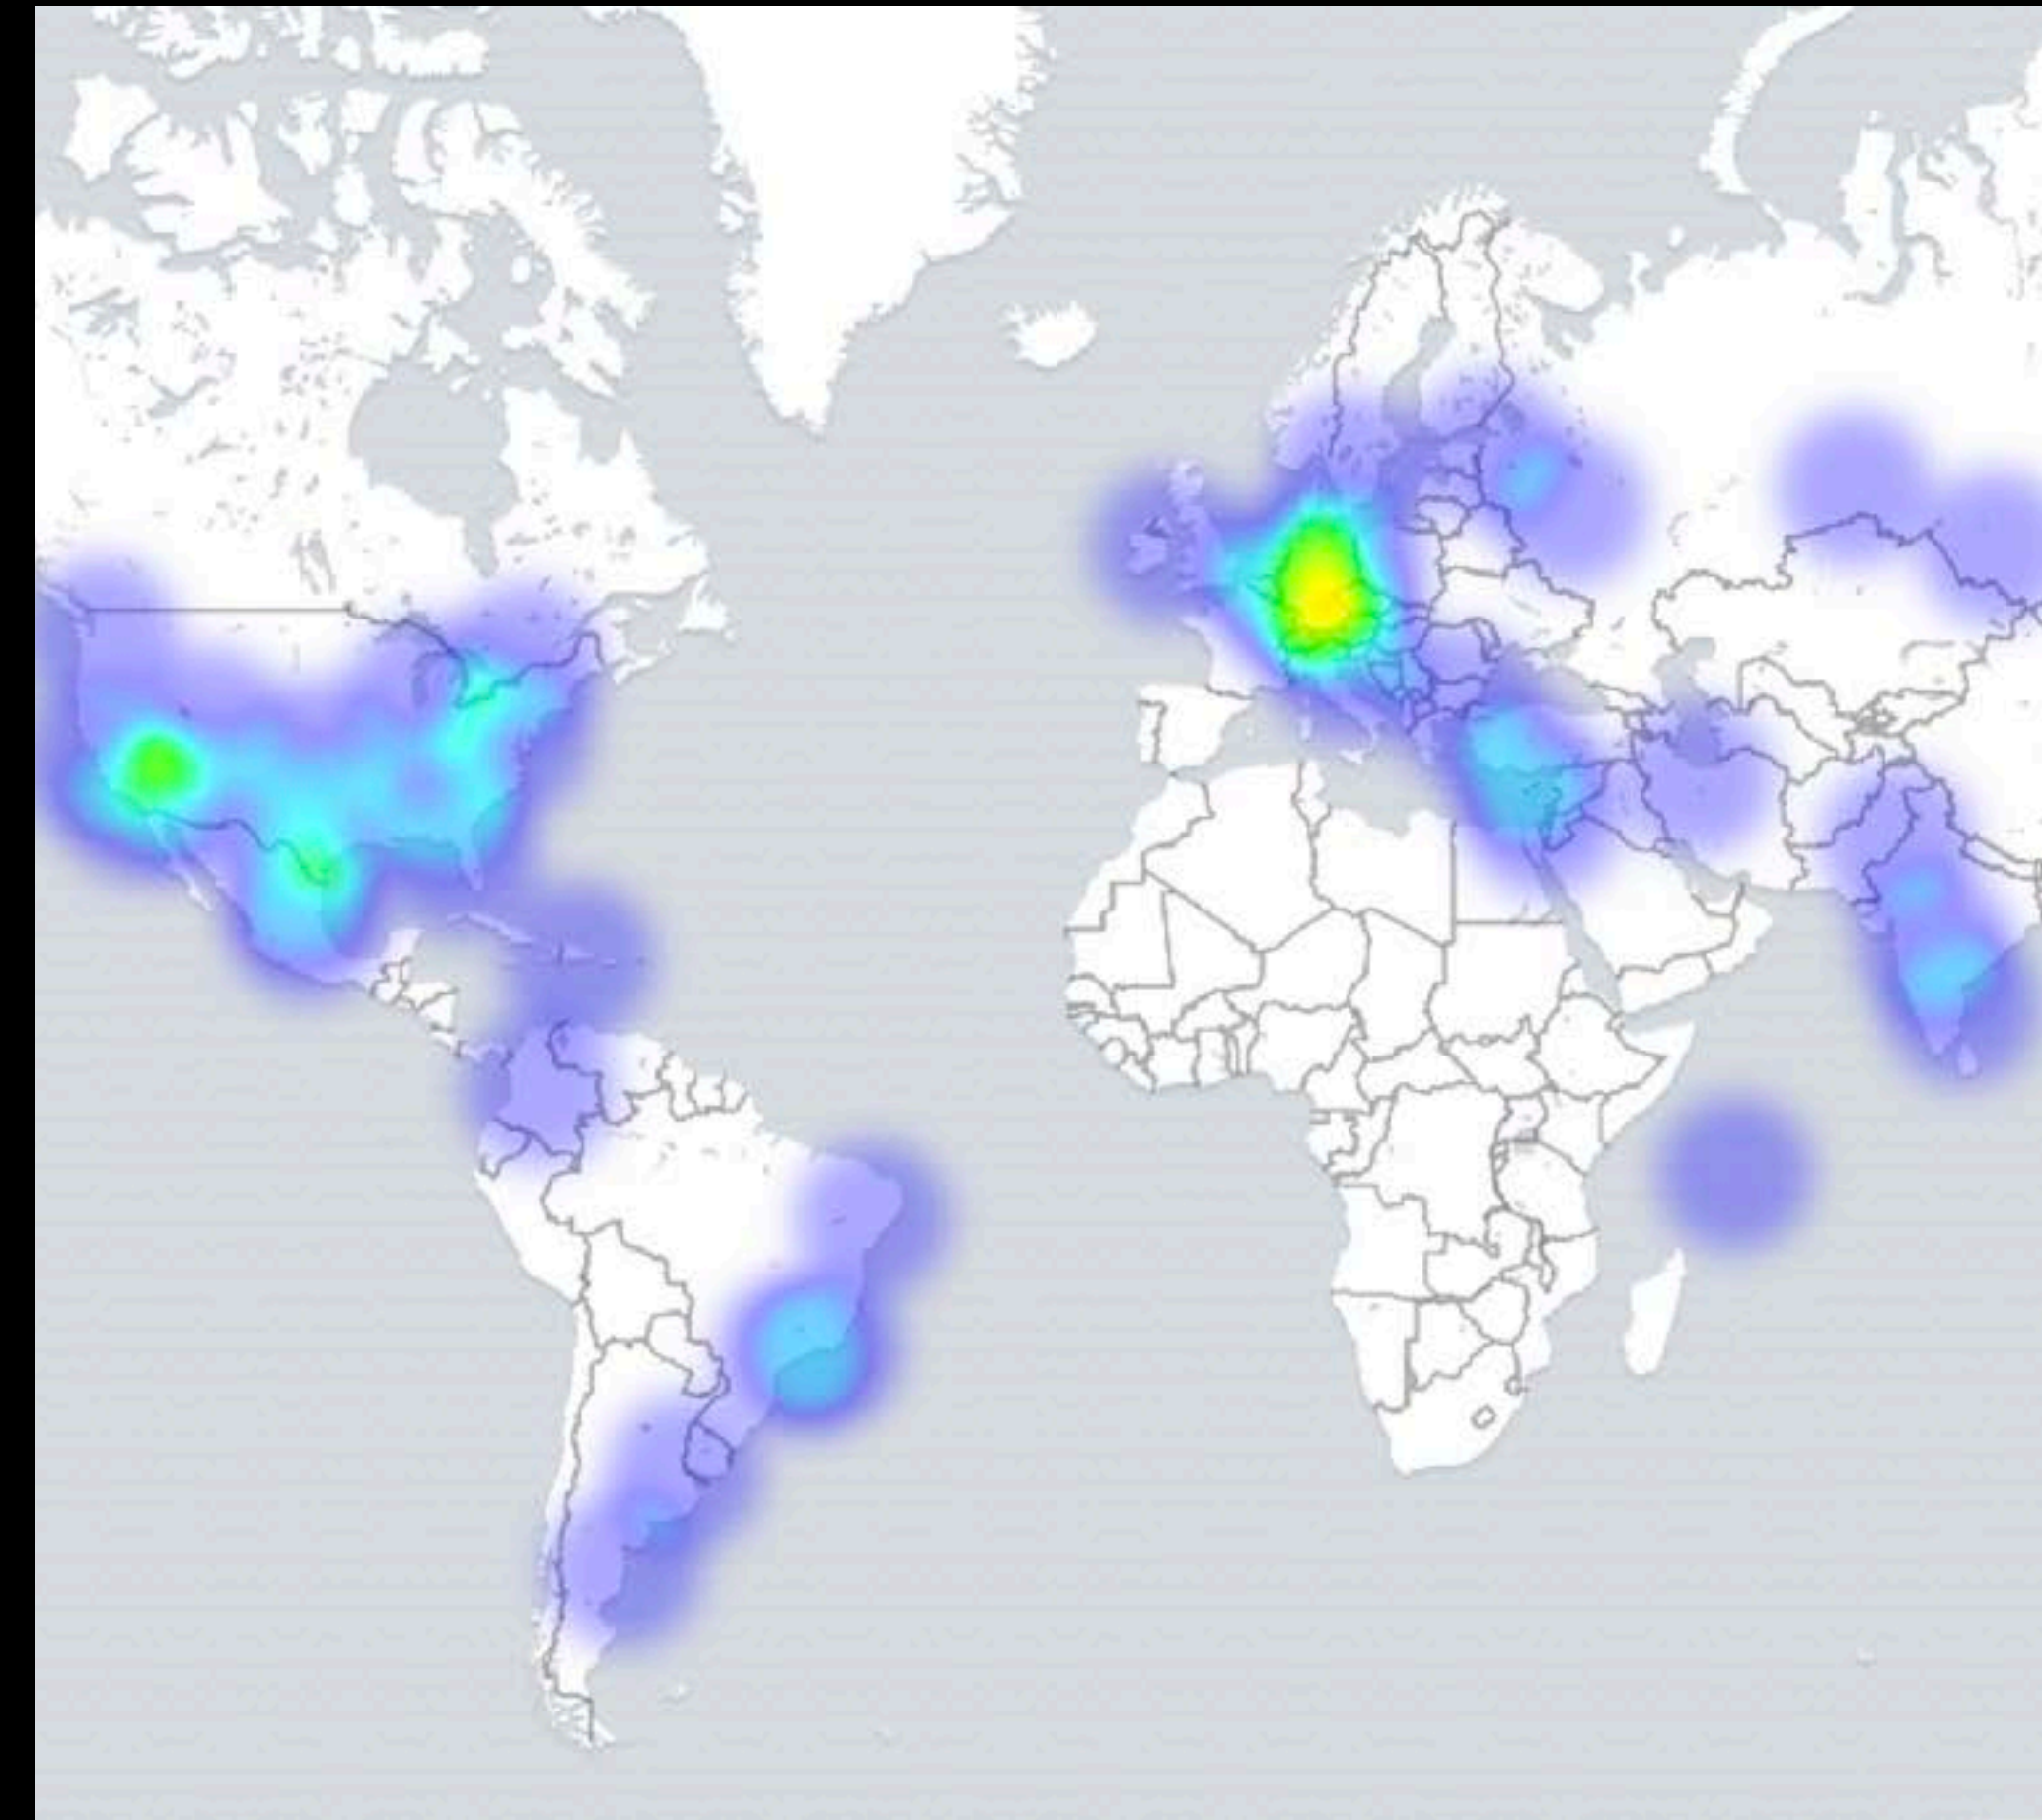
- source
- destination
- packet/byte counts
- timestamps
- type of Service (ToS)
- application ports
- input/output interfaces
- and more



# NETFLOW – DATA COLLECTION

NetFlow generation at strategic locations: e.g. Internet uplinks, core routers, data centre links

Potential implementations: Cisco kit, Elastic PacketBeat on Linux with PF\_RING, or NTOPNG



GeoIP based graph of NetFlow sources



# INTRUSION DETECTION ISN'T DEAD



*"IDS as a security technology is going to disappear."*

Richard Stiennon, Gartner Research Director, June 2003.



Deployment considerations:

- Inline blocking or passive detection?
- Location, location, location!  
e.g. Internet uplink, Extranet, data centre links, client networks



# DNS, THE ONE TRUE KING



Logging and analysing DNS traffic can be a challenge, particularly if you host your own DNS services or have a large network.

Potential implementations:

- Log client requests and server responses on your DNS servers (not popular with DNS server administrators).
- Set-up passive DNS collectors like lionmsg or ncap



# GET A HANDLE ON YOUR DATA

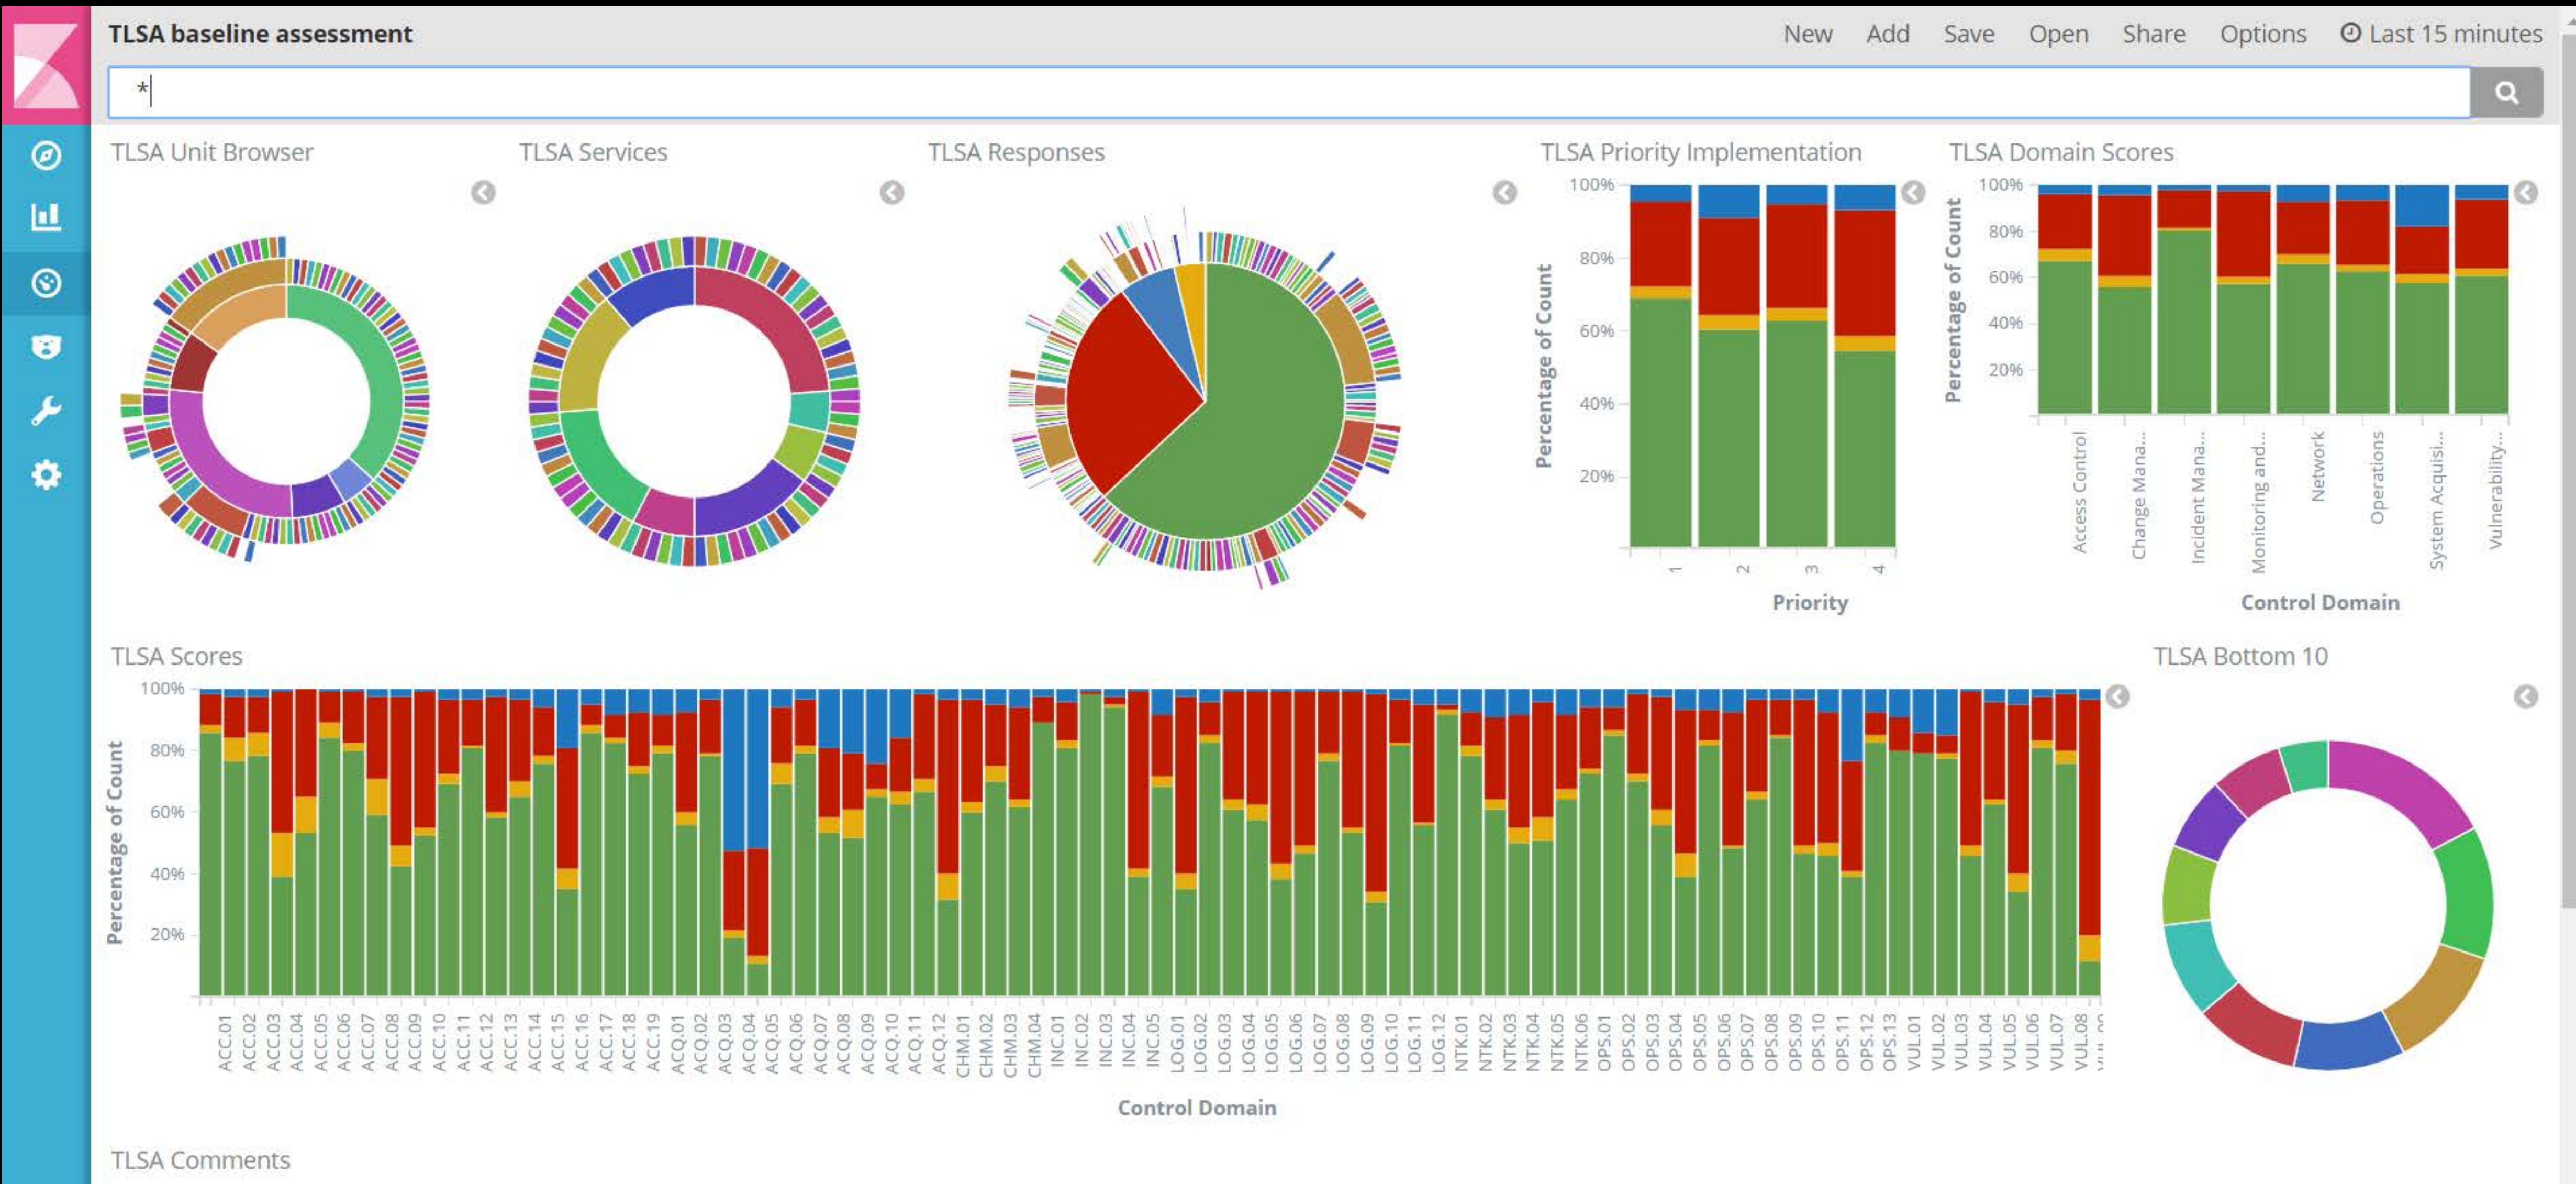


Just collecting all possible logs, events, and alarms does not help making sense out of them! Key lessons learned so far:

- NTP time source, UTC based logging, and ISO 8601 date
- Just the facts – filter and prepare your data
- Normalise your data  
e.g. 2001:420:1101:1::A vs 2001:420:1101:1:0:0:0:a
- Maintain consistent keys and curate metadata



# COMPLIANCE INFORMATION





# COMPLIANCE INFORMATION

Systems report compliance violations to your SIEM solution:

- Endpoint visibility (e.g. osquery, Santa, zentral)
- Server auditing (e.g. OpenSCAP, Lynis)
- Business application auditing (e.g. financial transactions)
- Automated vulnerability Scans (e.g. Nessus)





| Host          | Vulnerabilities ▲   |
|---------------|---|
| 172.26.21.251 | <div><div></div><div>23</div><div></div><div>219</div></div>              |
| 172.26.21.100 | <div><div></div><div>23</div><div>66</div><div>8</div><div>74</div></div> |
| 172.26.21.103 | <div><div></div><div>9</div><div>72</div><div></div><div>70</div></div>   |
| 172.26.21.220 | <div><div></div><div>21</div><div>7</div><div>99</div></div>              |
| 172.26.21.106 | <div><div></div><div>6</div><div>47</div><div></div><div>59</div></div>   |
| 172.26.21.148 | <div><div></div><div>11</div><div></div><div>81</div></div>               |
| 172.26.21.10  | <div><div></div><div>14</div><div>7</div><div>70</div></div>              |
| 172.26.21.160 | <div><div></div><div>10</div><div></div><div>74</div></div>               |
| 172.26.21.2   | <div><div></div><div>7</div><div></div><div>79</div></div>                |
| 172.26.21.18  | <div><div></div><div>17</div><div>6</div><div>64</div></div>              |
| 172.26.21.159 | <div><div></div><div>6</div><div>22</div><div>7</div><div>50</div></div>  |
| 172.26.21.17  | <div><div></div><div>14</div><div>6</div><div>64</div></div>              |
| 172.26.21.155 | <div><div></div><div>6</div><div>20</div><div>7</div><div>52</div></div>  |
| 172.26.21.219 | <div><div></div><div>8</div><div></div><div>72</div></div>                |
| 172.26.21.104 | <div><div></div><div>6</div><div>31</div><div>5</div><div>37</div></div>  |
| 172.26.21.109 | <div><div></div><div>7</div><div></div><div>72</div></div>                |
| 172.26.21.147 | <div><div></div><div>7</div><div></div><div>71</div></div>                |

Scan Details

Name:

Comprehensive Scan

Folder:

My Scans

Status:

Completed

Policy:

this is the scan to use for PCI

Shared with:

1 user

Scanner:

US Cloud Scanner

Targets:

172.26.21.0/24

Start time:

Wed May 14 12:52:57 2014


End time:

Wed May 14 14:49:40 2014

Elapsed:

2 hours

Vulnerabilities





| www.jordan.ox.ac.uk |   |        |     |      |  |  |       |
|---------------------|---|--------|-----|------|--|--|-------|
| Results Summary     |   |        |     |      |  |  |       |
| Critical            | High  | Medium | Low | Info |  |  | Total |
| 0                   | 2   | 16     | 2   | 75   |  |  | 95    |
| Results Details     |   |        |     |      |  |  |       |
| 0/tcp               |   |        |     |      |  |  |       |
| <div></div>         | 11936 - OS Identification   |        |     |      |  |  | [-/+] |
| <div></div>         | 12053 - Host Fully Qualified Domain Name (FQDN) Resolution        |        |     |      |  |  | [-/+] |
| <div></div>         | 19506 - Nessus Scan Information                                   |        |     |      |  |  | [-/+] |
| <div></div>         | 25220 - TCP/IP Timestamps Supported                               |        |     |      |  |  | [-/+] |
| <div></div>         | 45590 - Common Platform Enumeration (CPE)                         |        |     |      |  |  | [-/+] |
| <div></div>         | 54615 - Device Type   |        |     |      |  |  | [-/+] |
| <div></div>         | 66334 - Patch Report  |        |     |      |  |  | [-/+] |
| <div></div>         | 84239 - Debugging Log Report                                      |        |     |      |  |  | [-/+] |
| 0/udp               |   |        |     |      |  |  |       |
| <div></div>         | 10287 - Traceroute Information                                    |        |     |      |  |  | [-/+] |
| 80/tcp              |   |        |     |      |  |  |       |
| <div></div>         | 42479 - CGI Generic SQL Injection (2nd pass)                      |        |     |      |  |  | [-/+] |
| <div></div>         | 48926 - CGI Generic 2nd Order SQL Injection Detection (potential) |        |     |      |  |  | [-/+] |
| <div></div>         | 39466 - CGI Generic XSS (quick test)                              |        |     |      |  |  | [-/+] |
| <div></div>         | 44136 - CGI Generic Cookie Injection Scripting                    |        |     |      |  |  | [-/+] |
| <div></div>         | 44670 - Web Application SQL Backend Identification                |        |     |      |  |  | [-/+] |



# THE TOOL MAKETH THE TEAM



**SIEM**

Security Information and  
Event Management

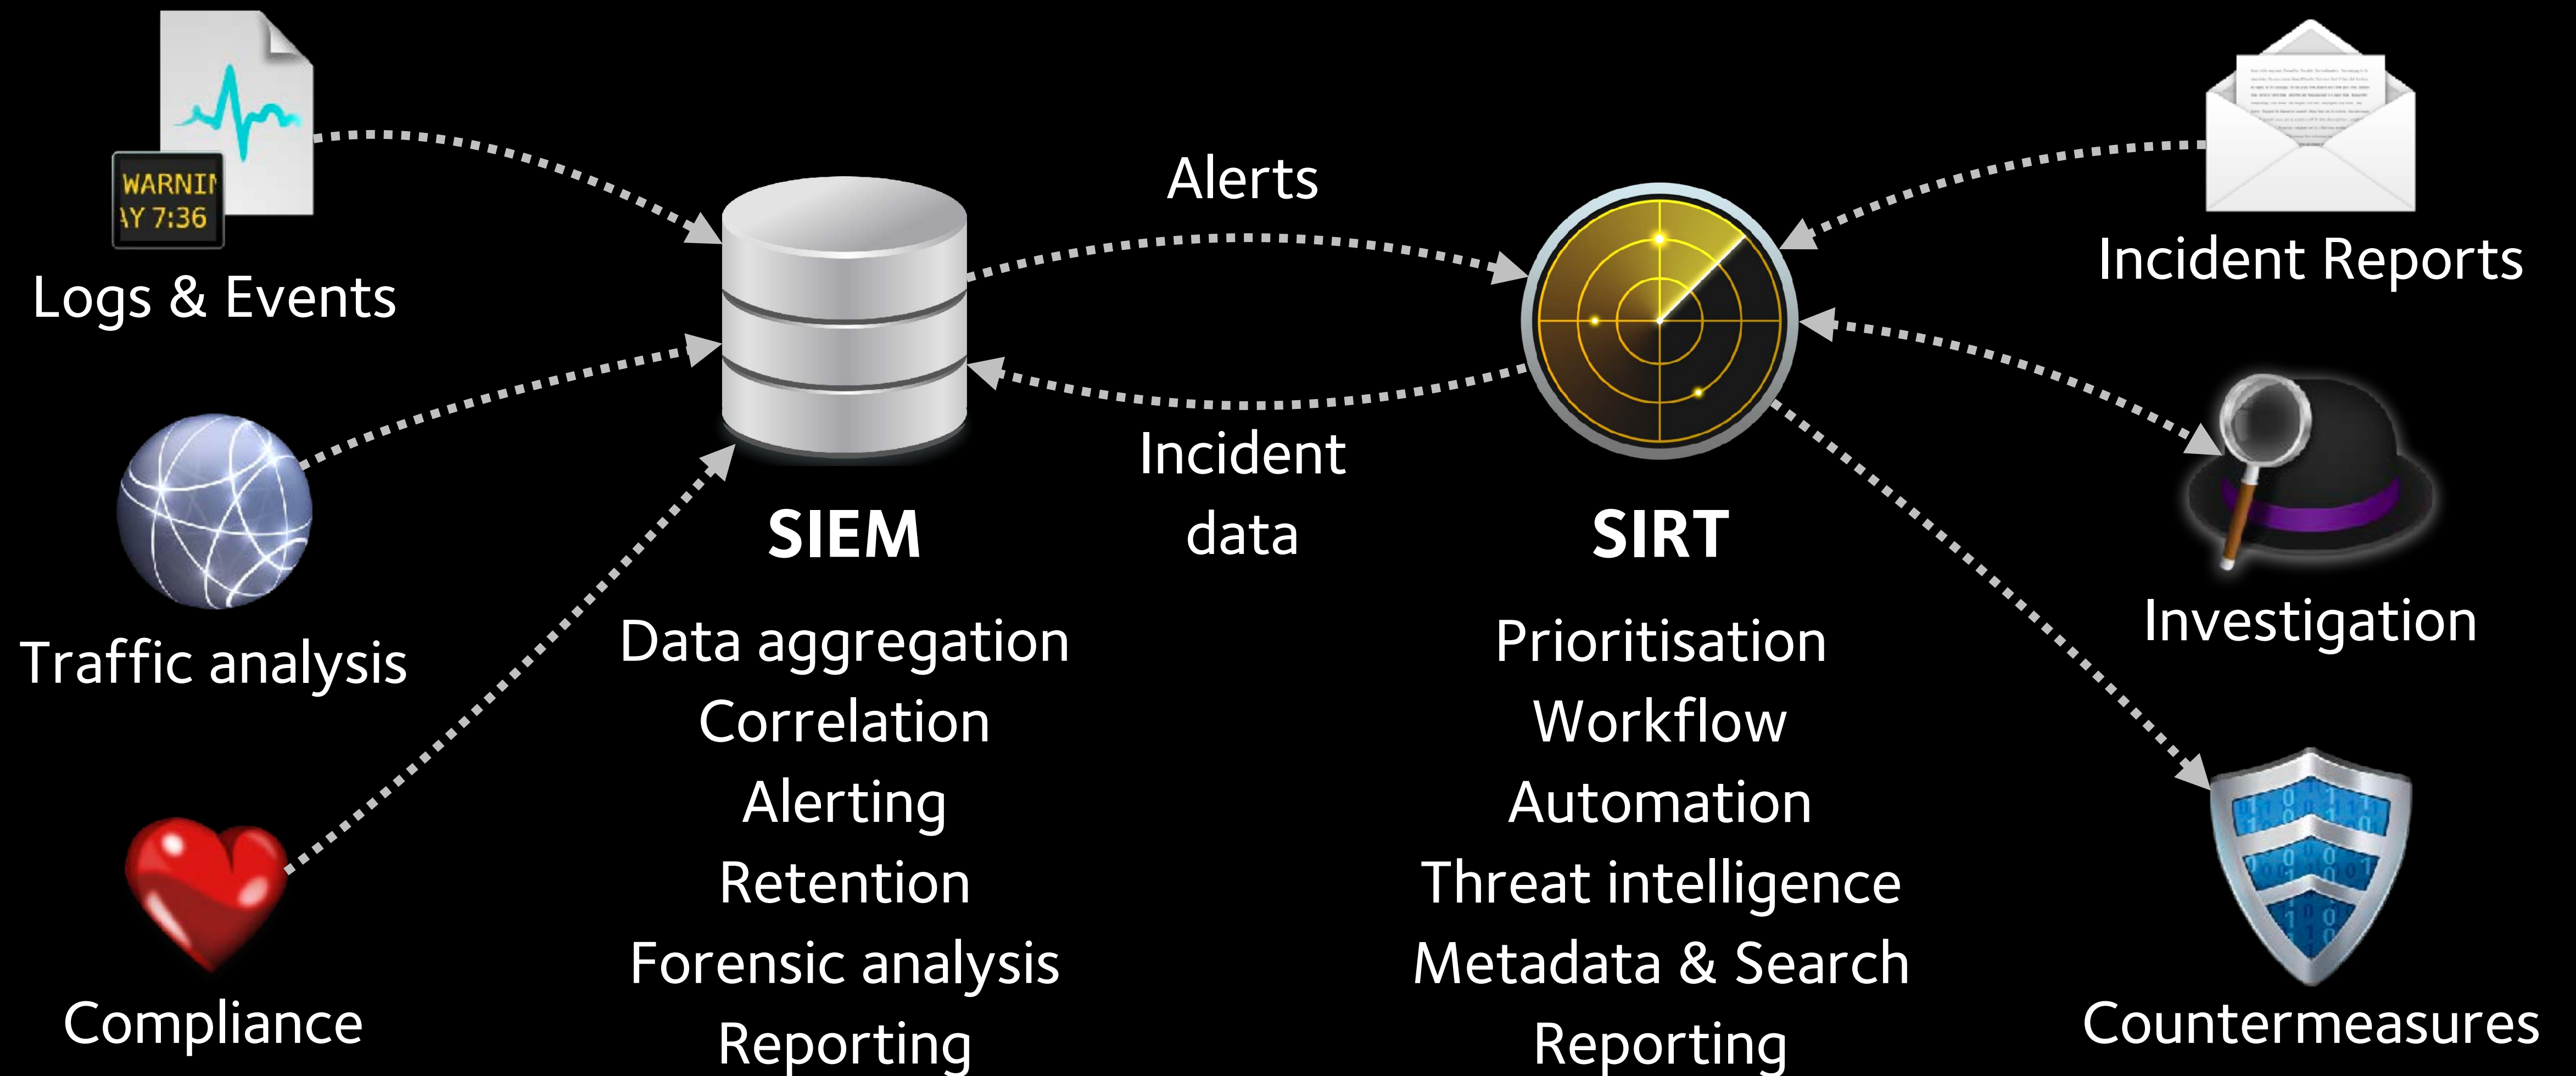


**SIRT**

Security Incident  
Response Tracker



# THE TOOL MAKETH THE TEAM





# THE TOOL MAKETH THE TEAM



**SIEM**



**SIRT**




<https://elastic.co>



<https://bestpractical.com/rtir/>



# RTIR: REQUEST TRACKER FOR INCIDENT RESPONSE



UNIVERSITY OF  
OXFORD

Lookup '192.168.1.2'

New ticket in

Incident R

192.168.1.2

Current Incident: #5

| # Subject      | Status<br>Owner | Last Updated<br>Told | Created<br>Due        | Priority<br>Time Left |
|----------------|-----------------|----------------------|-----------------------|-----------------------|
| 5 Possible DoS | open<br>root    | 1 minute ago         | 6 weeks ago<br>2 days | 50                    |

Incidents: 192.168.1.2

Search

| # Subject      | Status   | Priority | Actions              |
|----------------|----------|----------|----------------------|
| 2 a problem!   | resolved | 50       | [Merge][Investigate] |
| 5 Possible DoS | open     | 50       | [Investigate]        |

Investigations: 192.168.1.2

Search

Link

| # Subject       | Status | Priority | Actions |
|-----------------|--------|----------|---------|
| 10 Possible DoS | open   | 0        | [Link]  |

Incident Reports: 192.168.1.2

Search

Link

| # Subject      | Status   | Priority | Actions |
|----------------|----------|----------|---------|
| 1 a problem!   | resolved | 0        | [Link]  |
| 4 Possible DoS | resolved | 0        | [Link]  |

Blocks: 192.168.1.2

Search

Link

Create

| # Subject      | Status        | Priority | Actions |
|----------------|---------------|----------|---------|
| 3 a problem!   | removed       | 0        | [Link]  |
| 6 Possible DoS | post incident | 0        | [Link]  |



# RTIR: REQUEST TRACKER FOR INCIDENT RESPONSE



Incident Tracker shall be also used for

- Phishing
- Third party thread intelligence
- Bulletins
- Vulnerability Scanning
- General advise and guidance
- ...

Lookup '192.168.1.2'

^ Current Incident: #5

| # | Subject      | Status | Owner |
|---|--------------|--------|-------|
| 5 | Possible DoS | open   | root  |

^ Incidents: 192.168.1.2

| # | Subject      | Status   | Priority | Actions      |
|---|--------------|----------|----------|--------------|
| 2 | a problem!   | resolved | 50       | [Merge][Inv] |
| 5 | Possible DoS | open     | 50       | [Investigat  |

^ Incident Reports: 192.168.1.2


| # | Subject      | Status   | Priority |
|---|--------------|----------|----------|
| 1 | a problem!   | resolved | 0        |
| 4 | Possible DoS | resolved | 0        |



# YES WE ARE OPEN!

Oxford's network security model has never been based on a trusted internal network.

- The University acts as an ISP to its colleges and departments.
- There is no perimeter firewall for the organisation.
- User services are deployed to the Internet (exceptions apply).





# BEYOND CORP

*“BeyondCorp is an enterprise security model that builds upon 6 years of building zero trust networks at Google”*

## PRINCIPLES

- Connecting from a particular network must not determine which services you can access.
- Access to services is granted based on what we know about you and your device.
- All access to services must be authenticated, authorized and encrypted.

## BeyondCorp A New Approach to Enterprise Security

RORY WARD AND BETSY BEYER



Rory Ward is a senior reliability engineering manager in Google Ireland. He previously worked in Ireland at Valista, in Silicon Valley at AOL, Netscape, Kiva, and General Magic, and in Los Angeles at Retix. He has a BSc in computer applications from Dublin City University. [rcryward@google.com](mailto:rcryward@google.com)



Betsy Beyer is a technical writer specializing in virtualization software for Google SRE in NYC. She has previously provided documentation for Google Data Center and Hardware Operations teams. Before moving to New York, Betsy was a lecturer in technical writing at Stanford University. She holds degrees from Stanford and Tulane. [bbeyer@google.com](mailto:bbeyer@google.com)

Virtually every company today uses firewalls to enforce perimeter security. However, this security model is problematic because, when that perimeter is breached, an attacker has relatively easy access to a company's privileged intranet. As companies adopt mobile and cloud technologies, the perimeter is becoming increasingly difficult to enforce. Google is taking a different approach to network security. We are removing the requirement for a privileged intranet and moving our corporate applications to the Internet.

Since the early days of IT infrastructure, enterprises have used perimeter security to protect and gate access to internal resources. The perimeter security model is often compared to a medieval castle: a fortress with thick walls, surrounded by a moat, with a heavily guarded single point of entry and exit. Anything located outside the wall is considered dangerous, while anything located inside the wall is trusted. Anyone who makes it past the drawbridge has ready access to the resources of the castle.

The perimeter security model works well enough when all employees work exclusively in buildings owned by an enterprise. However, with the advent of a mobile workforce, the surge in the variety of devices used by this workforce, and the growing use of cloud-based services, additional attack vectors have emerged that are stretching the traditional paradigm to the point of redundancy. Key assumptions of this model no longer hold: The perimeter is no longer just the physical location of the enterprise, and what lies inside the perimeter is no longer a blessed and safe place to host personal computing devices and enterprise applications.

While most enterprises assume that the internal network is a safe environment in which to expose corporate applications, Google's experience has proven that this faith is misplaced. Rather, one should assume that an internal network is as fraught with danger as the public Internet and build enterprise applications based upon this assumption.

Google's BeyondCorp initiative is moving to a new model that dispenses with a privileged corporate network. Instead, access depends solely on device and user credentials, regardless of a user's network location—be it an enterprise location, a home network, or a hotel or coffee shop. All access to enterprise resources is fully authenticated, fully authorized, and fully encrypted based upon device state and user credentials. We can enforce fine-grained access to different parts of enterprise resources. As a result, all Google employees can work successfully from any network, and without the need for a traditional VPN connection into the privileged network. The user experience between local and remote access to enterprise resources is effectively identical, apart from potential differences in latency.

### The Major Components of BeyondCorp

BeyondCorp consists of many cooperating components to ensure that only appropriately authenticated devices and users are authorized to access the requisite enterprise applications. Each component is described below (see Figure 1).




# MOST COMMON INCIDENTS





# INCIDENTS 09/10/2016-26/09/2017





## cybergangs

Scientific research targeted by hackers


EXCLUSIVE

---

Peter Yeung, Rosemary Bennett

---


September 5 2017, 12:01am, The Times





## cybergangs

Scientific research targeted by hackers



Source: Times research



The University of Oxford said there had been 515 cases of unauthorised access to its accounts or machines last year and UCL said that it experienced 57 successful attacks in 2016-17.





# UNAUTHORISED ACCESS

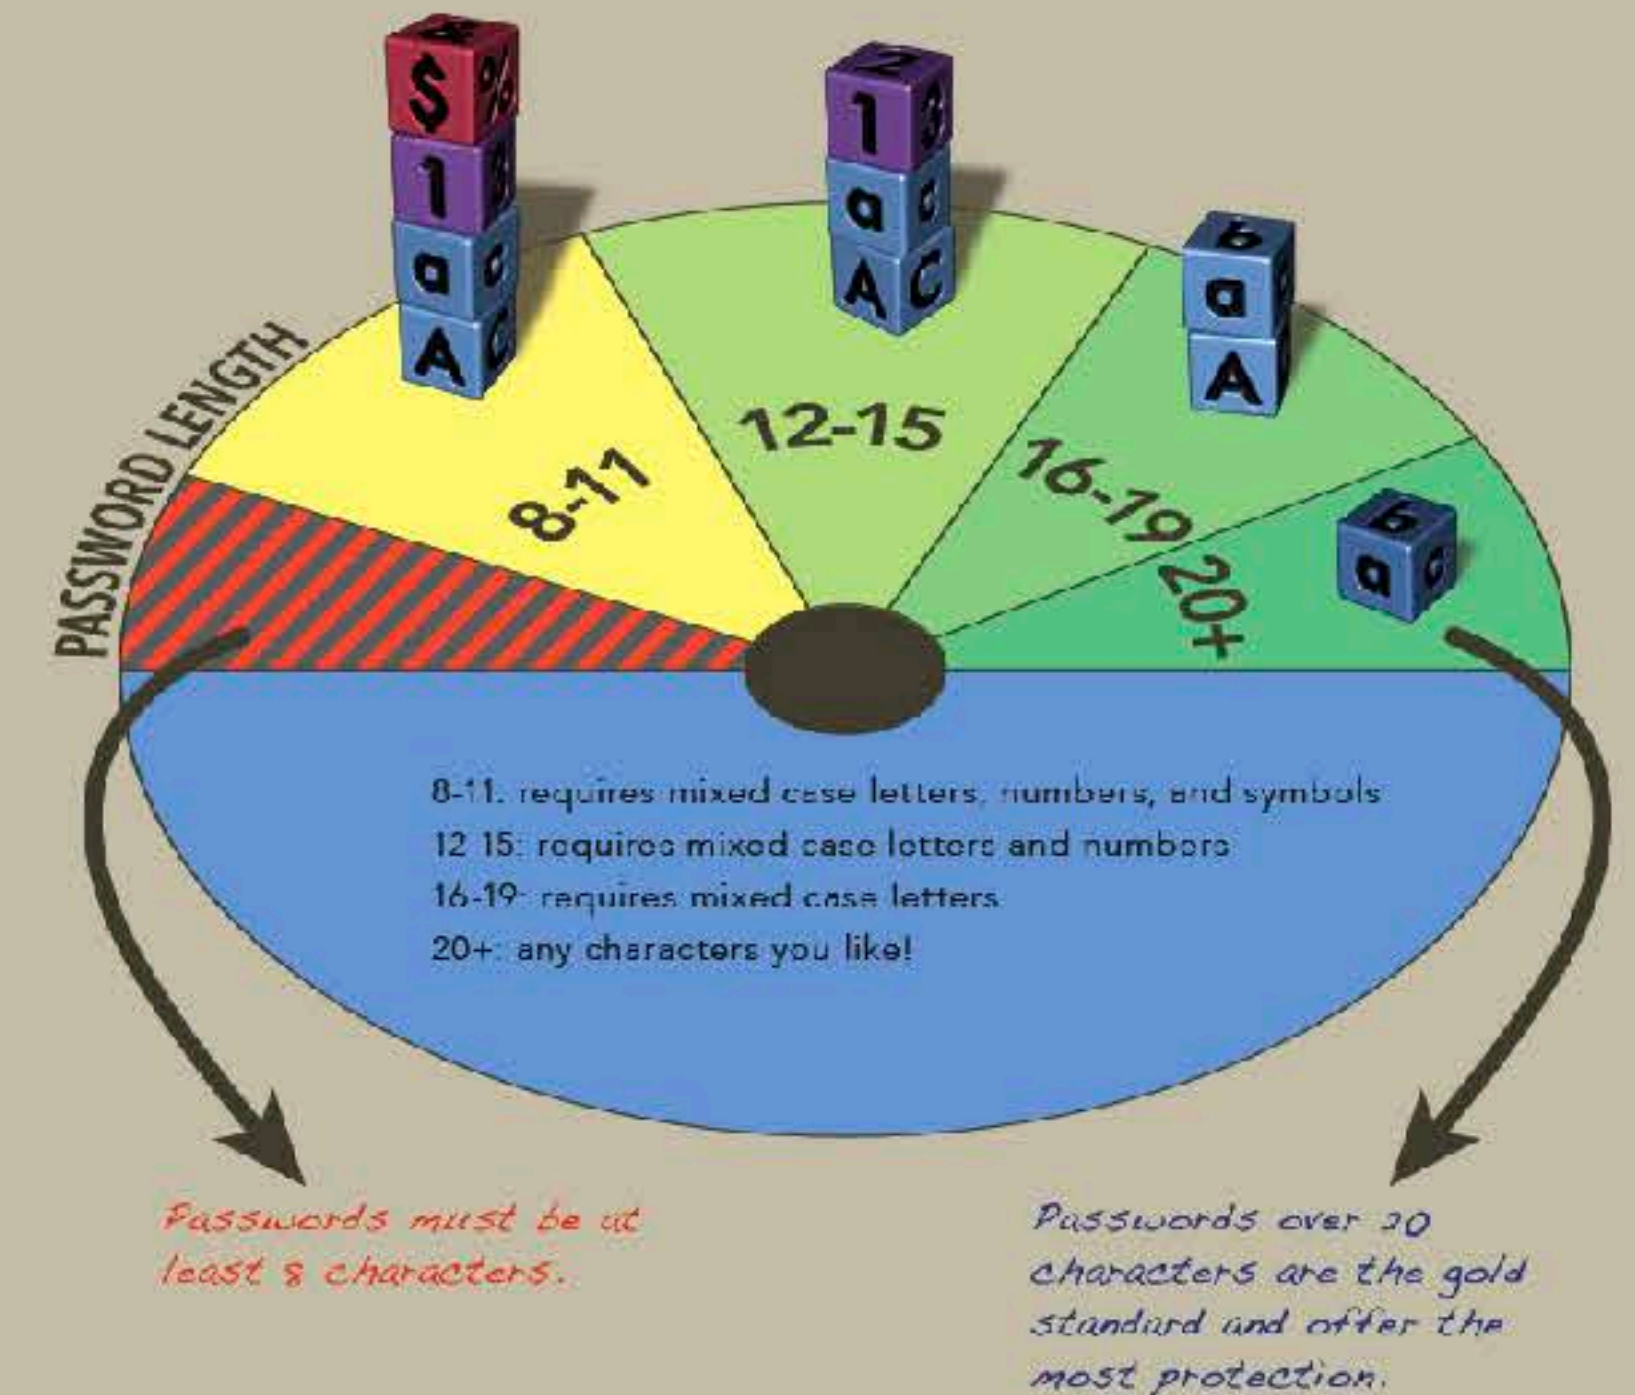
We observe the majority of accounts being phished as opposed to hacked.

A good password policy helps.

Ability to lock accounts and audit access is of prime importance.

## WHICH CHARACTERS ARE REQUIRED IN MY **PASSWORD**?

*HINT: it depends on password length!*



MAY WE RECOMMEND...

**16**  
OR MORE

Longer passwords are inherently more secure because it takes hackers longer to guess them when employing a brute force method. So make your password 16 characters or longer!



# UNAUTHORISED ACCESS

We observe the majority of accounts being phished as opposed to hacked.

A good password policy helps.

Ability to lock accounts and audit access is of prime importance.



Longer passwords are inherently more secure because it takes hackers longer to guess them when employing a brute force method. So make your password 16 characters or longer!

Because they only require upper and lower case letters, passwords that are 16 characters or longer are much easier to type on a mobile device.



*How on Earth can I come up with a password that long?!*



**EASY!**  
Select 4 random words.





# UNAUTHORISED ACCESS

We observe the majority of accounts being phished as opposed to hacked.

A good password policy helps.

Ability to lock accounts and audit access is of prime importance.

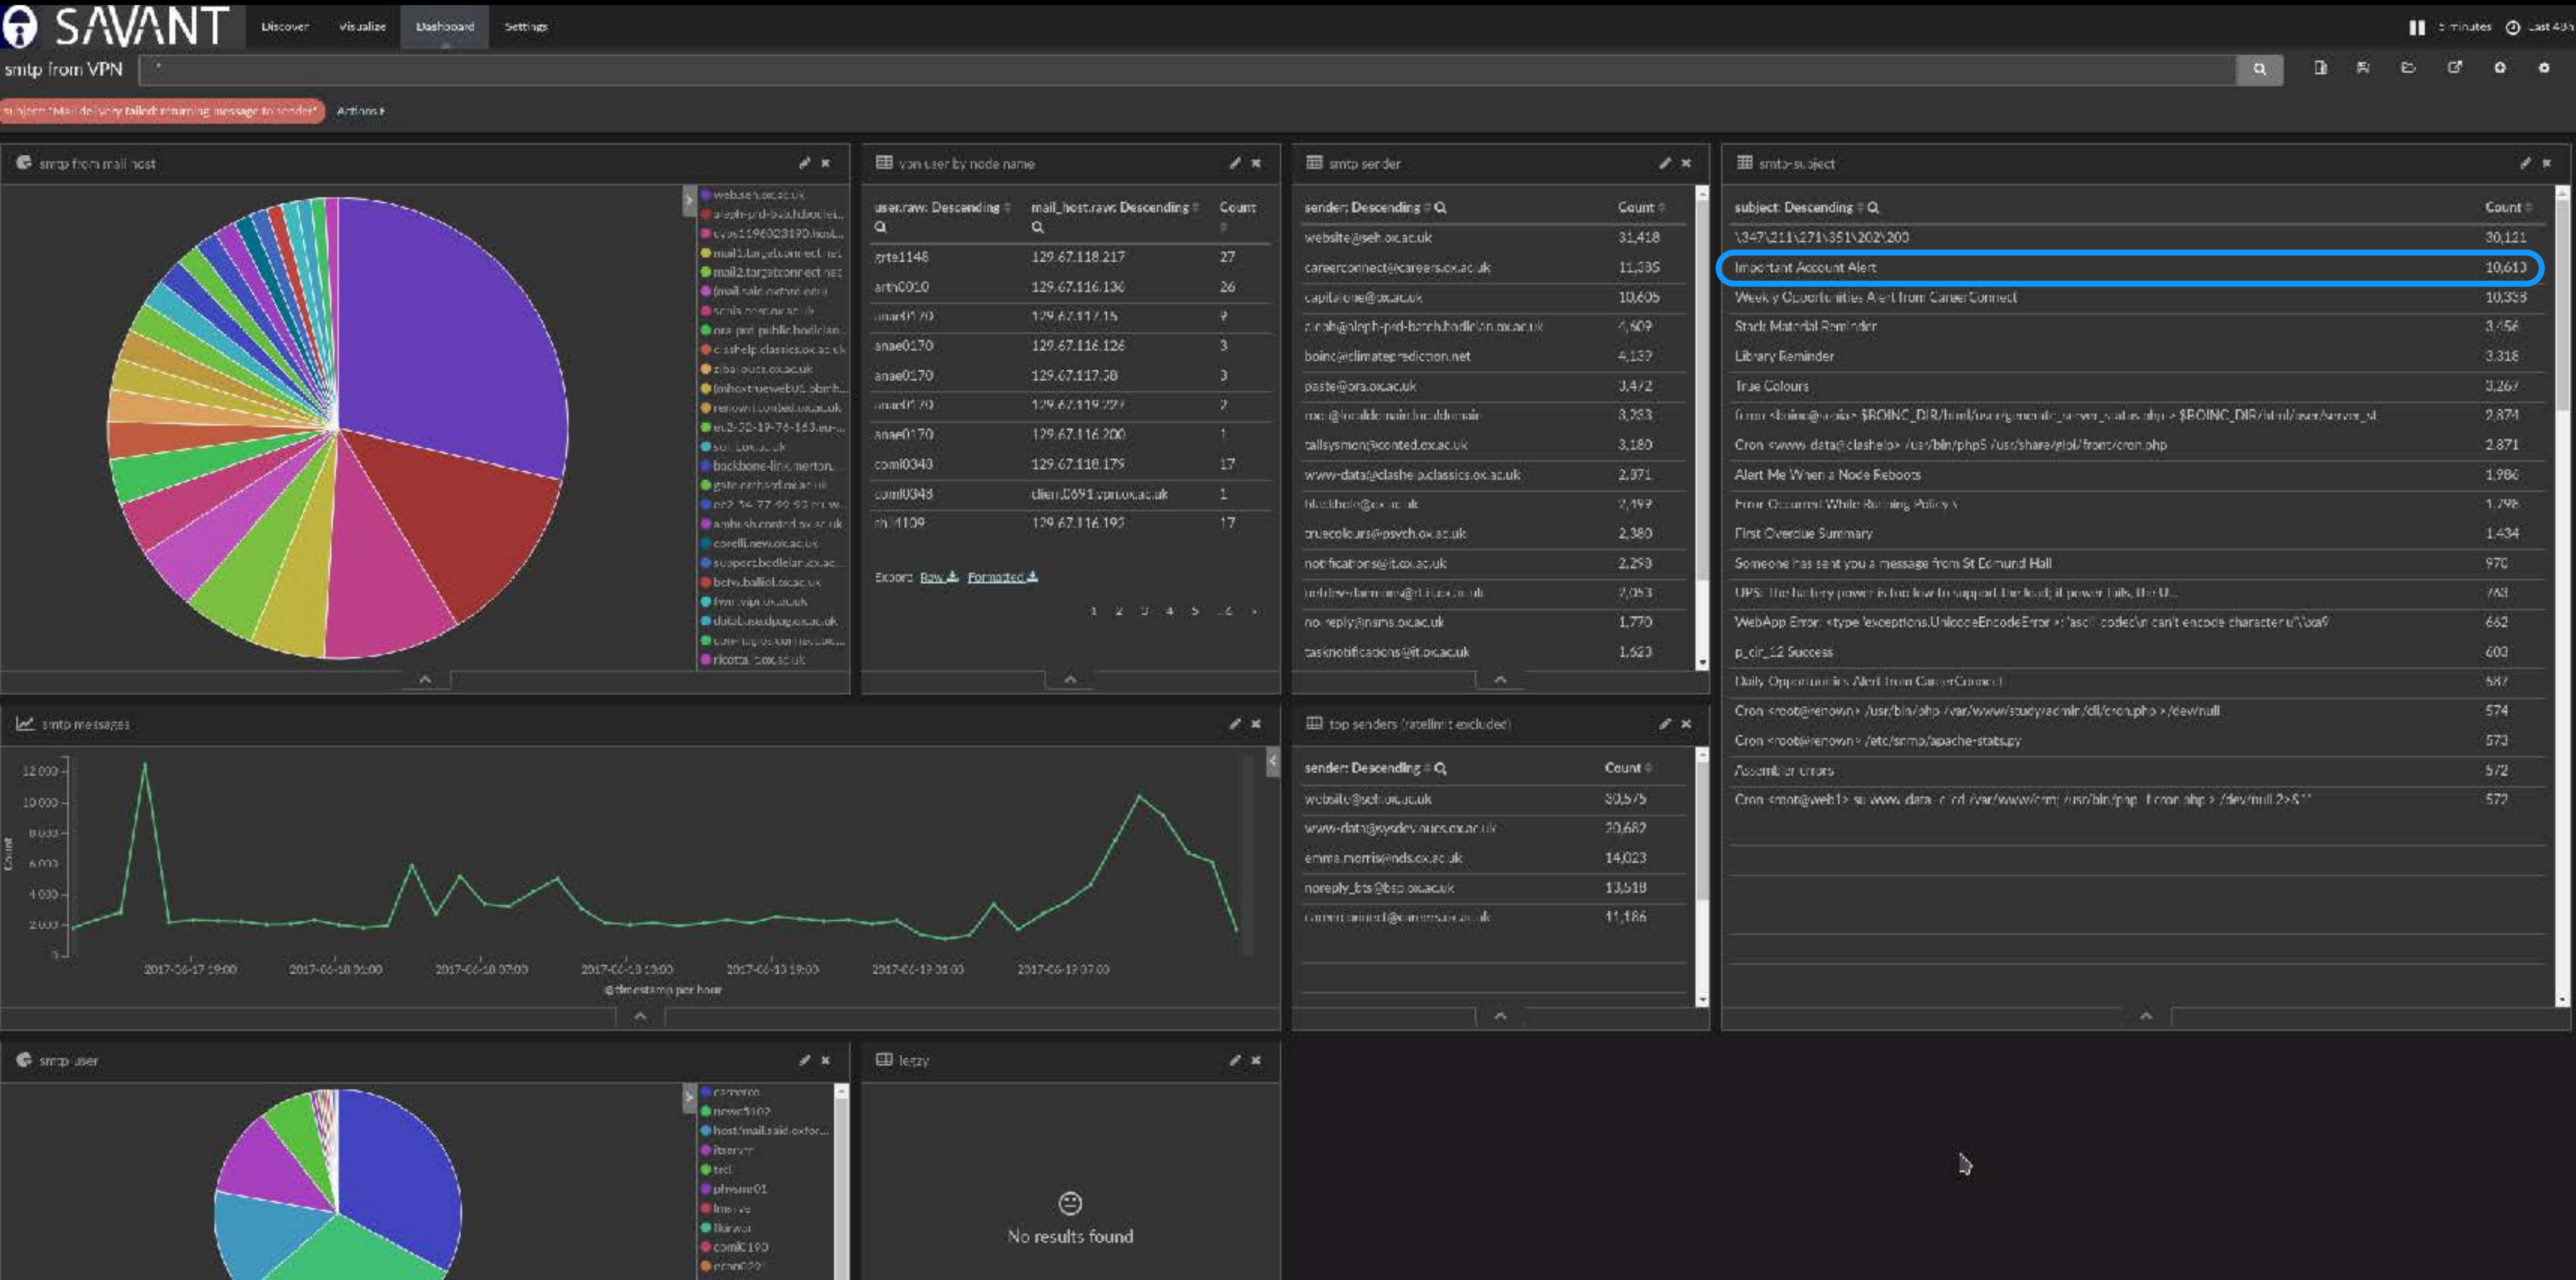




# EXAMPLE: SPAMMER

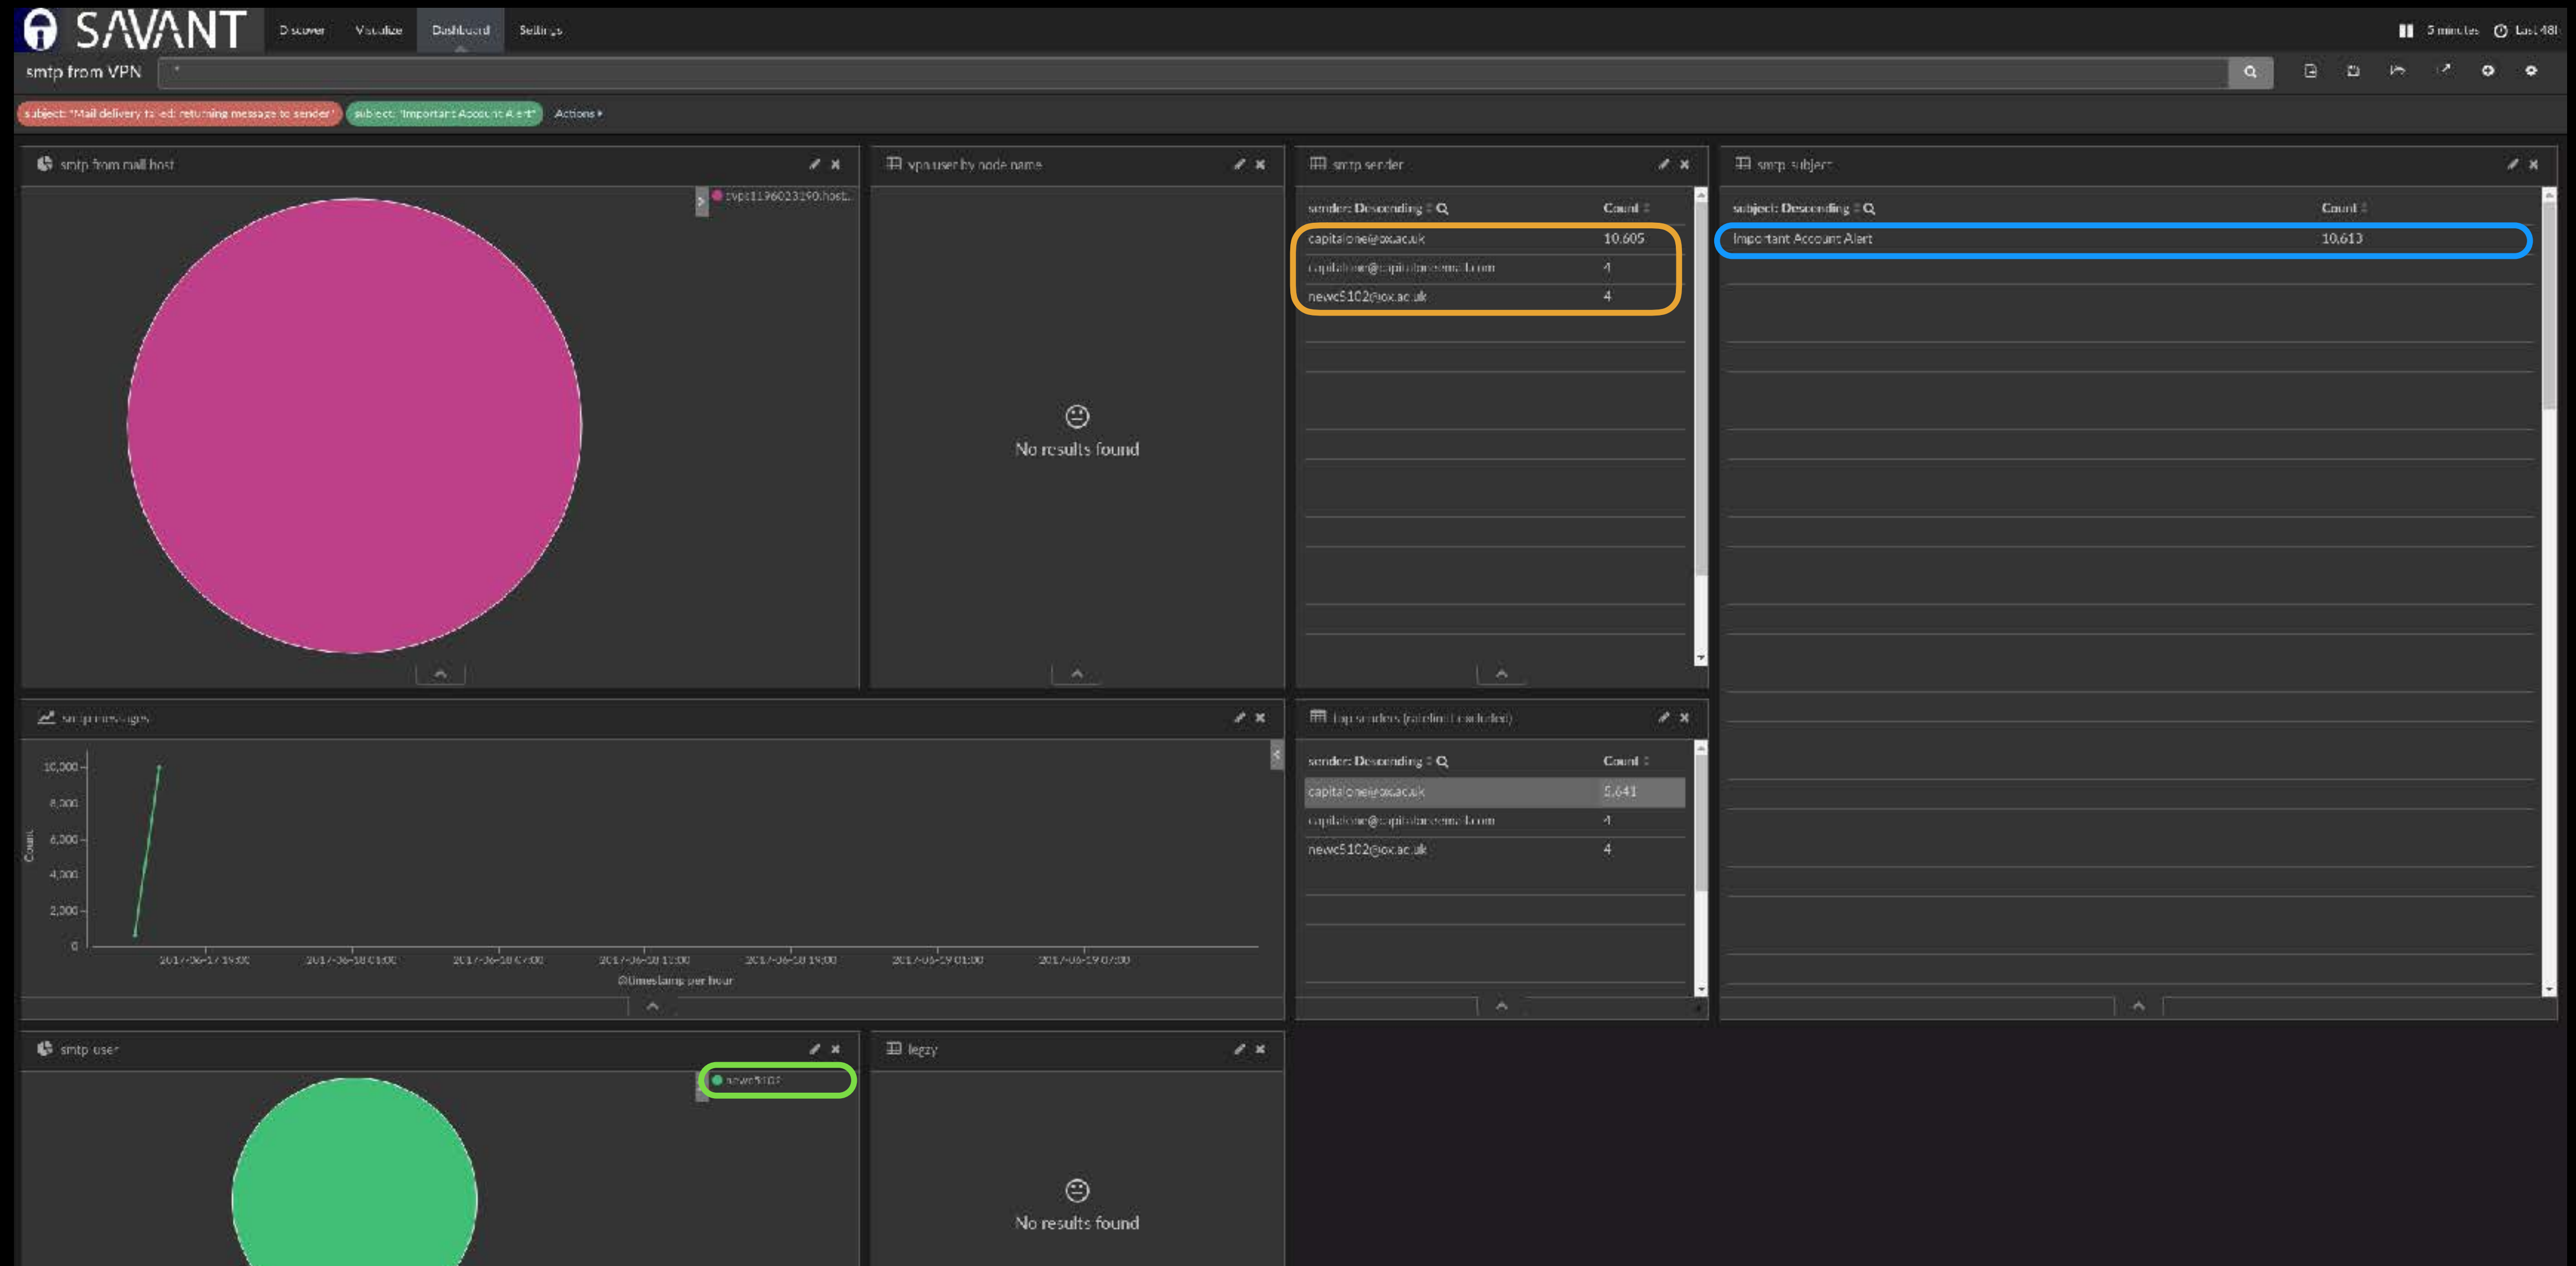


UNIVERSITY OF  
OXFORD



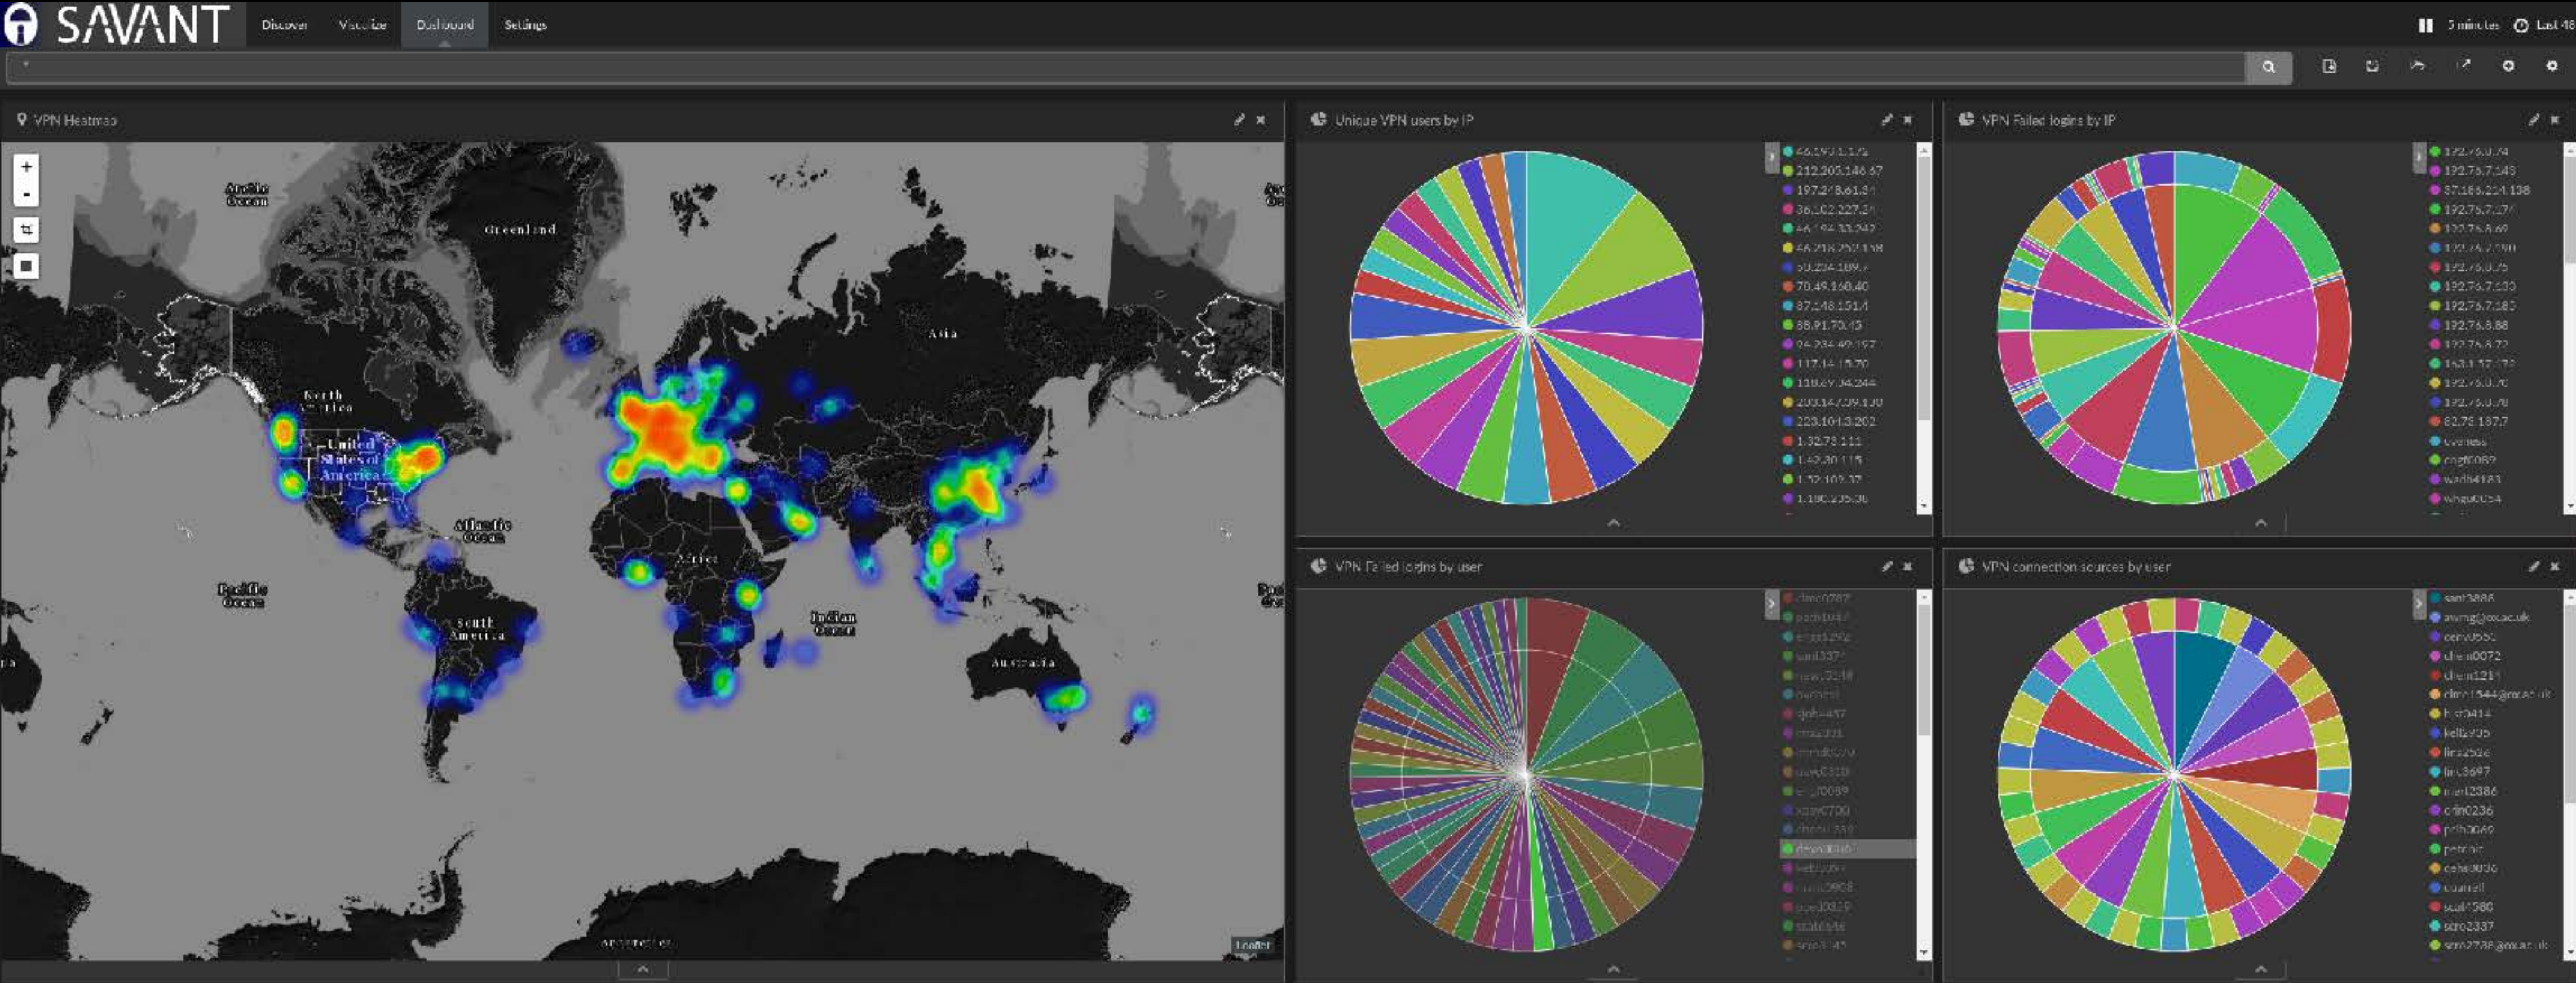


# EXAMPLE: SPAMMER

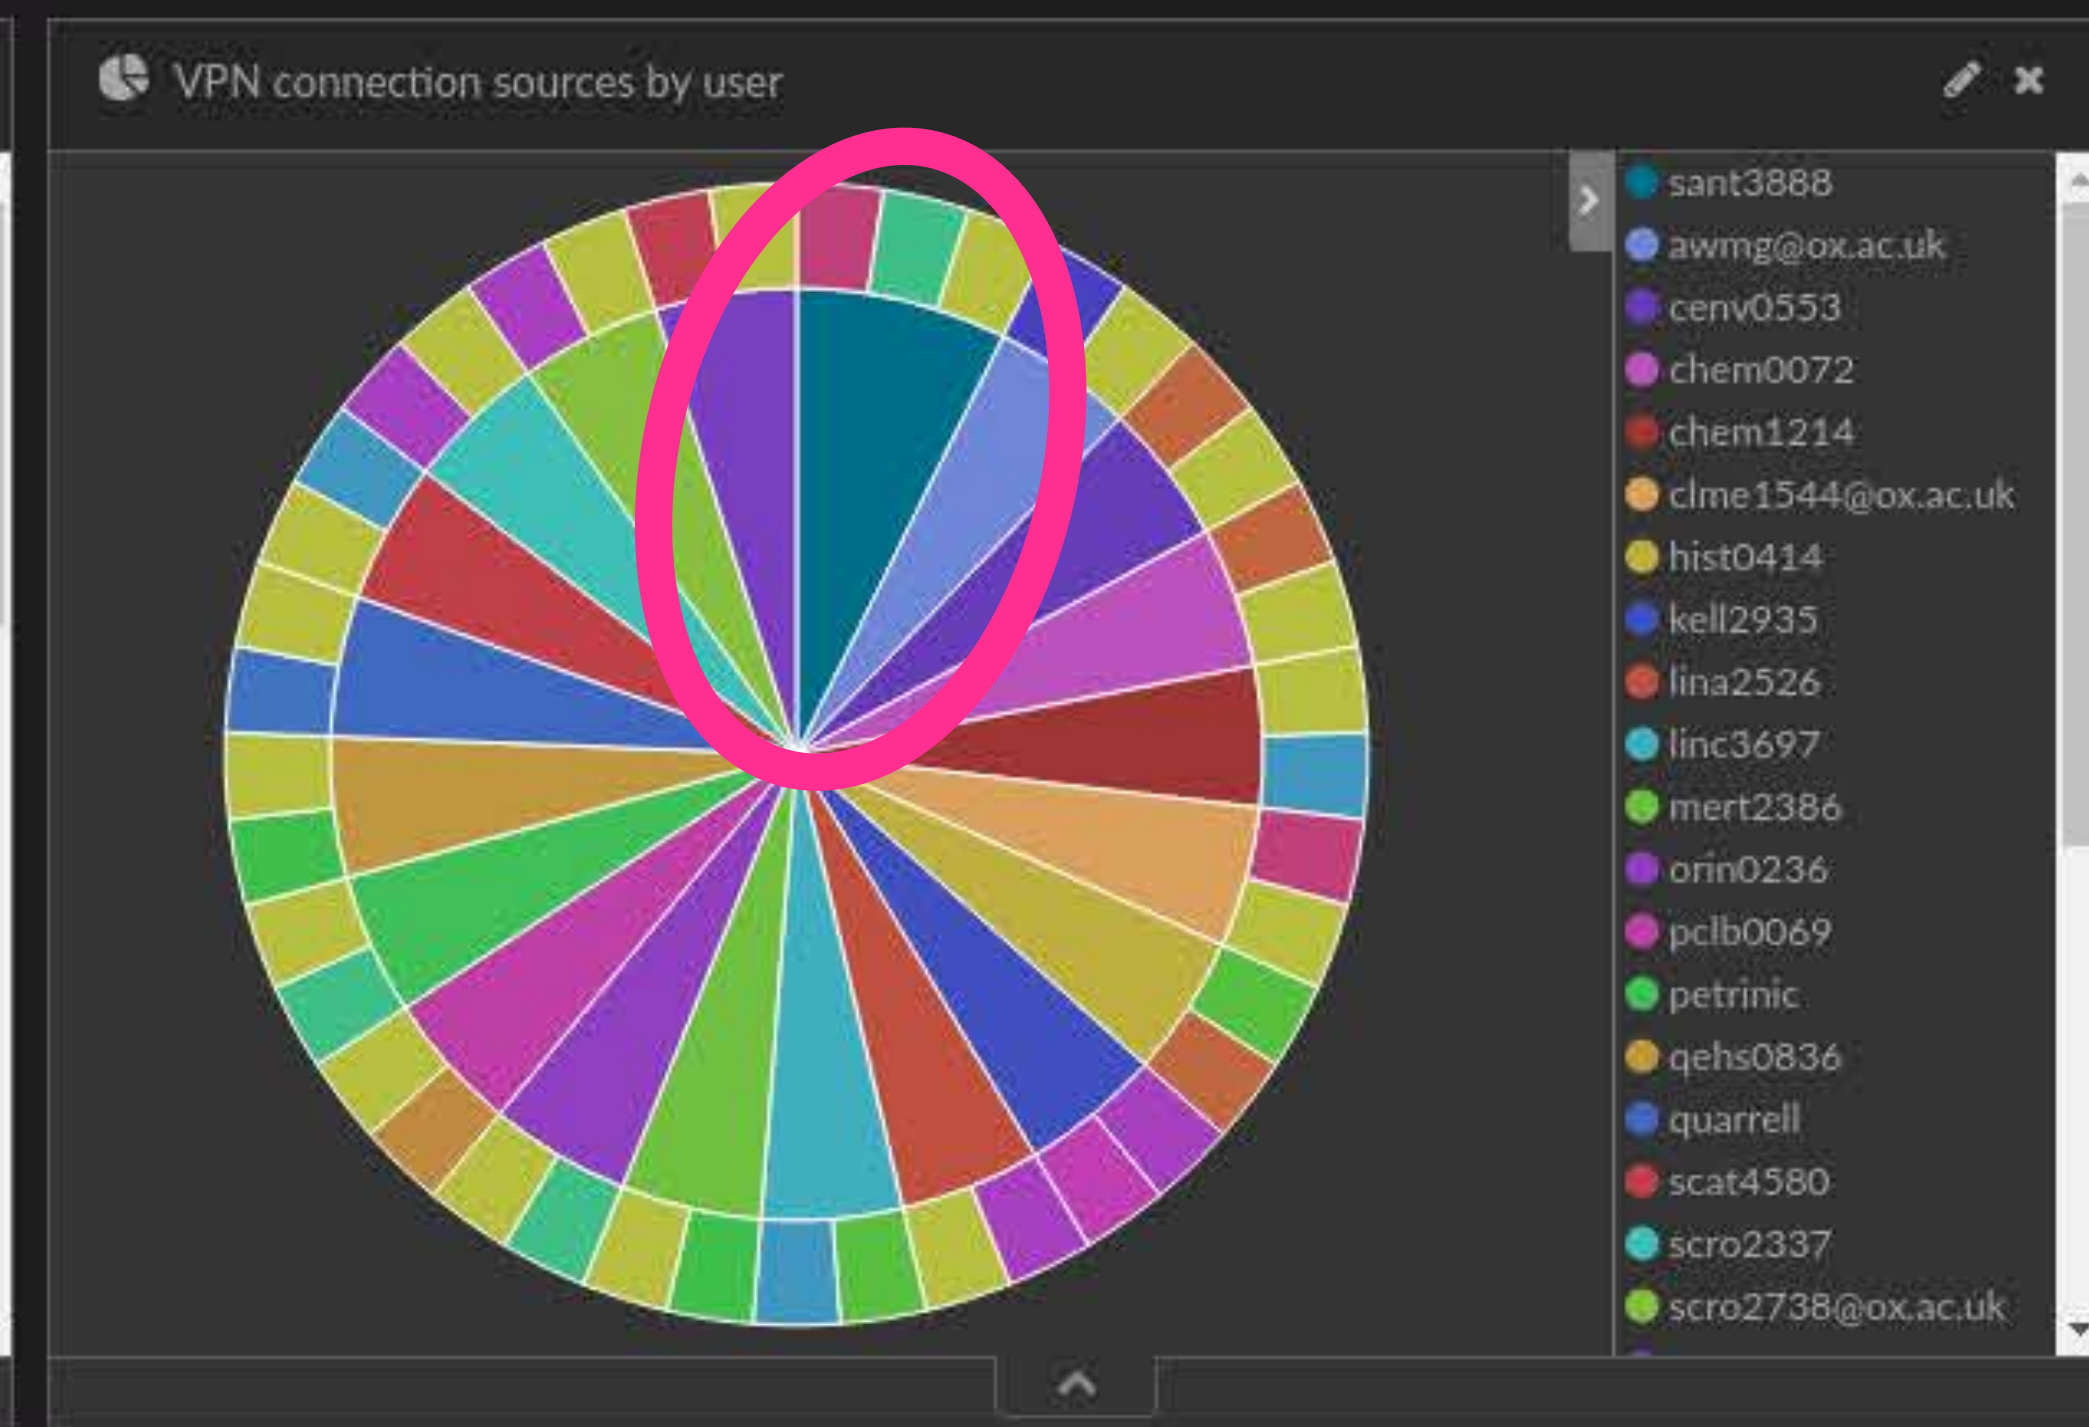
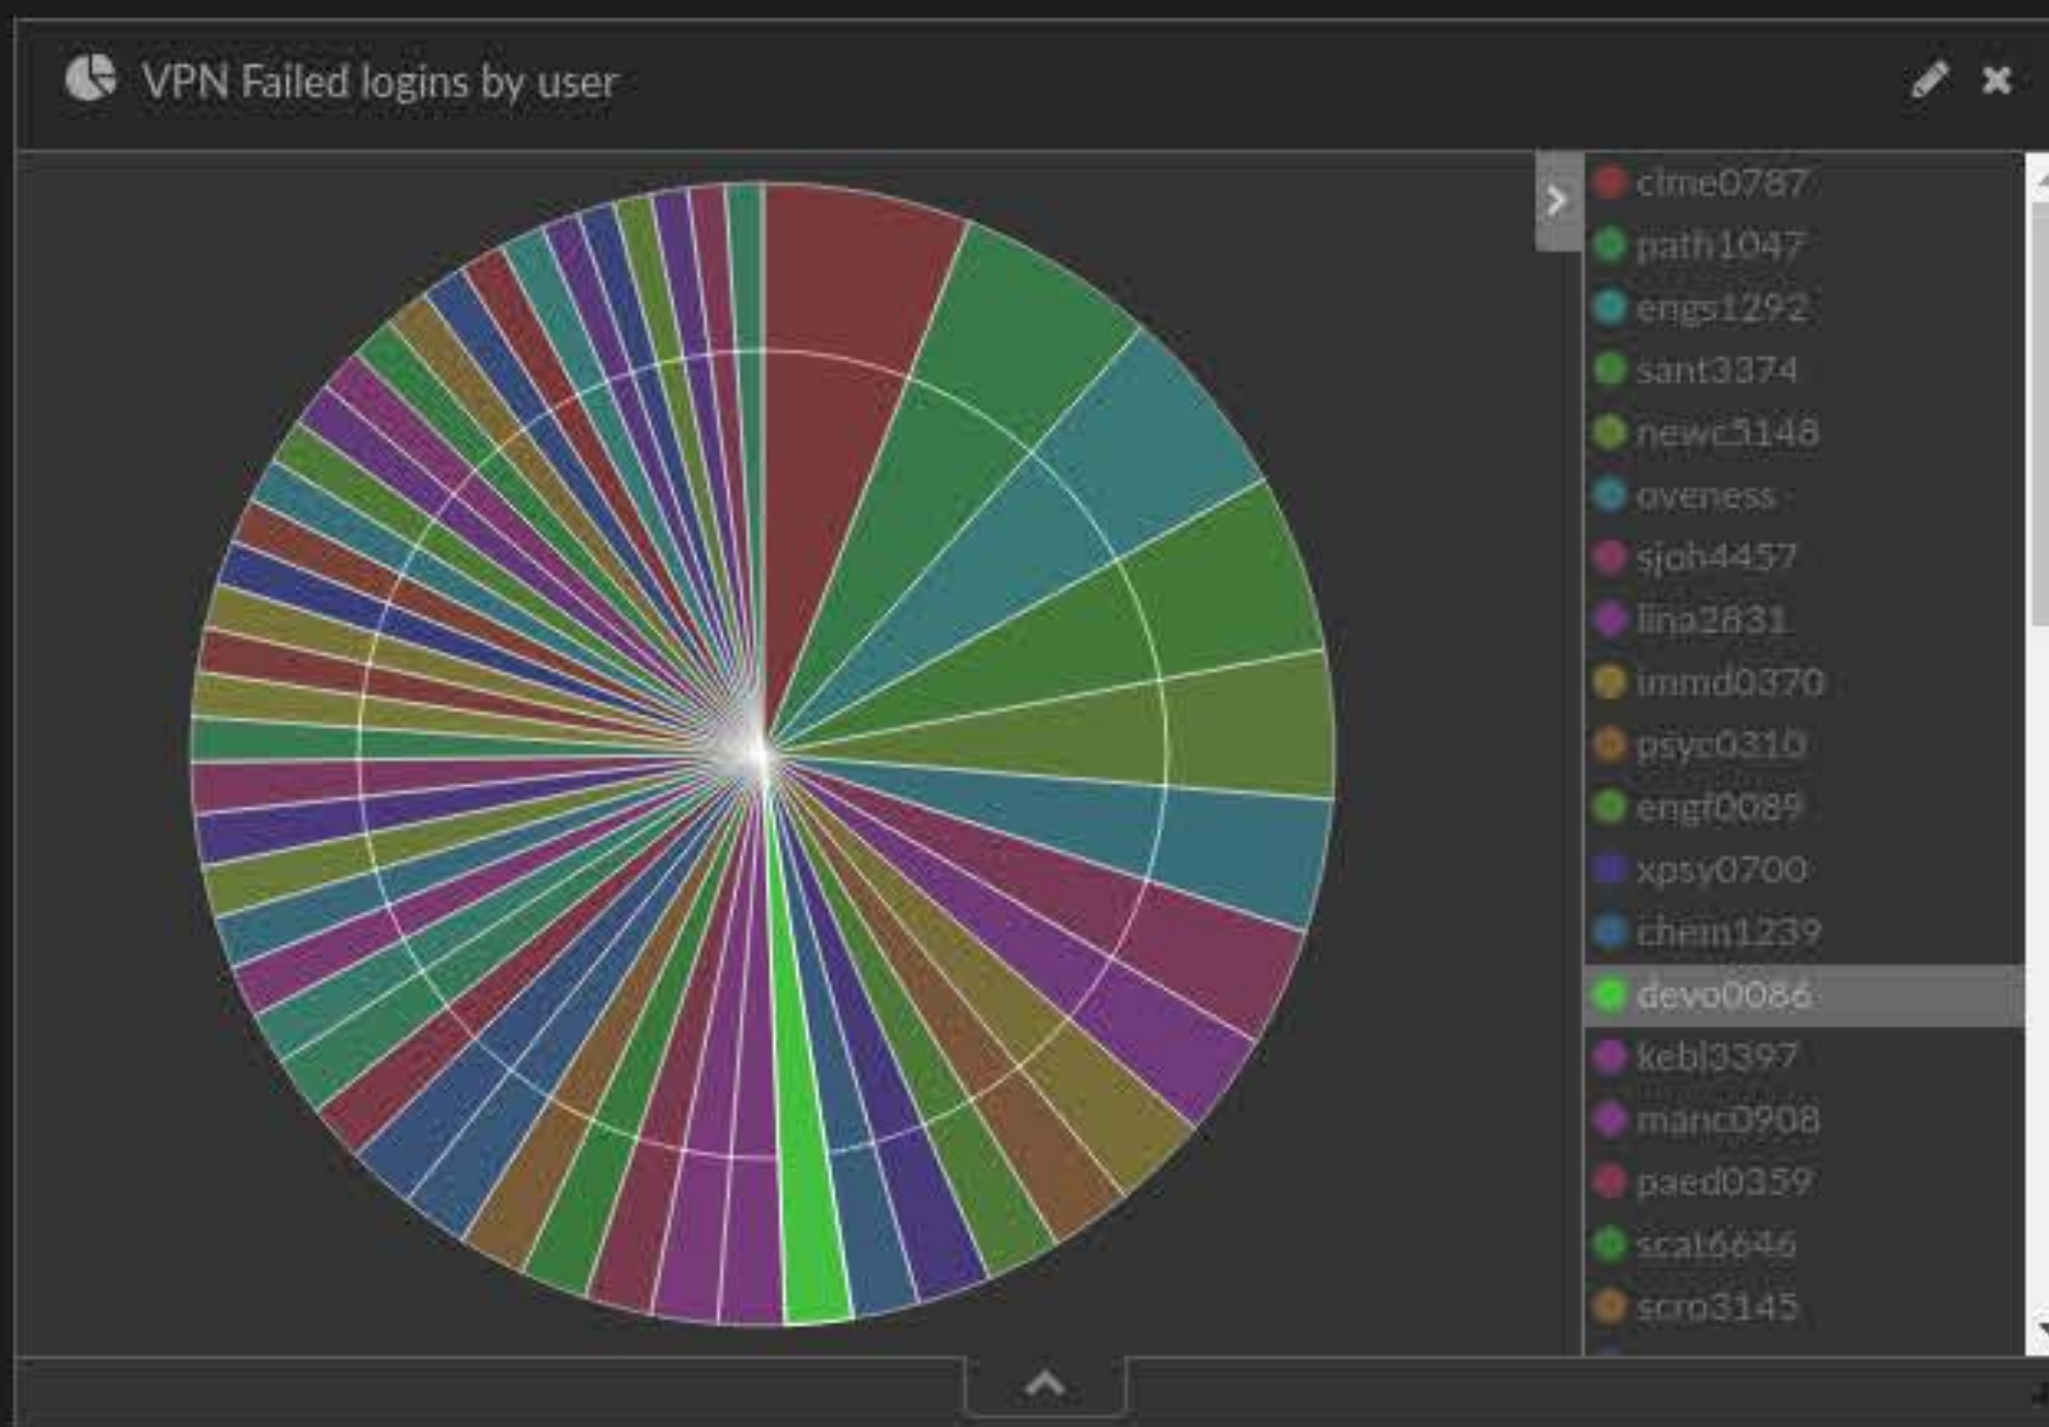

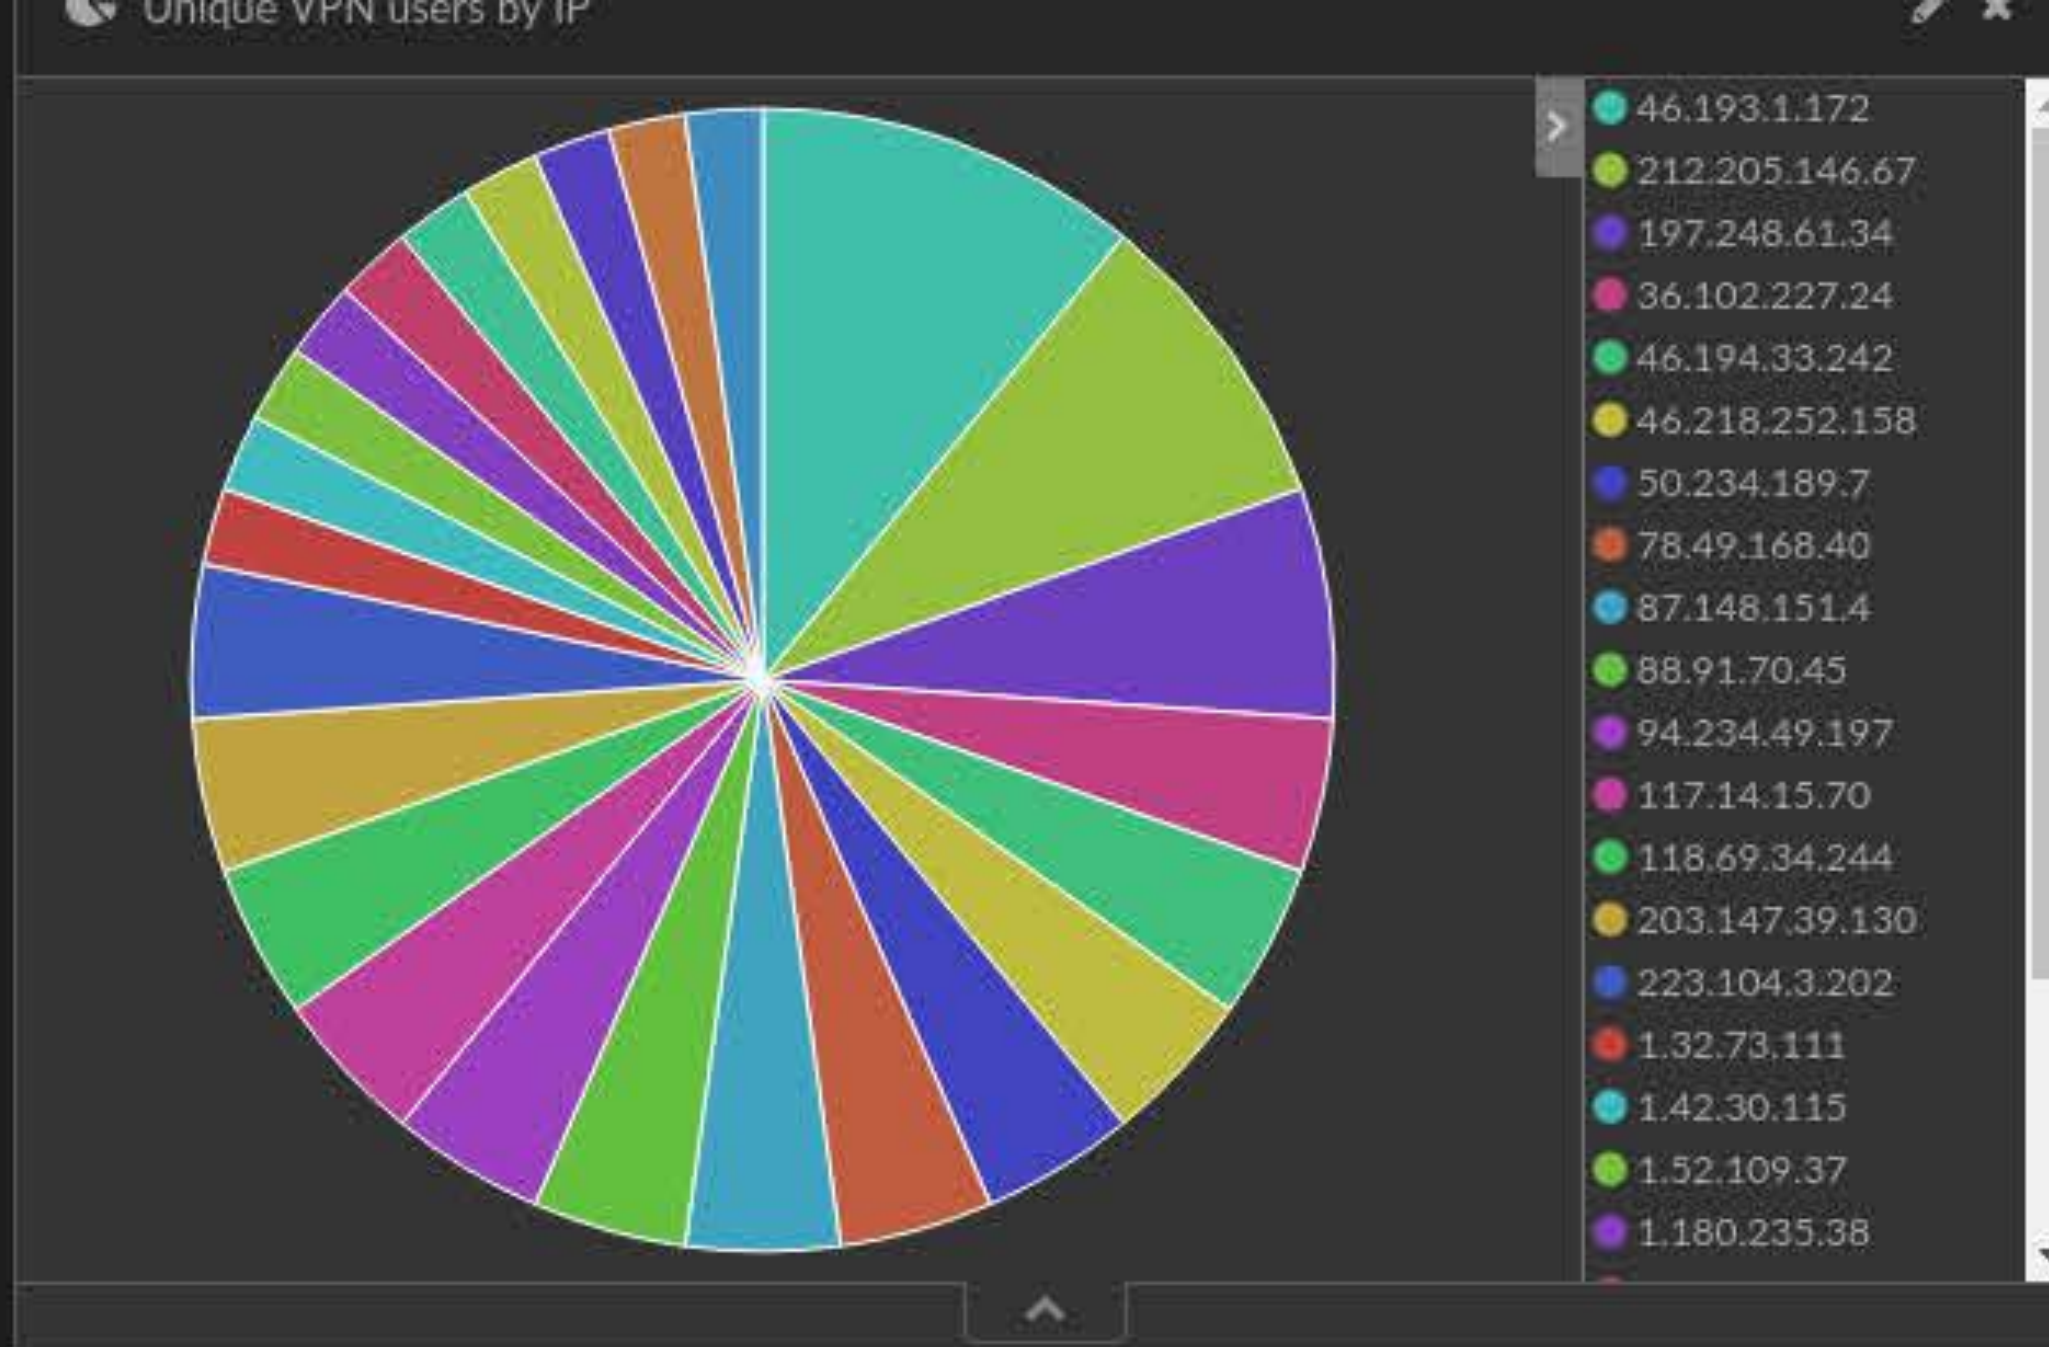





# EXAMPLE: VPN ABUSE

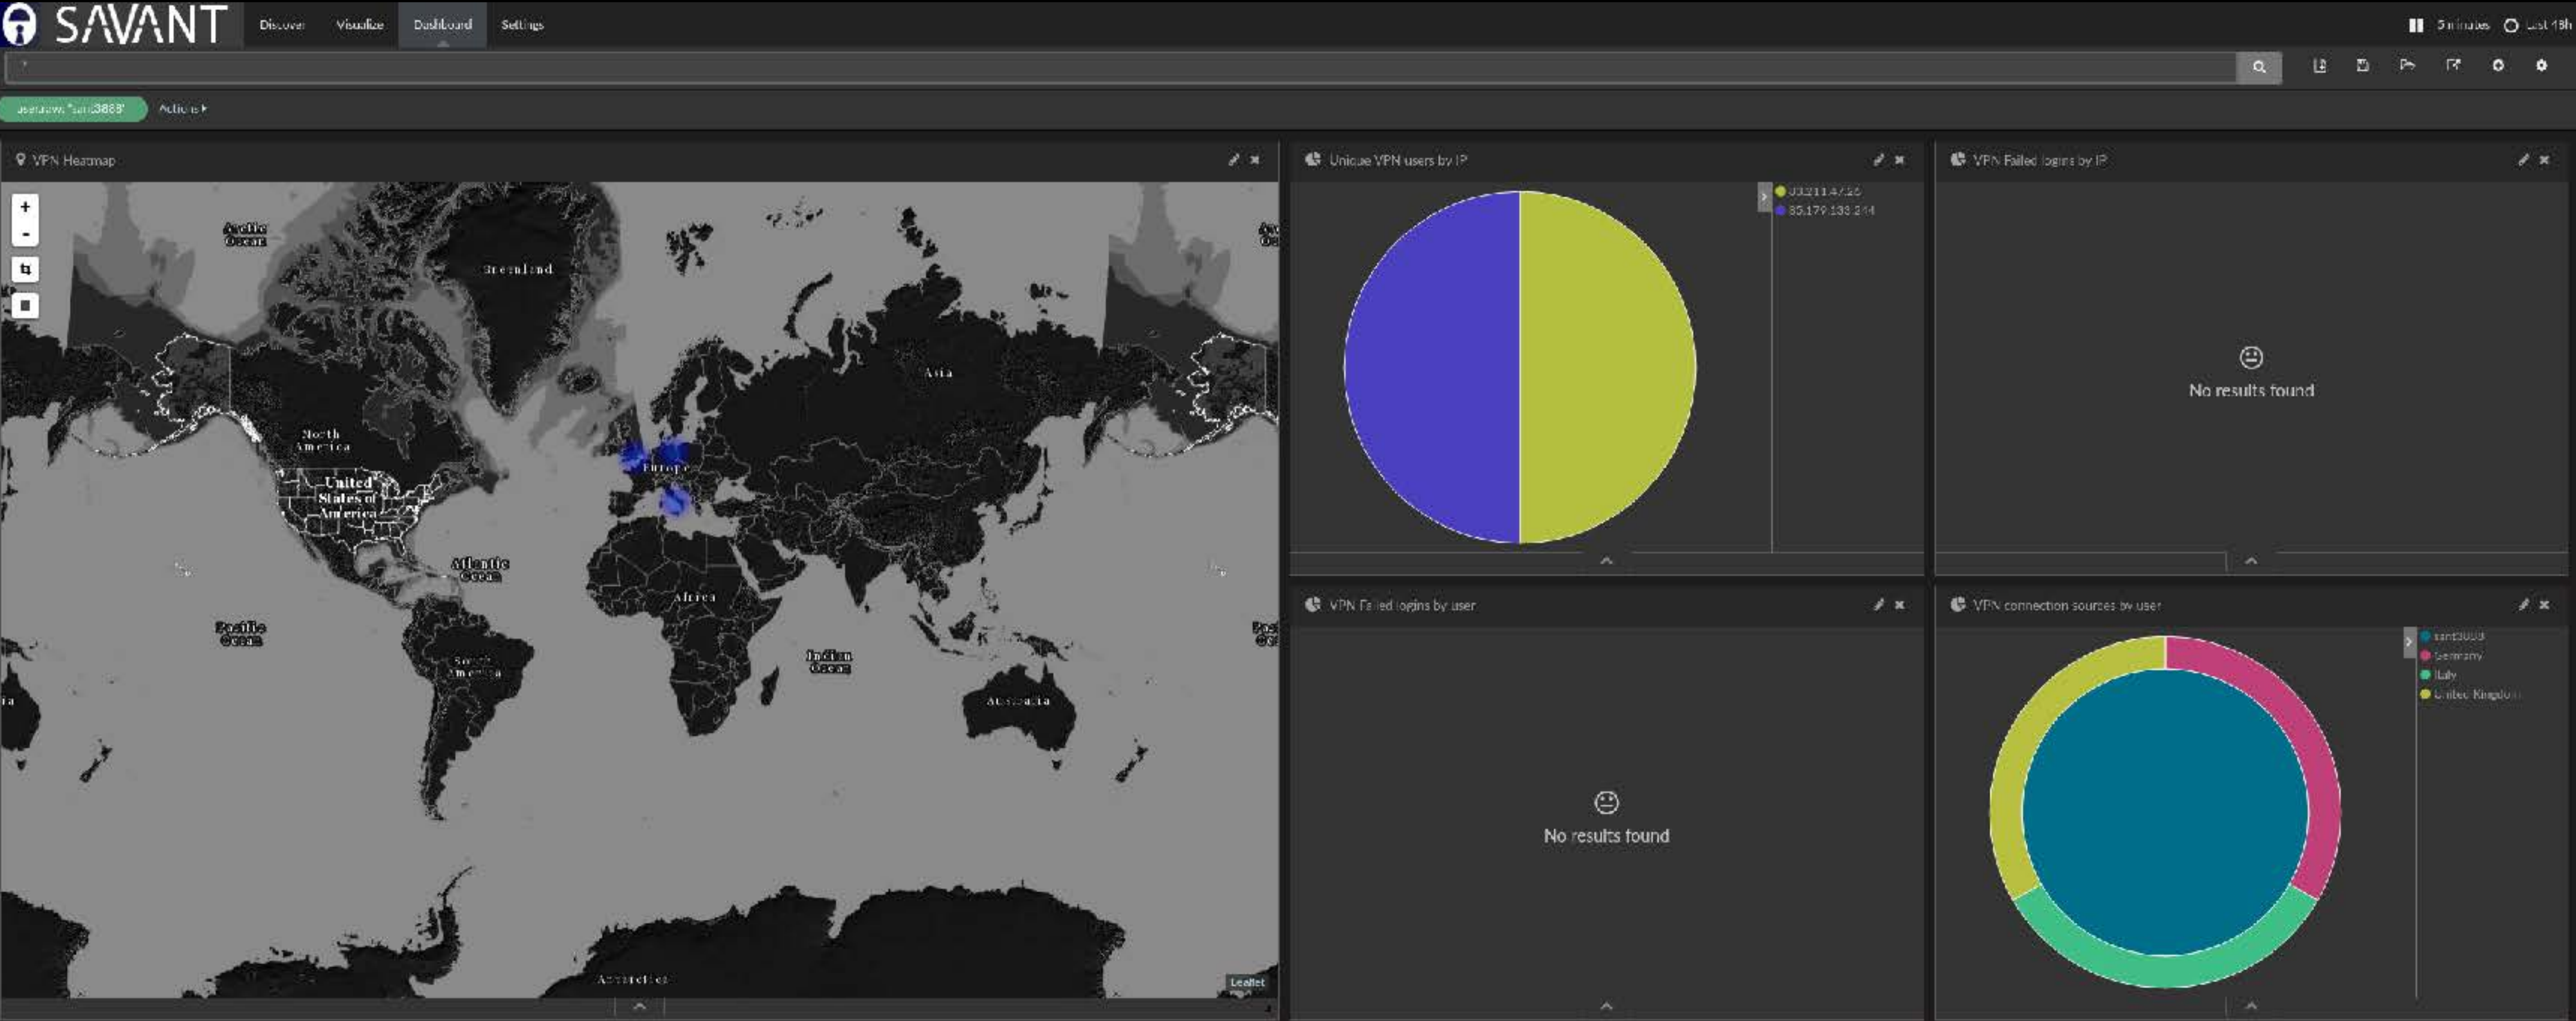








# EXAMPLE: VPN ABUSE (CONTINUED)





# STRATEGIES FOR PHISHING?



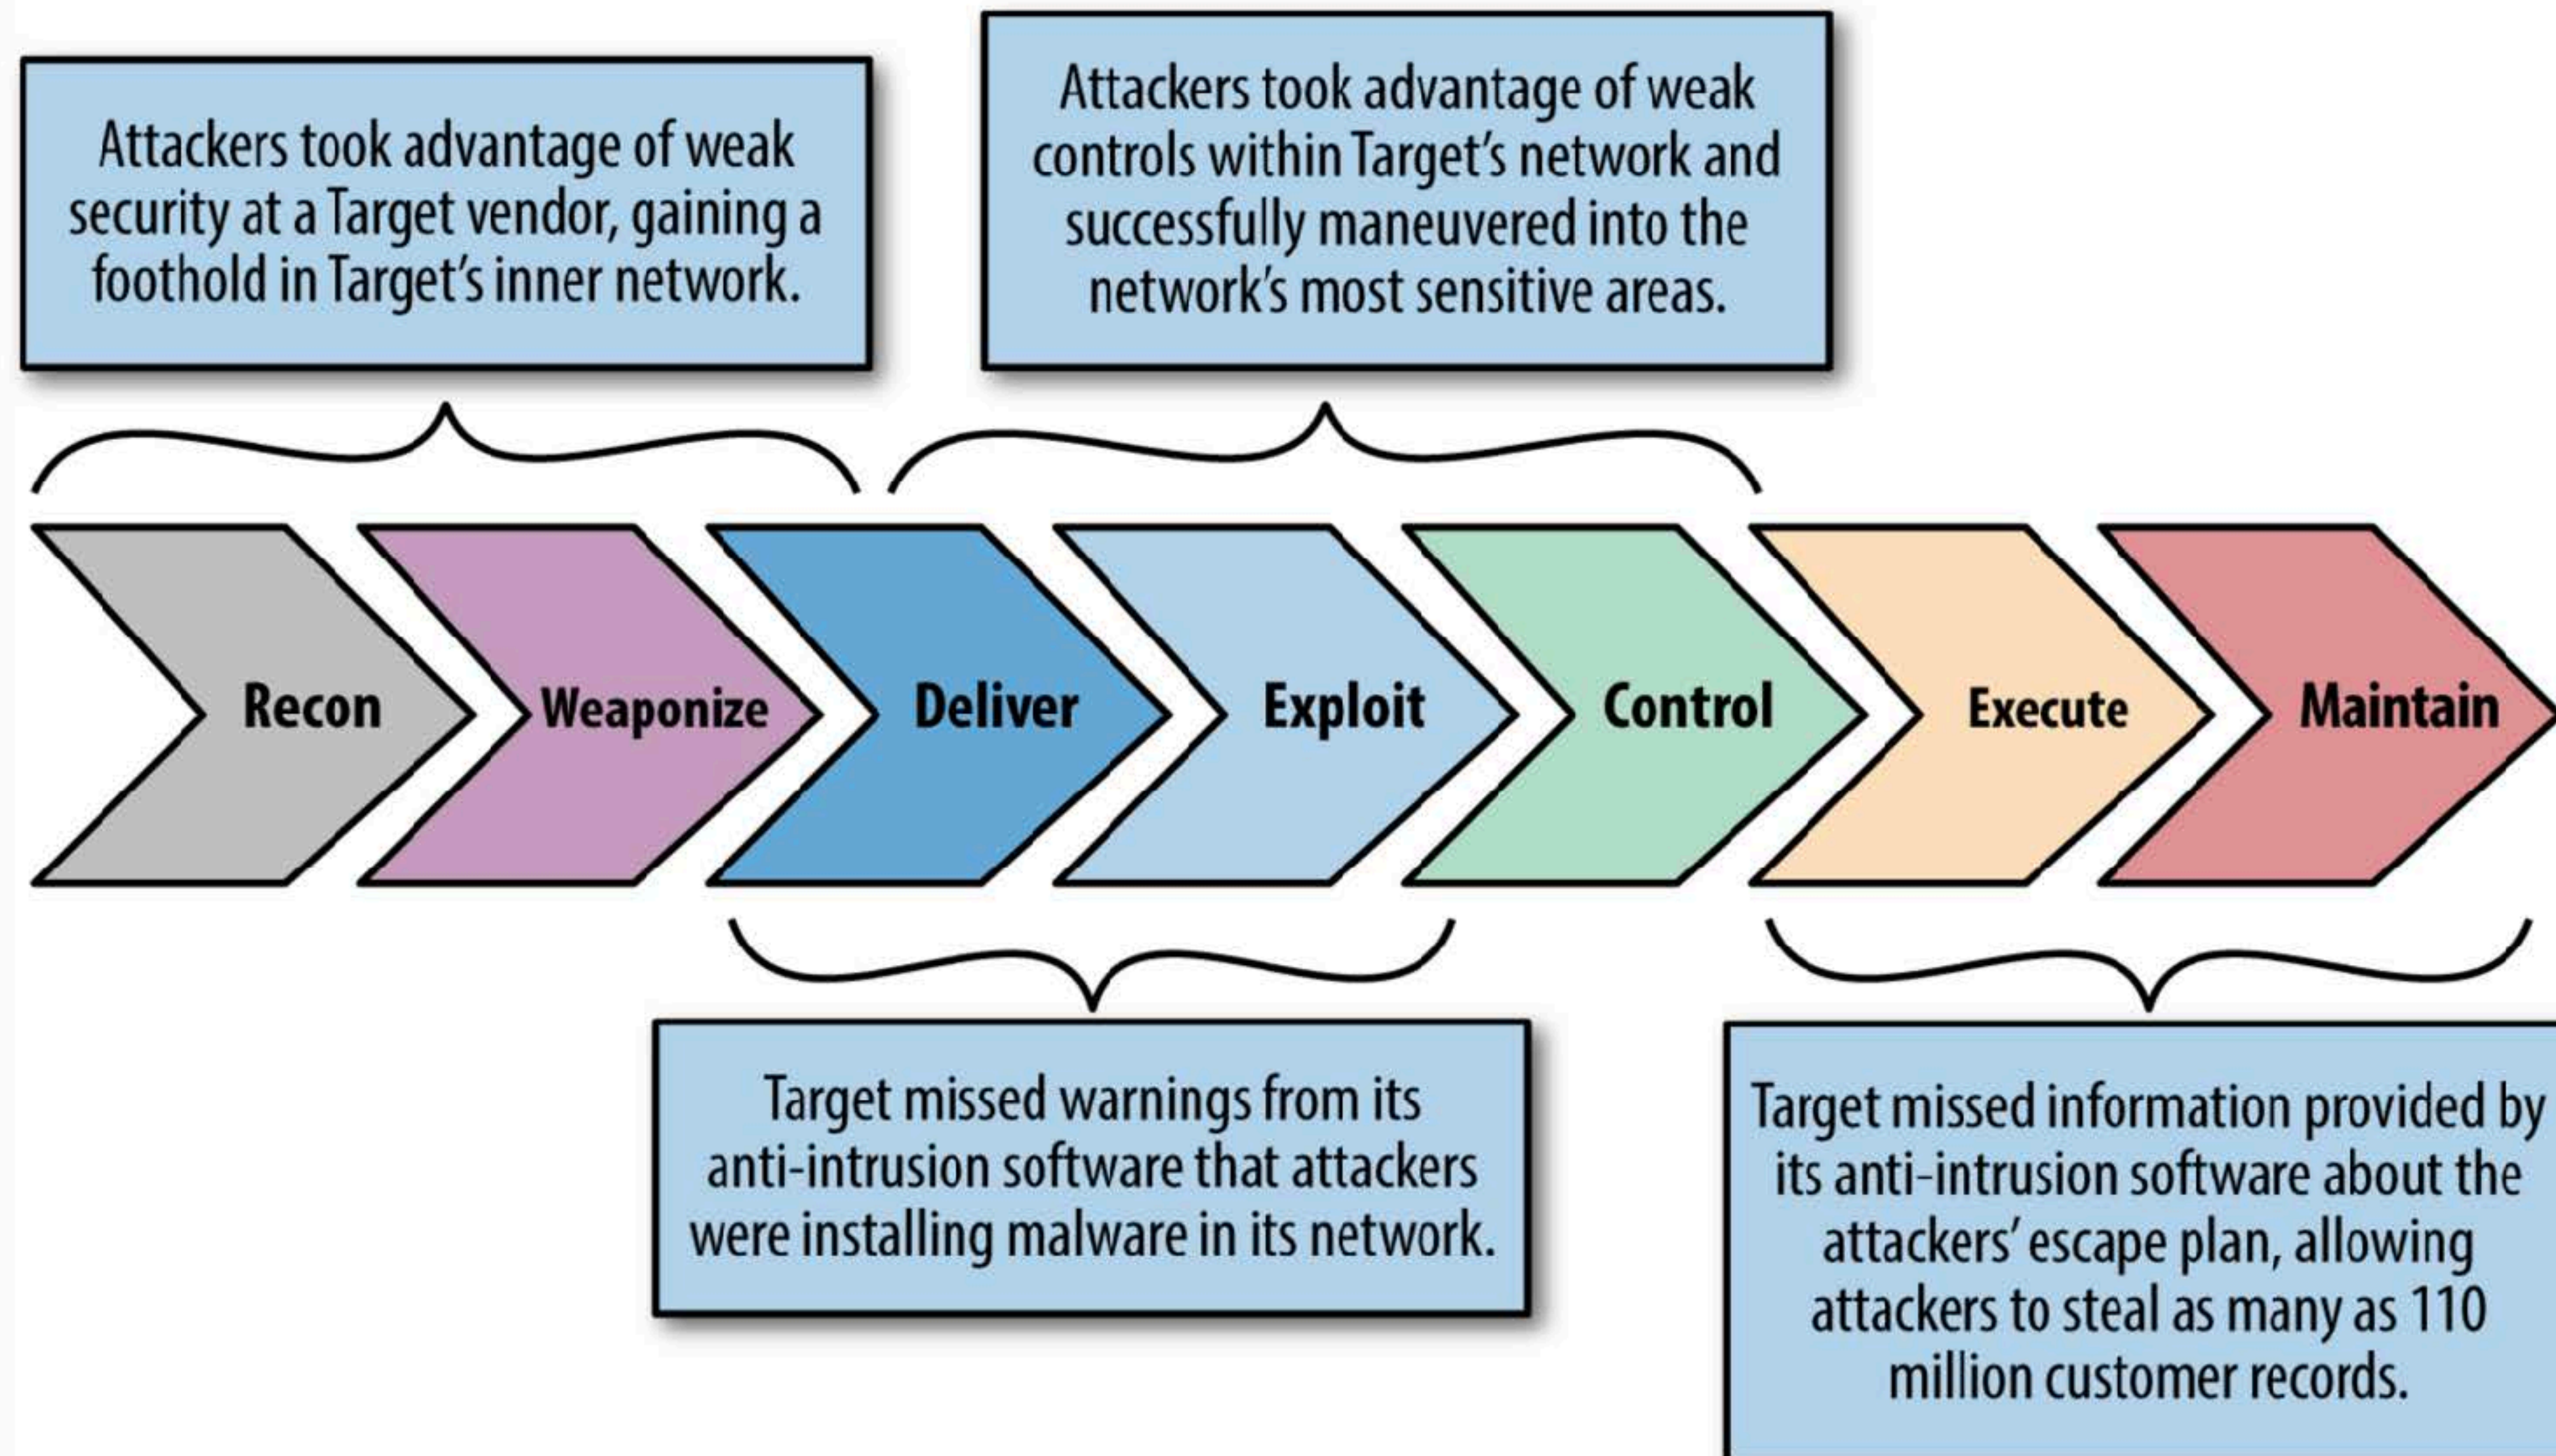
User education on phishing:  
continuous awareness campaigns,  
training, internal phishing exercises.

Response Policy Zone (RPZ) to  
subvert phishing sites to a sinkhole.

Email security products, e.g.  
Mimecast, Advanced Email Threat  
Protection, Hosted Email Security, ...



# MALICIOUS CODE: THE KILL CHAIN





# MEASURE UP



There are several ways to measure a team's detection efficiency with a few simple metrics such as the following:

- How long it takes to detect an incident after it occurred?
- How long it takes to contain an incident after its detection?
- How long it takes to analyse an alert or solve an incident?
- How many infections are blocked or avoided?
- How well are playbook reports performing?



# FURTHER READING



## Jeff Bollinger, Brandon Enright, and Matthew Valites: Crafting the InfoSec Playbook

O'Reilly; 1st edition (6 May 2015)

<http://oreilly.com/catalog/errata.csp?isbn=9781491949405>

## Aaron Bradley: OS X Incident Response

Syngress/Elsevier; 1st edition (6 May 2016)

<https://www.elsevier.com/books/os-x-incident-response/bradley/978-0-12-804456-8>





# THANK YOU!

 <https://github.com/mjung/publications>

**MARKO JUNG**

GALACTIC VICEROY OF RESEARCH EXCELLENCE

 [m@mju.ng](mailto:m@mju.ng)

 [@mjung](https://twitter.com/mjung)

 [fb.com/markohjung](https://fb.com/markohjung)