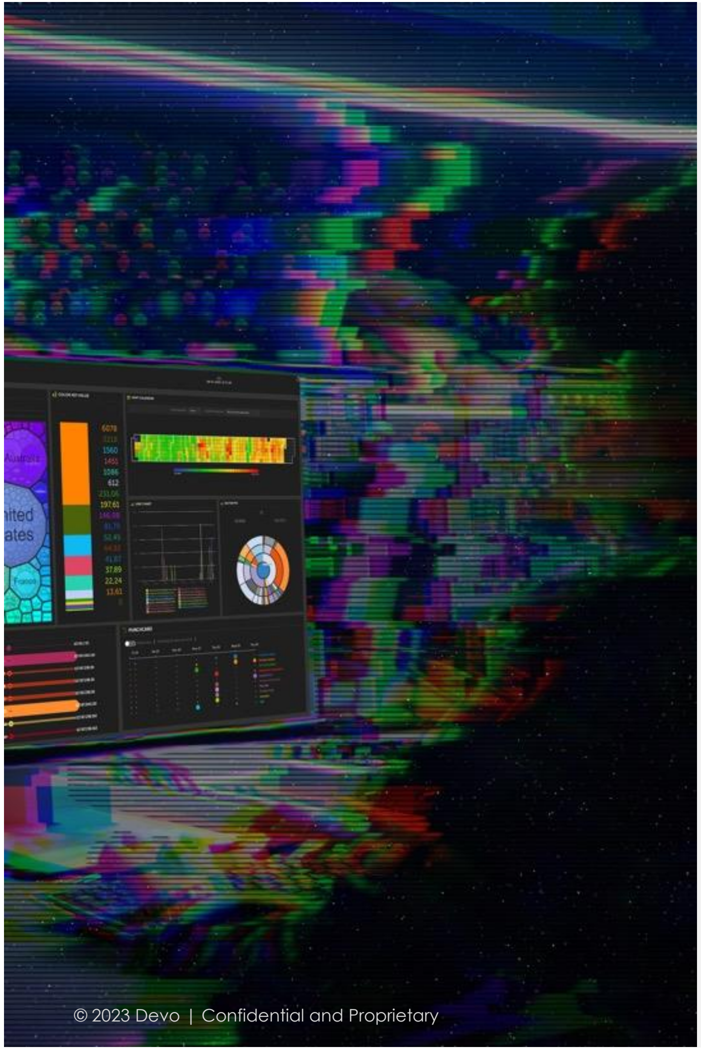

Quit Fussing over All Those Alerts:

Using Automation to identify Leads

February, 2023

Dan Pistelli

> whoami dan.pistelli

- Security Researcher at Devo
- Former offensive security professional turned Blue
- Currently focusing on finding threats in large data sets

Alert Fatigue Demands Immediate Attention

Common Scenario

SOCs deal with a sea of data and 1000s of alerts per day

- The average SOC sees over 10,000 alerts per day
- Even if each alert only takes 5 minutes to work, it would take over 800 FTE to triage all the alerts each day

Common Symptoms

Across the industry, all SOCs face alert fatigue driven by

- Lack of alert aggregation/poor correlation
- Inconsistent alert review and escalations
- Overwhelming volume/backlog of alerts

Common Results

Struggling to keep up with the backlog, analyst will

- Ignore or sometimes even turn off alerts
- Not investigate deep enough or make incorrect assumptions to get through more alerts
- Miss incidents, which in turn will amplify the situation

Out with the Old: A Better Approach

- Content funnelling is the process reducing data from left to right as data works its way down the funnel
- Data can be reduced in various ways, such as pattern matching, statistical models, and aggregation
- In a SOC, for example, if you start with 1 MM events per day, ideally you end up with 1000's of alerts but only 10's of cases

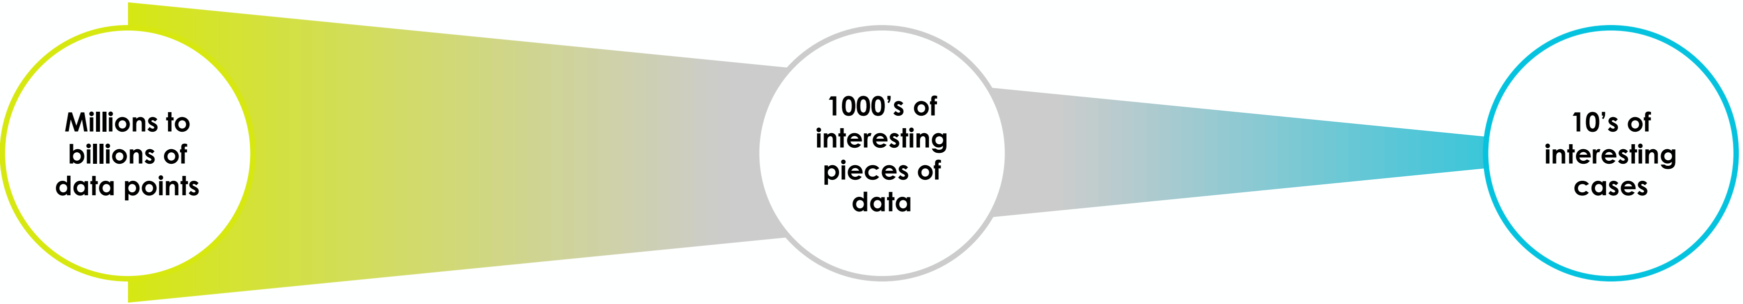
Millions to billions of data points
1000's of interesting pieces of data
10's of interesting cases

How Do We Implement a Successful Content Funnel

Such an approach should include **automated**:

- **Aggregation**
- **Correlation**
- **Data Science**
- **Remediation**

Let’s Take a Deep Dive into an Alert

Account Name	Alert Name	IP Address	Location
Charlie	Brute Force Attempt	23.88.190.126	Kansas

Aggregation and Correlation

A common step in alert triage is to answer the following questions

- Did this entity do anything else within a given time frame?
- Are there any other entities that may have been involved?
- Can I build a story with the data I have collected?

The answers to these questions should provided to the analyst before they start

- We can aggregate the alerts by entities over different time periods so that an analysts does not need to perform these extra steps
- Using graph theory, we can build relationships between the entities such that related entities can also be group together within a time frame

The Example: Visualizing a Single Alert

Single alerts don’t have sufficient context

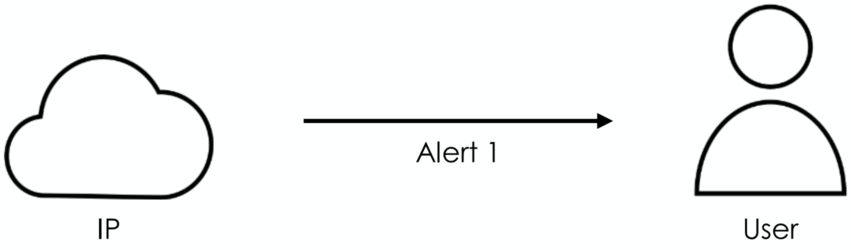

Account Name	Alert Name	IP Address	Location
Charlie	Brute Force Attempt	23.88.190.126	Kansas

The Example

We can give the analysts a head start using graph theory

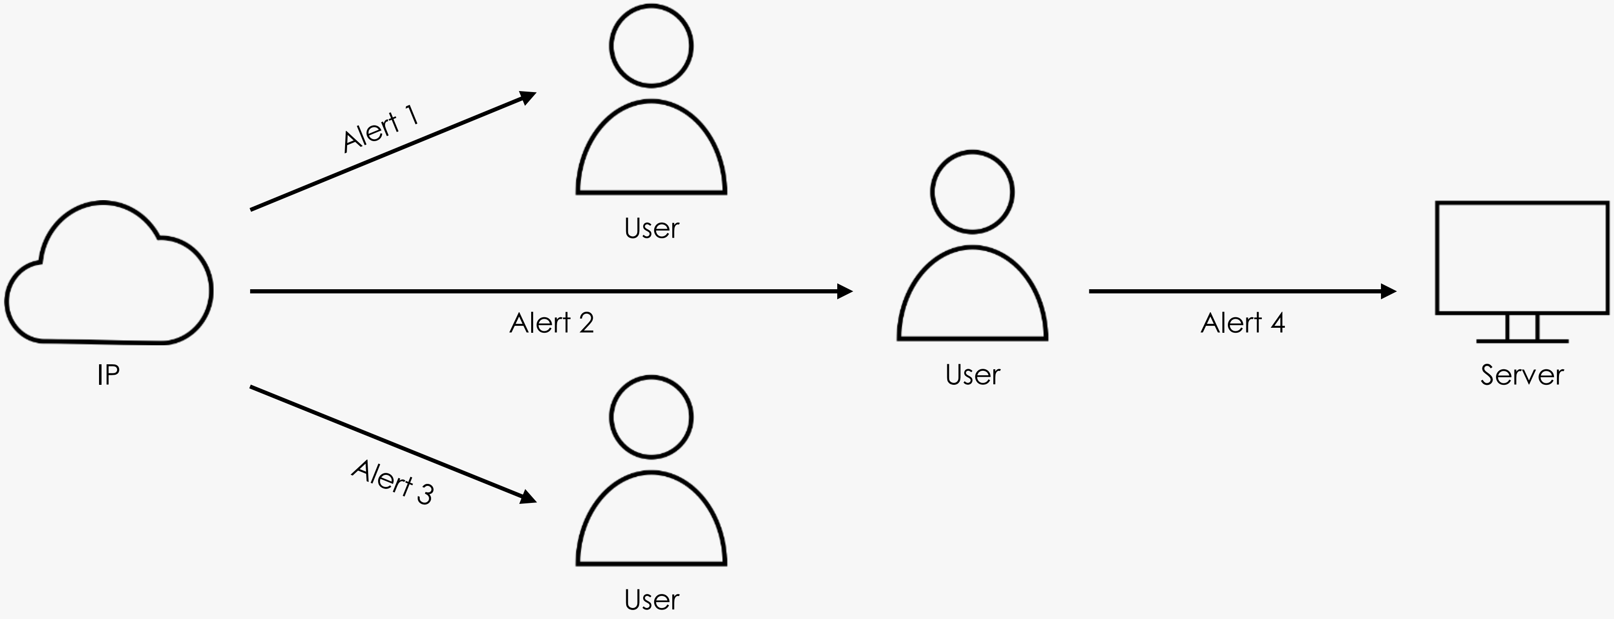

Using Data Science to Improve Alerting

Situation 1:

Charlie C-level logs onto system while on his first business trip to Kansas

Situation 2:

Attacker compromises and authenticates to Charlie C-level's account from Kansas

Both situations may provide the exact same data to a analyst

Account Name	Event	Location
Charlie	Successful Login	Kansas

Adding More Context with Data Science

- Working from our previous example, we are still missing key answers
  - Is it normal for Charlie to log in from Kansas?
  - Has he ever logged into from Kansas
- Simple math can be used to have these questions answered before the analysts begins an investigation
- Preprocess statistics can be enhance our data along the entire funnel

Using Data Science to Improve Alerting

Situation 1:

Charlie C-level logs onto system while working in Kansas

Situation 2:

Attacker compromises and authenticates to Victor Vip's account from Denver

Account Name	Event	Location	Count	Consecutive_days
Victor	Successful Login	Denver	1	1
Charlie	Successful Login	Kansas	300	27

Automated Remediation

Situation 2:

Attacker compromises and authenticates to Victor Vip's account from Denver

Remediation:

1. Notify user
2. Reset user's password
3. Send Jira ticket to IR team to follow up
4. Write incident report

Account Name	Event	Location	Count	Consecutive_days
Victor	**Successful Login**	**Denver**	**1**	**1**
Charlie	Successful Login	Kansas	300	27

Improving Mean Time Metrics with Automation

- **Notify user**
  - Analyst do not need to spend time re-writing this message
- **Reset user’s password**
  - Analyst do not need to spend time switching screens or remembering this process*
- **Send Jira ticket to IR team to follow up**
  - Automation: Analyst do not need to spend time re-writing this ticket
- **Write incident report**
  - Automation: Analyst do not need to spend time re-writing this report

* With more intrusive remediations, implementing human-in-the-loop style automations can help reduce miss-fire

Putting it all together…

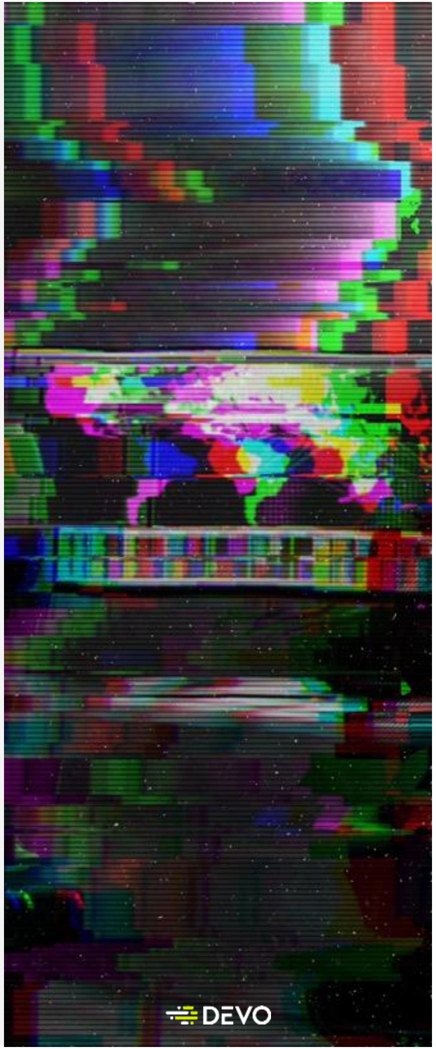

Correlated ATT&CK Techniques

1. Attacker authenticates to Charlie C-level account from Tokyo

2. Charlie logs into his account from the New York HQ 3 minutes later

3. Charlie’s machine begins to make short lived connections out to a Tokyo IP every 15 minutes

4. Google Drive logs an IP with 100’s of failed login attempts, except one successful one for Charlie

5. Google Drive logs charlie downloading 10x documents than he typically would from Google Drive

Correlated ATT&CK Techniques

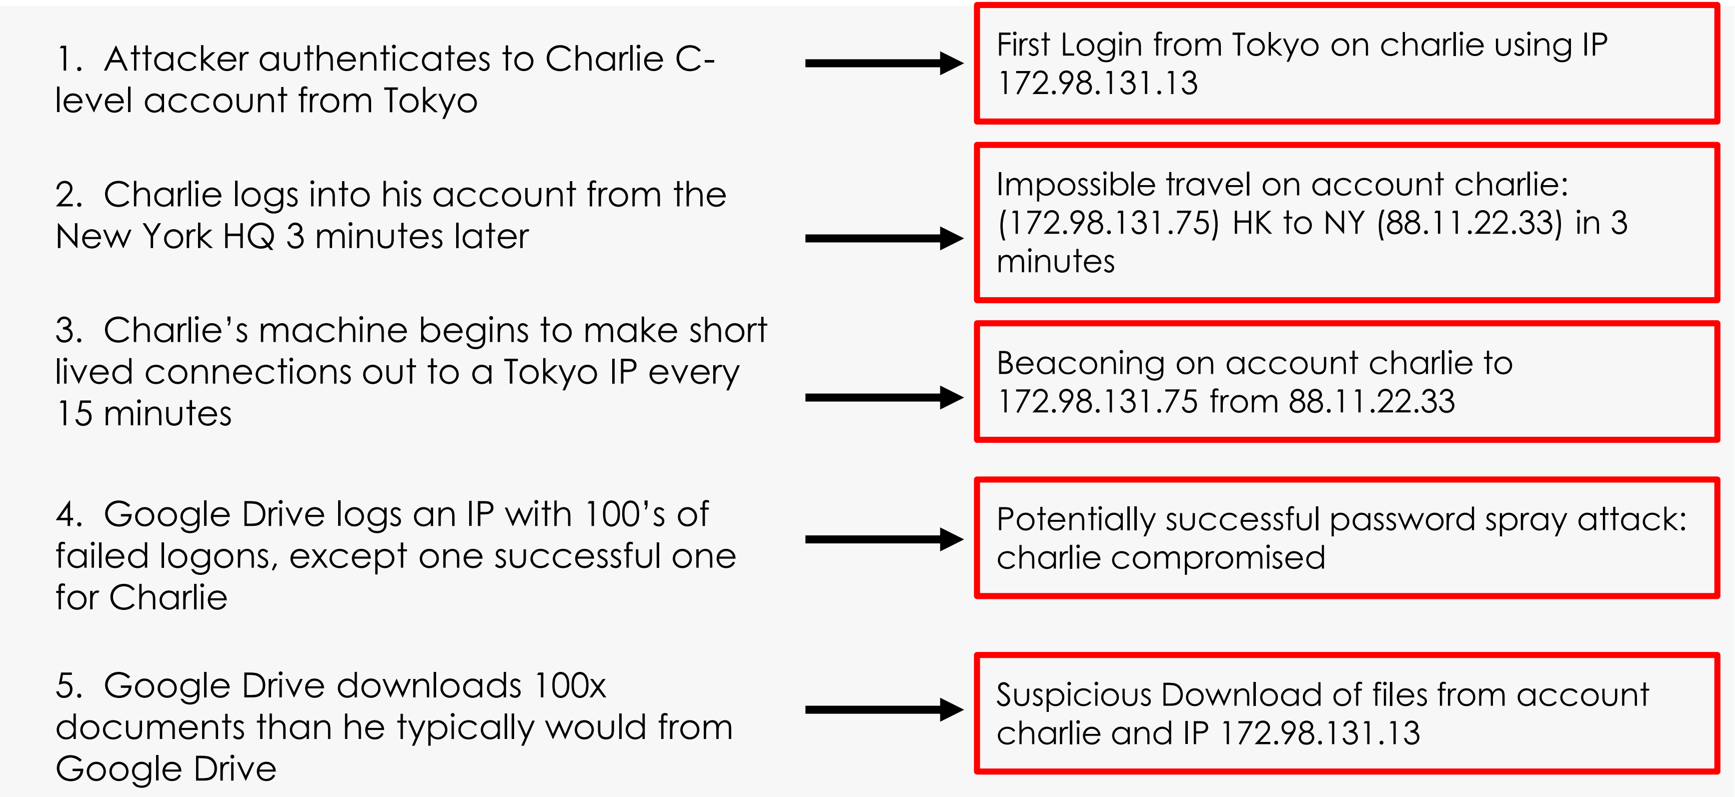

1. Attacker authenticates to Charlie C-level account from Tokyo

First Login from Tokyo on charlie using IP 172.98.131.13

2. Charlie logs into his account from the New York HQ 3 minutes later

Impossible travel on account charlie: (172.98.131.75) HK to NY (88.11.22.33) in 3 minutes

3. Charlie's machine begins to make short lived connections out to a Tokyo IP every 15 minutes

Beaconing on account charlie to 172.98.131.75 from 88.11.22.33

4. Google Drive logs an IP with 100's of failed logons, except one successful one for Charlie

Potentially successful password spray attack: charlie compromised

5. Google Drive downloads 100x documents than he typically would from Google Drive

Suspicious Download of files from account charlie and IP 172.98.131.13

Correlated ATT&CK Techniques

1. Attacker authenticates to Charlie C-level account from Tokyo → First Login from Tokyo on ccharlie using IP 172.98.131.13

2. Charlie logs into his account from the New York HQ 3 minutes later → Impossible travel on account ccharlie: (172.98.131.75) HK to NY (88.11.22.33) in 3 minutes

3. Charlie's machine begins to make short lived connections out to a Tokyo IP every 15 minutes → Beaconing on account ccharlie to 172.98.131.75 from 88.11.22.33

4. Google Drive logs an IP with 100's of failed logons, except one successful one for Charlie → Potentially successful password spray attack: ccharlie compromised

5. Google Drive downloads 100x documents than he typically would from Google Drive → Suspicious Download of files from account ccharlie and IP 172.98.131.13

Correlated ATT&CK Techniques

1. Attacker authenticates to Charlie C-level account from Tokyo

2. Charlie logs into his account from the New York HQ 3 minutes later

3. Charlie's machine begins to make short lived connections out to a Tokyo IP every 15 minutes

4. Google Drive logs an IP with 100's of failed logons, except one successful one for Charlie

5. Google Drive downloads 100x documents than he typically would from Google Drive

Confirmed Case	
Target	User: ccharlie
Correlated Indicators	1) Potentially successful pwd spray involving account: ccharlie 2) Unusual login by account: ccharlie 3) Impossible travel from account: ccharlie 4) Suspicious Files Downloads 5) Beaconing on account: ccharlie
Attacker	Entity out of Tokyo from IPs: ● 172.98.131.75 ● 172.98.131.13

More data. More clarity. More confidence.