Hurricane Labs

Supercharged SIEM:

Empowering your logs with the right data

Tom Kopchak

Logs, Logs, Logs

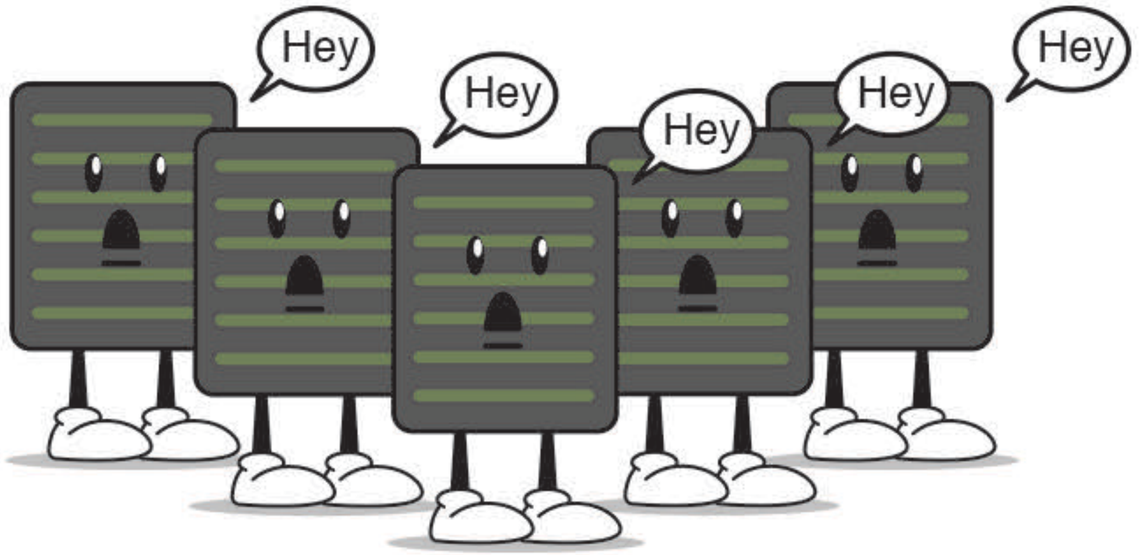

We can get logs from everywhere.

But how do we decide what logs we really need?

This is especially useful when we're paying for our SIEM by log volume.

Some Logs Are Better Than Others

- While almost every log can be used for something, we want to prioritize those that are important from a security perspective
- It's possible that some logs may have little to no security value; however, they're likely useful for another group

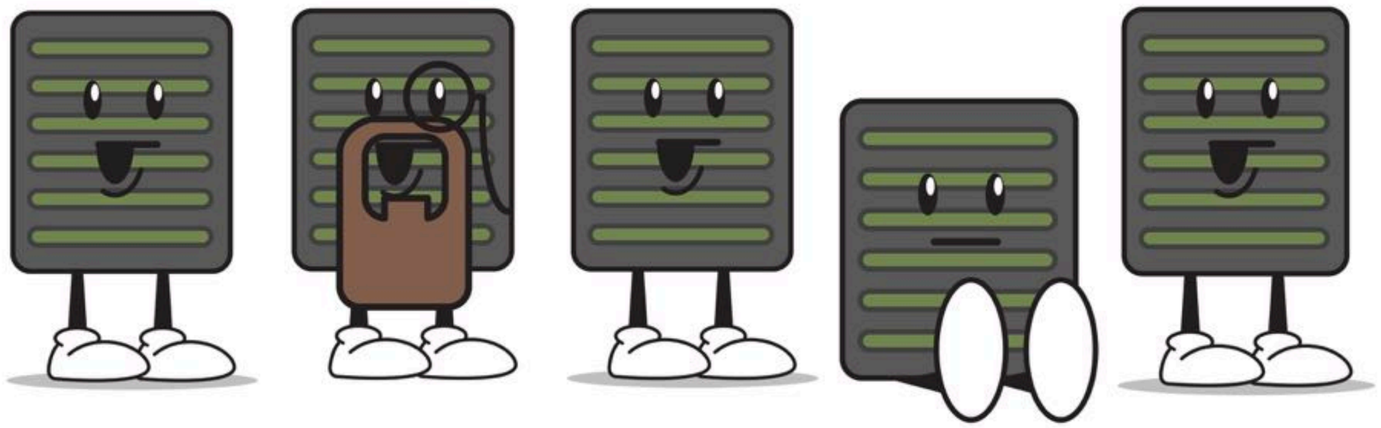

Hurricane Labs

Security Use Cases

- In the words of a Project Manager: “What problem are we trying to solve here?”
- When implementing a SIEM, start by identifying the most business-critical threat types (e.g. industry, assets, attack vectors, etc.)
- Address the largest issues first, then continuously evaluate, add, and improve

Designing Security Use Cases

Identify the type of threat you are concerned about.

Think like an attacker:

- What vectors are most likely to be used for an initial attack?
- How might an attacker pivot through your environment?
- What controls do you have in place to mitigate such an attack?

So, What Data Do I Really Need?

- Hurricane Labs has identified eight types of data that are most important to a successful SIEM implementation

- Some of these are more immediately useful than others, but all eight work together to form a comprehensive SIEM

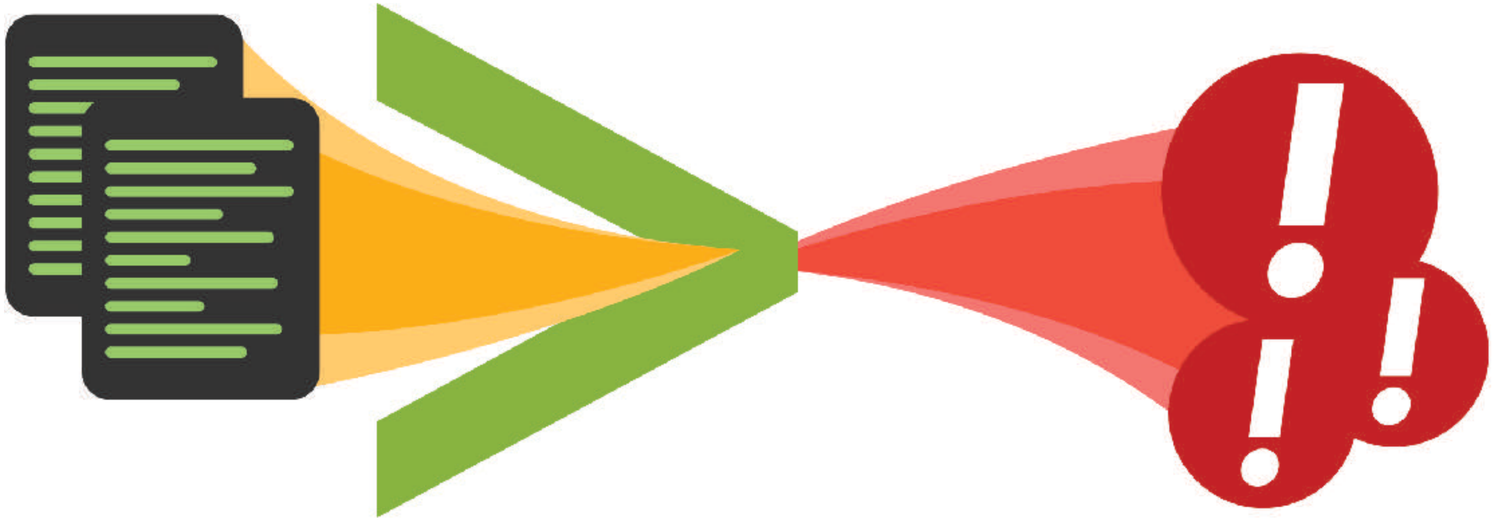

Hurricane Labs

How To Use This List

- Two birds, one stone: The data sources are broken into different categories and you may have sources that fall into multiple (e.g. IDS events from your super-duper magical next-gen firewall)
- One bird, multiple stones? Likewise, you may have various sources that contribute to a data type (e.g. multiple firewall or IDS vendors)
- We’re not mind readers: Knowing what technologies exist in your environment is critical

Hurricane Labs

Firewall

- Let’s start with an easy one
- Firewalls show everything that enters and leaves the network - in theory
- Both traffic allowed and not allowed can be useful from a SIEM perspective
- Challenge: Not everything you want to see may be logged by default

DNS

- This is easily the most useful, but often the most overlooked
- No single source will provide as much security-actionable data
- With SSL, DNS remains an easy way to view system communication patterns in a human-readable way

Challenges: Identifying the true source of the request

- E.g. Knowing your DNS server made a DNS request isn't that useful; however, determining what endpoint made the DNS request is
- Volume of data is huuugggeee, as is almost every communication that requires a DNS request

Hurricane Labs

IDS/IPS

- IDS can identify threat types based on network activity
- May detect malware such as ransomware where exploit-kits are used

Challenges:

- Far from perfect when dealing with encrypted traffic
- Signature-based detection is inherently limited

Hurricane Labs

Authentication

Provides insight into account usage and user behavior:

- What is being accessed and modified in your environment?
- What users are connecting successfully, and from what locations?
- Are there any anomalous login patterns, or a significant increase in activity?

Hurricane Labs

Anti-Virus

- Another source that’s far from perfect, yet still useful
- Can identify sources of malware infections, as well as threats that are flagged but unable to be cleaned

Hurricane Labs

Proxy

- Provides greater insight into web activity than firewall logs alone (remember, most malware activity is equivalent to web browsing from an application level)
- Can identify an initial compromise
- If the proxy decrypts SSL, it provides significantly more insight into communications and payload (but this also complicates matters and creates privacy concerns)

Hurricane Labs

E-mail

- Header information provides details on message origin and recipients without having to ingest message content
- Can identify potential phishing/spam e-mails based on this information alone, and proactively block attack attempts
- Can also be used for forensics purposes after a security incident

Hurricane Labs

Vulnerability Scanner

- Can be used to identify hosts running outdated or vulnerable software
- Can provide insight into the likelihood of a successful attack on a given host
- Can be used to increase or reduce the severity of alerts

Hurricane Labs

Runners Up

Everyone gets a trophy!

Some of the useful sources didn't quite make it into The Big 8, including:

- SSL Certificate Data
- Asset and Identity (server and user) Info:
  - Who is who, what is what, and what does it do
  - Actually, this should be item “0” on the list
- Audit Logs (for sensitive data)

We Have These Logs, Now What?

- Use your SIEM for correlation: Let it find events that will require human interaction and verification.
- Meaningful tuning: Alerts should be timely, actionable, and relevant. Report on events that do not require immediate action, suppress events that are not applicable to your environment.
- Your MSSP should be measured in terms of "noise" decrease on your end and how many alerts they viewed (so you don't have to).

Let’s see how these logs are useful in a potential attack scenario.

Scenario 1: Ransomware

Hurricane Labs

The attack:

- A user opens a phishing email that leads to an exploit kit page containing ransomware
- The ransomware successfully encrypts the target machine and attempts to pivot laterally through the network

What is your SIEM doing?

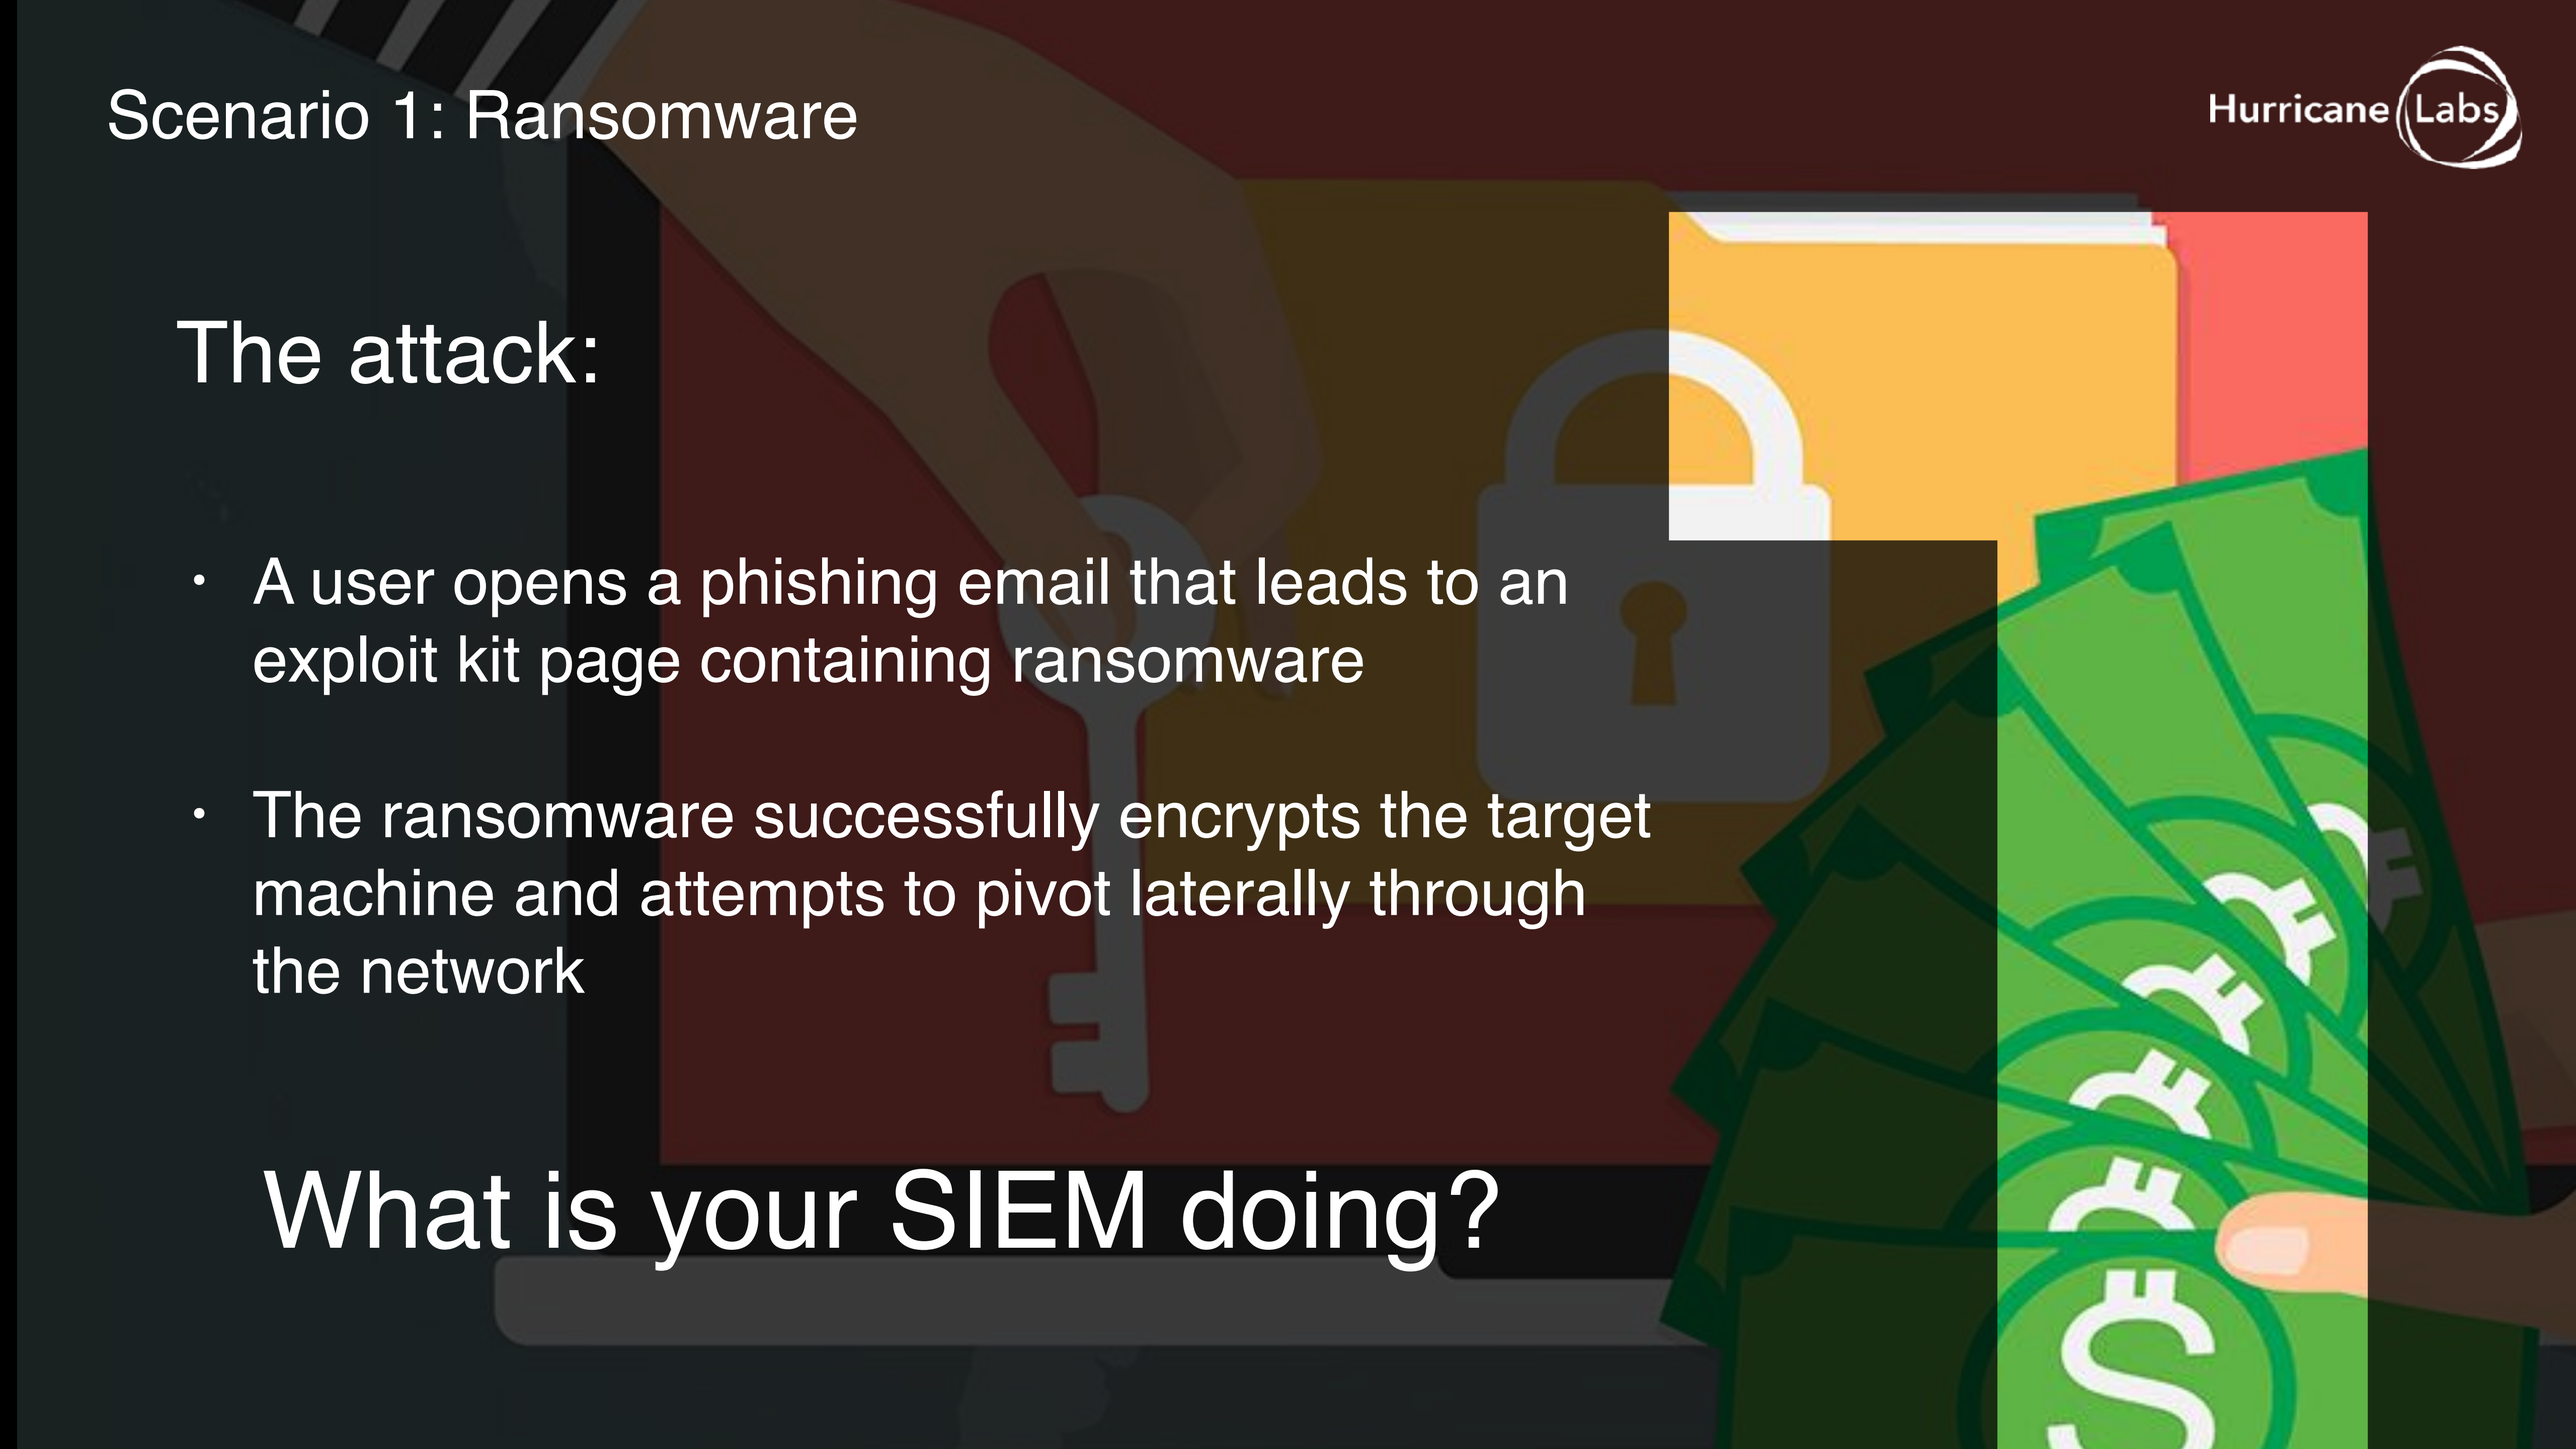

Scenario 1: Ransomware

Email:
- Email logs identify the initial phishing email, the target, and patient 0

DNS:
- Most malware relies on DNS requests to domain generation algorithm (DGA) domains
- Watch for requests for new, uncommon, or seemingly random domains
- May also see DNS requests for tor2webgateway
- Watch for requests for suspicious TLDs like .top and .xyz

Antivirus:
- May detect presence of ransomware installation, or ransom note (even if it doesn't stop ransomware from running)

Proxy:
- May detect outbound traffic, but less likely with SSL/HTTPS
- Useful for identifying patient 0

Scenario 1: Ransomware

Firewall:

- Will show outbound traffic from compromised machine, may be connecting to known malicious IPs
- Useful for identifying patient 0

Authentication:

- Activity associated with accessing/modifying a large number of files
(or attempting to modify files that a user doesn't have access to) - audit logs are even more helpful here

IDS:

- May detect ransomware activity/spread or initial compromise
- If ransomware is spread from exploit kit, IDS is more likely to pick that up

Vulnerability Scan:

- Can help in the case of exploit kits, identify hosts running outdated versions of software
susceptible of these attack vectors

Scenario 2: Unpatched Vuln Leads to Data Breach

Hurricane Labs

The attack:

- You're a major credit bureau
- You leave an unpatched server exposed to the Internet
- You lose data on 145 million Americans

What is your SIEM doing?

Scenario 2: Data Breach

Email:
- Potentially used for data exfiltration, you can watch for large emails to non-corporate domains or unusual email addresses

DNS:
- DNS tunneling can be used for exfiltration, look for excessively long DNS queries (especially TXT)

Antivirus:
- Less useful, though depending on the mechanism of breach, a RAT may be deployed

Proxy:
- Can be used to detect large outbound file transfers, access to pastebin sites

Scenario 2: Data Breach

Firewall:
- Vulnerability scanners will often scan consecutive IP addresses in a subnet
- Exfiltration can happen over unusual ports or to foreign countries

Authentication:
- Less relevant in this attack, why authenticate when you can exploit?
- Some software will log processes spawned via RCE, but most won’t

IDS:
- If your IDS is up-to-date, the exploited vulnerability probably has a signature

Vulnerability Scan:
- Can help identify hosts running software susceptible to these attack vectors

Scenario 3: Intellectual Property Monitoring

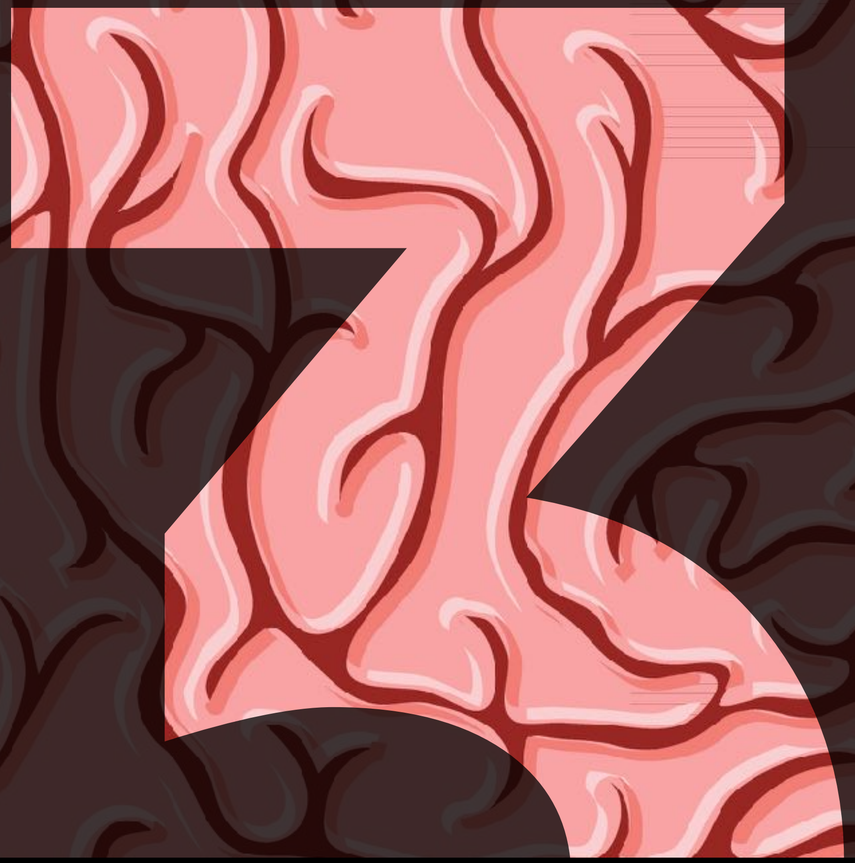

The attack:

- A user is maliciously attempting to access, exfiltrate, or modify intellectual property
- The user is an authorized employee, and may legitimately be allowed to interact with this data as part of their normal job function

What is your SIEM doing?

Scenario 3: Intellectual Property Monitoring

Email:

- Employee may use personal e-mail or the e-mail of an outsider for exchanging information

DNS:

- User may attempt to access unauthorized sites, or bypass controls using TOR or other tunneling

Antivirus:

- User may attempt to run unauthorized or otherwise malicious software, or tools to encrypt data to bypass DLP controls

Proxy:

- User may generate proxy access events or attempt to access resources restricted by the proxy
- There may be an increase in traffic to cloud storage providers

Scenario 3: Intellectual Property Monitoring

Firewall:
- Will show increased outbound traffic to external storage providers, as well as increased traffic to internal systems where IP is stored

Authentication:
- There may be excessive or abnormal authentication patterns, significantly higher than what is considered typical for one's normal job function

Vulnerability Scan:
- Can help detect if access was somehow accidentally left open

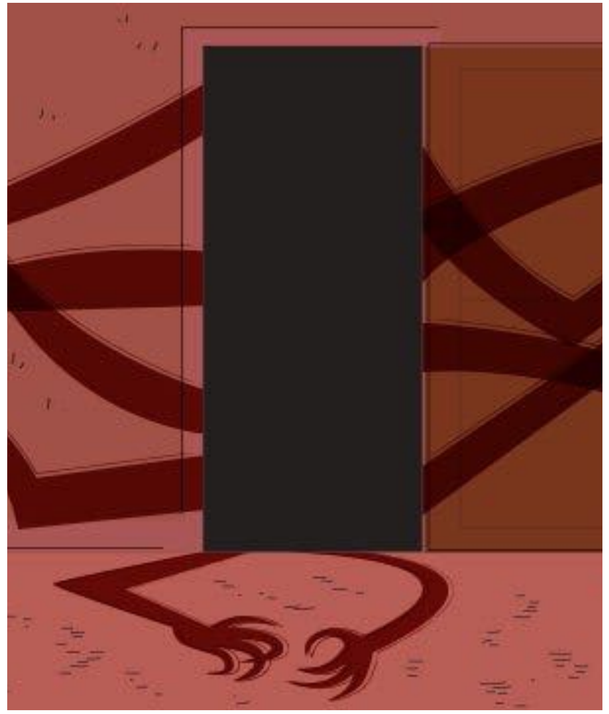

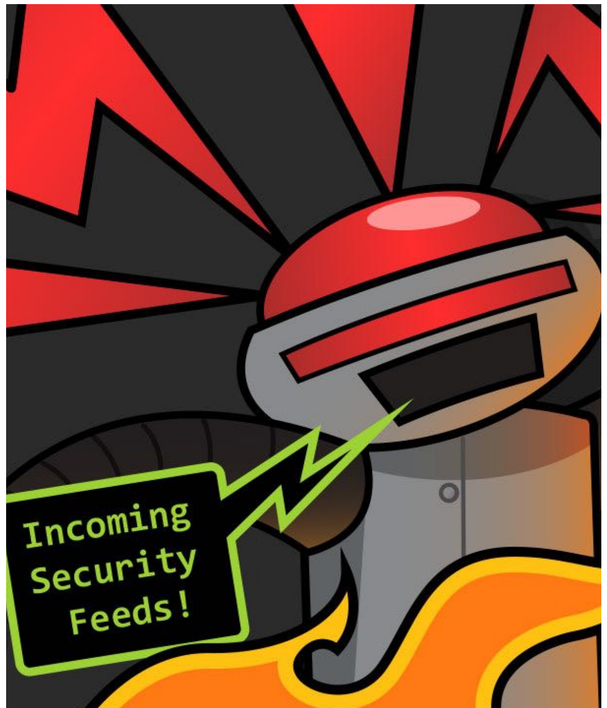

What Can We Learn?

- Defense in depth concepts continue to be applicable to your security monitoring solution
- Build layered defenses, and monitor them!
- No solution is perfect, but try to give yours the best chance of success
- Continuous improvement

Hurricane Labs

Any Questions?

Contact Us

Hurricane Labs

Keep in touch!

Tom Kopchak
Director of Technical Operations
Hurricane Labs (@tomkopchak)

tom@hurricanelabs.com


Hurricane Labs