ExpressVPN

Beyond Security Operations

Yordan Ganchev, Security Operations Center Manager

Intro

Security @ ExpressVPN

$whoami?

Yordan Ganchev

- GCFA, GNFA, GCIA
- SOC Manager @ExpressVPN
- Threat Intel, DFIR
- Hobbies: cats (obviously)

Who are we?

ExpressVPN Security

- HTB Business CTF 2021 - #11/374 🎖️
- Splunk BOTS ASEAN 2021- #1/34 🏆
- Meta CTF 2021- #10/1346 👑
- HTB Cyber Apocalypse 2022 - #31/7024 🥇
- HTB Business 2022 - #8/656 💎

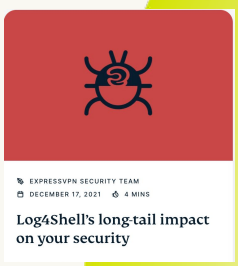

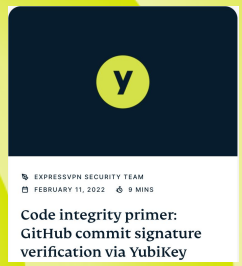

Intro

Agenda

- Traditional SOC operations vs Future of SOC
- “Security Operations as a Service”
- Intelligence in Security Operations
- Threat Modelling
- Threat Hunting
- Automation
- Case study using a world-renowned company

Story Time

Facing Uncertainty

Recognised vs Unrecognised Threats

Recognised vs Unrecognised Threats

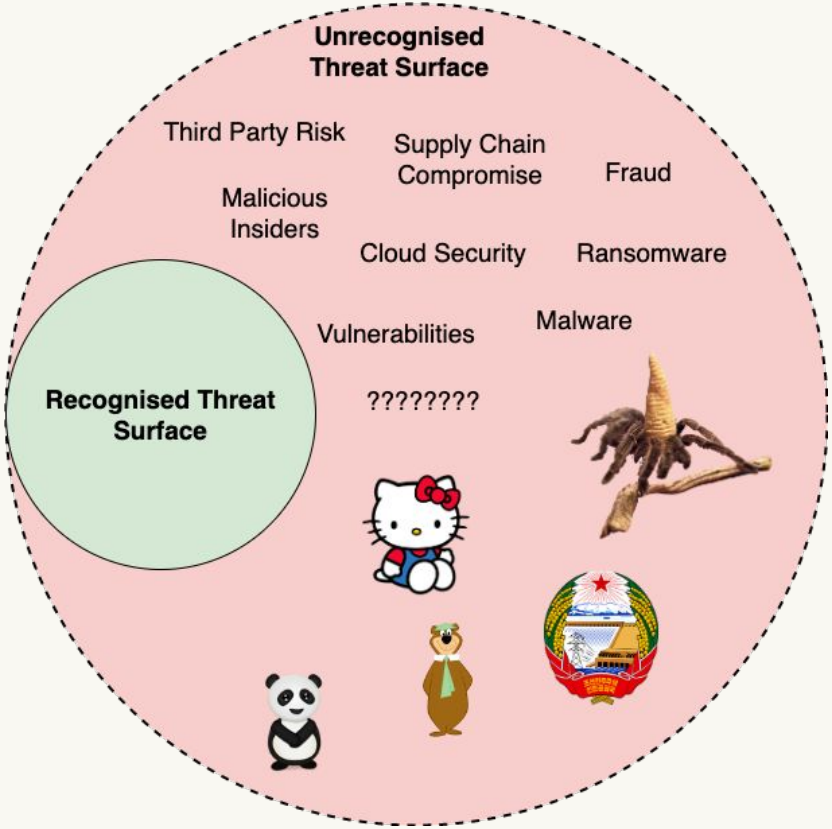

Recognised vs Unrecognised Threats

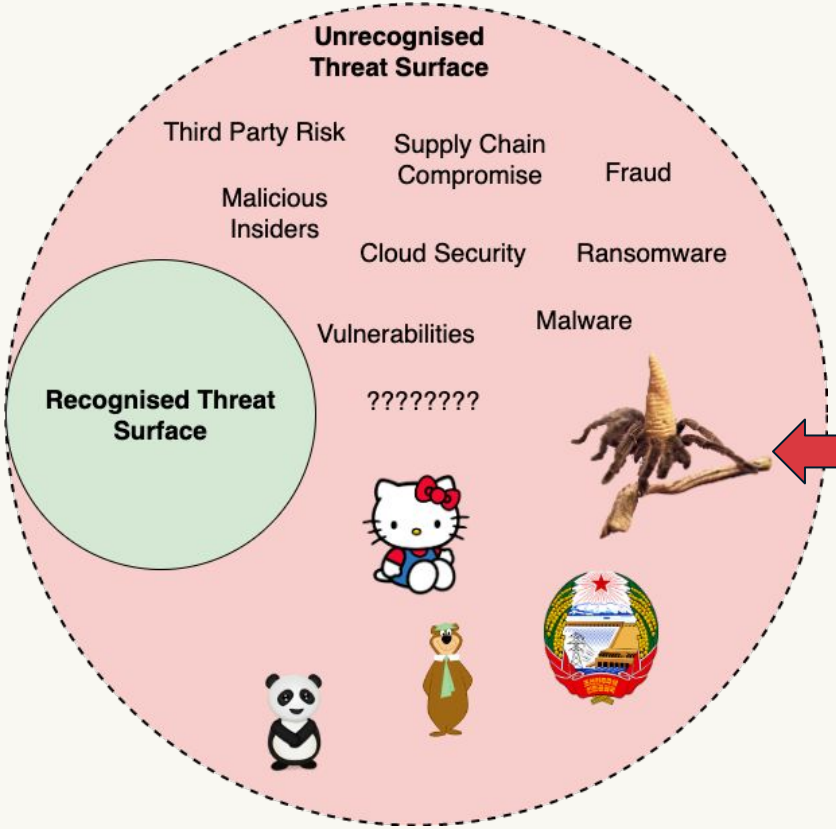

Confirmed Wizard Spider ring leader. Source - Interpol

Case Study - 4th-party email takeover

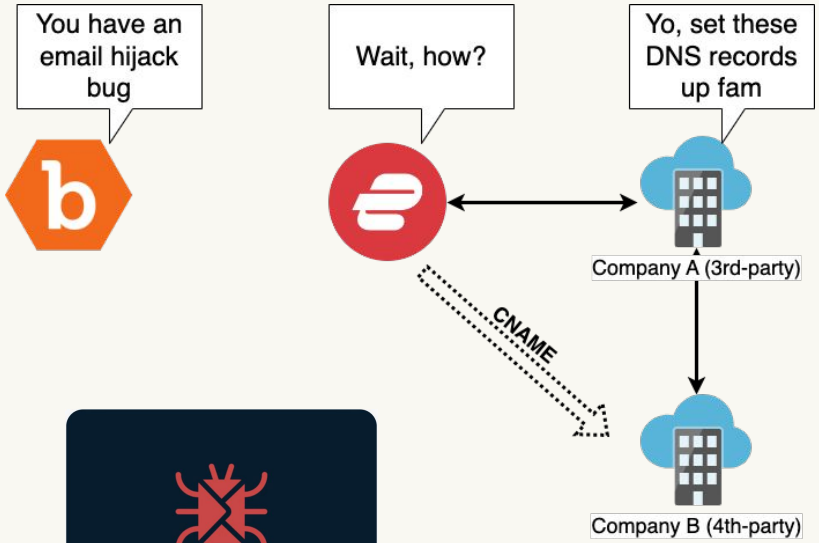

Why did this happen?

- CNAME record to Company B infra leads to MX/TXT inheritance (SPF)
- Loose account validation for Company B services through DNS records
- Undocumented SMTP feature flags allowed SPF bypassing

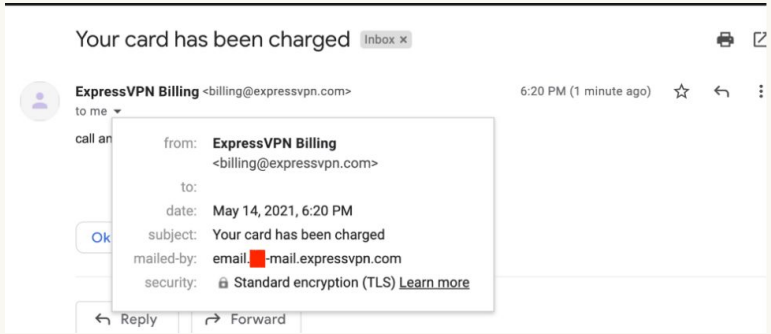

Context & Awareness

How do we fight these threats?

Traditional SOC Model

- Vendor tool-driven alert triage
- “Follow the sun” model with defined handoffs
- L1/2/3 Incident Triage and Response
- SOPs with formal incident declaration steps
- Vendor-centric Intelligence consumption (feeds of atomic IOCs that require integration)
- Point in time environment information gathering, asset enumeration, data flow assessment
- Outsourcing of operational components (hybrid SOC)

“SOC Of The Future” model

- Operationalizing Intelligence (*)
- Threat Modeling (*)
- Agile Dev/Sec Ops
- Data Science & Analytics (*)
- (Continuous) Threat Hunting (*)
- People - skills, not tiers. Diversity, not completeness
- Dichotomy - Procedural maturity vs creativity
- Automation
- Orchestration and enrichment

Challenges with each model

Traditional SOC Model

- Pros
  - Well established guidelines
  - Enterprise-ready products and services offering turnswitch capabilities
  - Divide and conquer strategy with L1/2/3 human resources
- Cons
  - SOC fatigue
  - Issues with scalability (more human resource and tech to solve capacity)
  - Strategic visibility limitation (big picture)
  - Compliance-oriented, disruptive at times in establishing positive inter-team relationships

“SOC Of The Future” model

- Pros
  - Objective & Key Result focused
  - Scalable, flexible and forward-looking
  - Generative and SDLC oriented
  - Metrics centric approach
  - Tactical & Operational correlation offering Strategic oversight
- Cons
  - Steep curve in initial stages of capability development and maturity
  - Bleeding edge of process and technology, you're more on your own
  - Org-level mindset transformation complexity (change is hard)

Our Target

Meet the Victim

ACME Corp. DaaS (Dynamite-as-a-Service)

- Leading company in the Explosives industry
- Big player that is looking to revolutionise the health and safety protocols for the entire mining, quarrying and construction industry
- Huge list of (**Mis**)Fortune500 customers, as well as a comprehensive B2C market share
- Offers a new product pipeline of IoT remote detonation devices, sensor networks and SCADA integrators **all from the Cloud (and a mobile app!)**
- PCI-DSS, SOC2, GDPR, ISO 27001 and more compliant

The network

Dynamite-as-a-Service

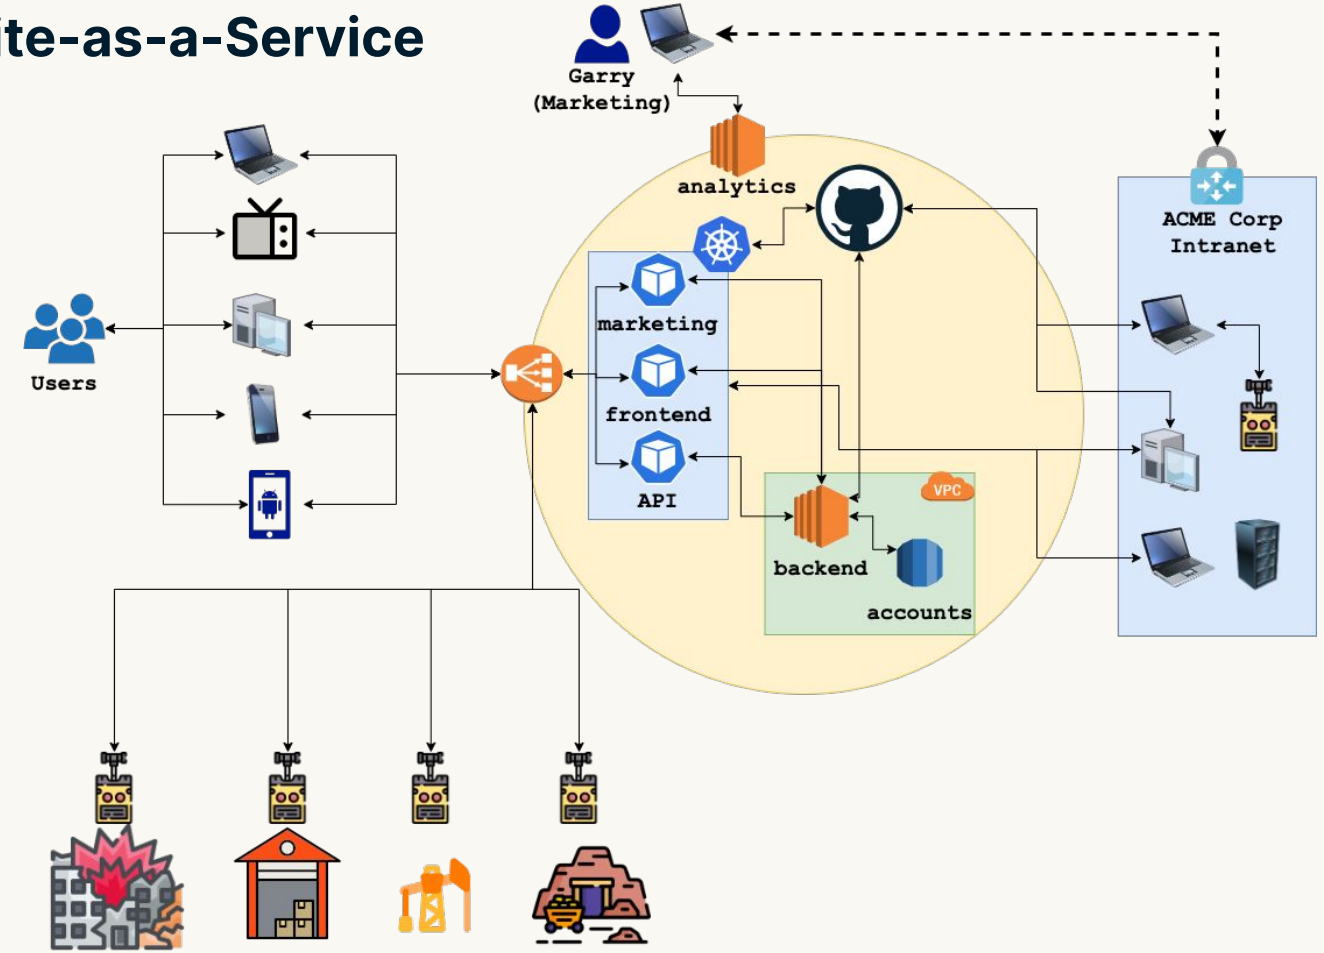

Our Adversary

Meet the Maker

Goodbye Kitty APT group:

- Believed to be associated with a foreign government entity in the Middle East
- Historic evidence of conducting broad scale credential harvesting, account takeover attacks and other forms of public infrastructure exploitation to aid in the next stages of exploitation
- Preferentially targeted critical infrastructure in Western countries
- Employs destructive capabilities in the form of Ransomware
- Leverages LOLBINs and generic malware in initial stages of exploitation, after which more advanced post-exploitation tools are downloaded
- Numerous intel sources suggest affiliation with the ROADRUNNER APT group

Our Adversary

The campaign

What happened?

The Goodbye Kitty APT group has been targeting ACME Corp over the course of a few months including:

- Malspam phishing targeting employees (**blocked** by Email Security Gateway)
- Targeted lures dropping generic trojans (**blocked** by EDR)
- Scanning/enumeration/exploitation of public facing services (**blocked** due to tight AWS Security Groups and WAF)
- **(1)** Bruteforcing of customer accounts through a fast-flux botnet (Layer 7 attacks, resulting in **successful compromise of a number of customer accounts, including those of developers**)
- **(2)** Private GitHub accounts compromised due to **credential reuse** resulting in secrets theft from **accidentally committed AWS secrets**
- **(3) K8s cluster compromise** due to mismanaged privileges
- **(4) Lateral movement to API Pod** and **scheduled detonation of all explosives** for all backend accounts

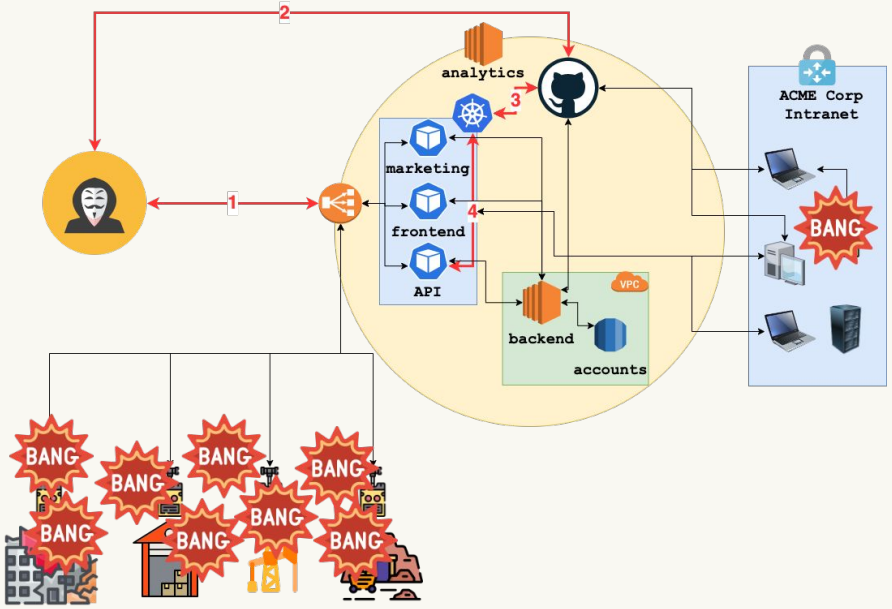

The path of more resistance

What else could have happened?

- What if Garry's offsite device was breached to attack ACME Corp's Intranet?
- What if the adversaries bought the IoT products and reverse engineered them in an attempt to exploit the API service and/or find vulnerabilities?
- What if a developer's workstation were compromised and their local AWS profile secrets were compromised
- What if the adversary pushed a malicious commit through Github and triggered a CI/CD build that provided them with admin access to the pods?
- What if there were an insider who sold access to the AWS infra?
- What if one of the build dependencies (latest) were compromised in a supply-chain attack?
- What if.....?

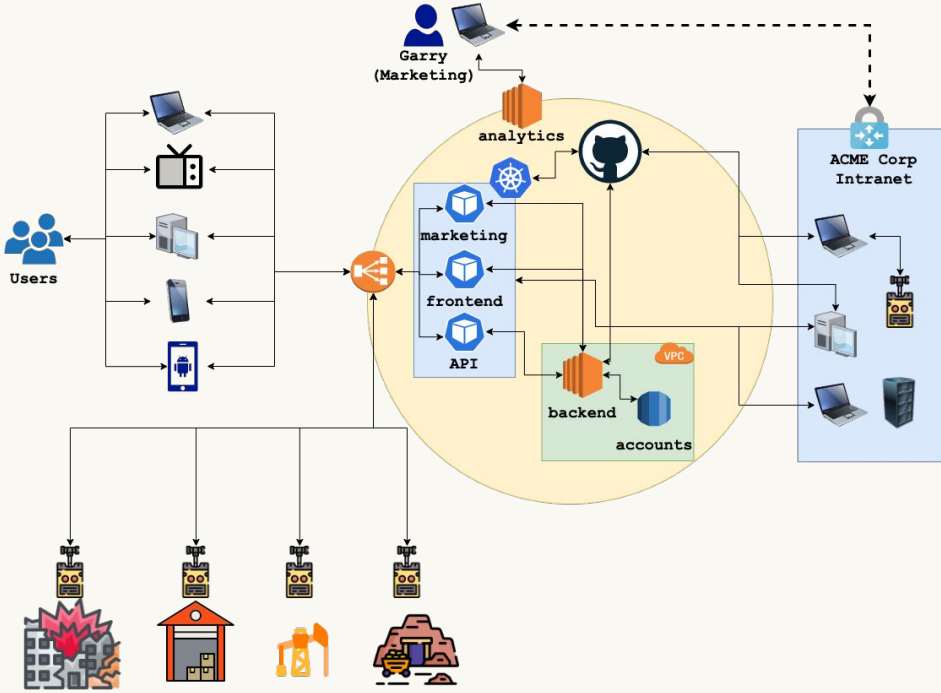

Let’s rewind the tape

Security Operations Pipeline

Security Operations as a Service

The funnel

Service *X* dev team

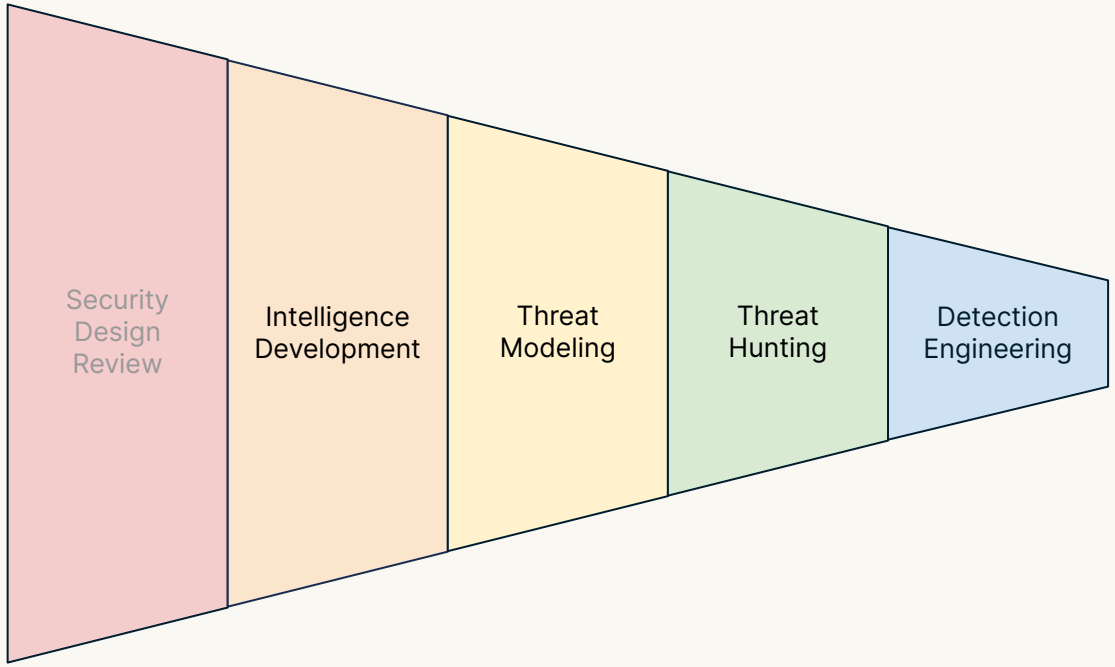

Rock Industries Field Engineer

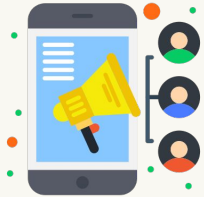

SME Customers

Rock Industries CEO

Intelligence

Know thy enemy

Definitions over the internet include:

- (Military) intelligence is a discipline that uses information collection and analysis approaches to provide guidance and direction to assist commanders in their decisions.
- the ability to learn or understand or to deal with new or trying situations
- the ability to apply knowledge to manipulate one's environment or to think abstractly as measured by objective criteria (such as tests)

Regardless of definitions it boils down to one thing - ***information is power, and it can give you leverage against your adversaries***.

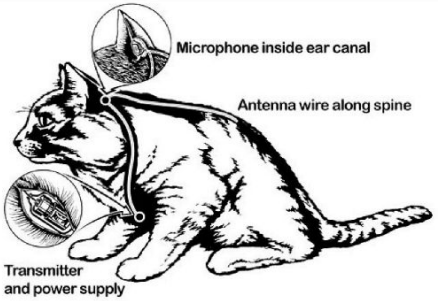

Operation acoustic kitty. CIA spent $15 million on this project. The cat was meant to walk up to Soviets and spy on them. First field experiment ended immediately when the cat got hit by a taxi and died

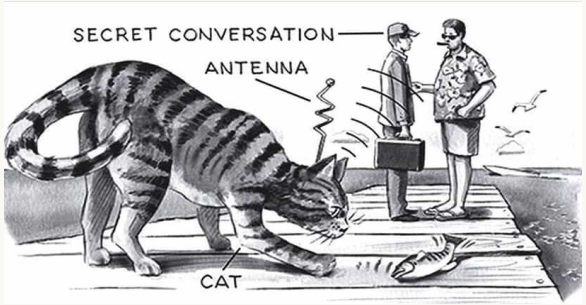

Domains of Intelligence

Threat Intelligence

- Who are my adversaries? How do they exploit my systems? What TTPs do they have? What do they aim to achieve?

Brand Intelligence

- Who is impersonating my brand in an attempt to cause monetary or reputational damage?

Fraud Intelligence

- Who is attempting to defraud us and how are they doing it?

Vulnerability Intelligence

- What is the vulnerability threat surface of my entire organisation? What dependencies and packages do all my systems have and what is their severity and priority for patching?

Operational Intelligence

- How are users leveraging my systems? What anomalies can I spot to improve my resilience, cut costs and reduce (detonator/client app/service) abuse?

Other

- .....

Building out your Intelligence program

1. Collating **Priority Intelligence Requirements** (PIRs) - solicit advice from all your stakeholders and partners internally on what are their biggest concerns about they systems, functions or processes.
2. Refining them into **Information Requirements** (IRs) - break down PIRs into more targeted questions which we can more accurately measure and answer
3. Create a **Collection, Integration, Action** (CIA) plan - Identify data that can support answering IRs
   - Identify data capable of answering IRs in the form of a **Collection Requirement** (CR)
   - **Integrate** the data into your systems for further analysis, enrichment and orchestration
   - Formulate a **specific plan of Action** in how the refined intelligence product will **drive decisions** that have impact

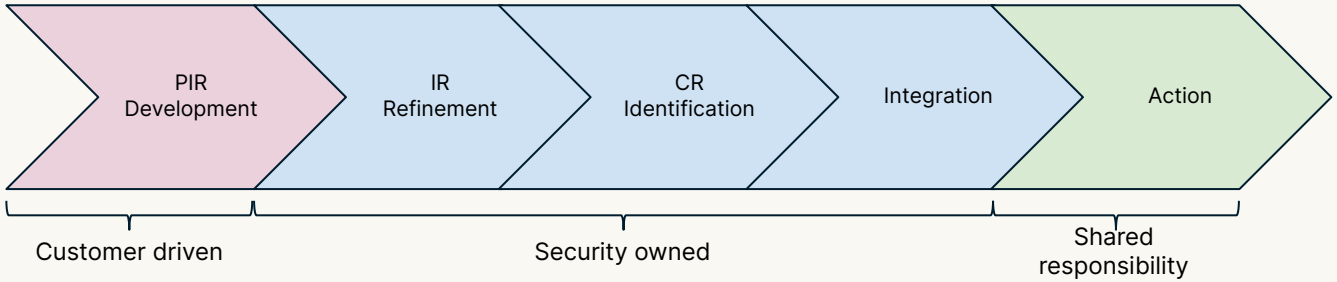

Intelligence

ACME Corp's Intelligence Program

(PIR) IntReq ProdOps-1: The Product Operations team requires information on the methods of which malicious actors are abusing our service and using it maliciously.

InfoReq ProdOps-1-a: What bruteforce scripts can we identify across D&DW, Telegram, Discord and other hacktivist channels with references to ACME Corp's public facing infrastructure?
InfoReq ProdOps-1-b: How many of our existing customers' credentials are exposed in public data breaches?
InfoReq ProdOps-1-c: What volumetric data, or patterns of access, across our public facing infrastructure constitute credential stuffing attempts?

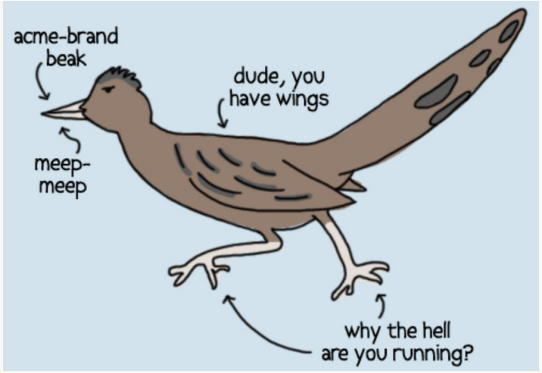

CR ProdOps-1-a

- Create a Yara signature for popular bruteforce script configuration files, like OpenBullet, amd deploy it on VirusTotal to collect uploaded bruteforcers
- Configure D&DW trigger notifications across vendor platforms for mentions of ACME Corp and account cracking mentions.

CR ProdOps-1-b

- Create continuous monitoring via HaveIBeenPwned for our customer and employee base

CR ProdOps-1-c

- Sample CloudFront and AWS WAF logs via the Pareto Principle with the expectation that the top 20% of IP addresses would account for 80% of the traffic that may be malicious
- Create baseline models for expected customer behaviour and construct an ML model that singles out anomalies based on HTTP request telemetry

Integration & Action - Feed the data into a centralized data pipeline where enrichment can happen and accounts or attacker infrastructure can be suspended

Intelligence != Threat prevention completeness

ACME Corp may have identified means of answering the most critical questions through PIRs, but does that provide **security completeness**?

In comes **Threat Modelling (TM)** to the rescue!
Structured process where we try to proactively look at all potential threats to a system, service, team or entity that a malicious threat actor, or adversary, might be interested in exploiting.

Existing TM methodologies include STRIDE, PASTA, Attack Trees, Playing Cards, etc.

ACME Corp needed something that does not just generate a list of bad things that can happen, but also to **integrate it into their SDLC practices**.

Similar to TDD, Threat Modeling is a core component of securely designing systems by engineers themselves through the Security team's years of expertise in studying adversaries.

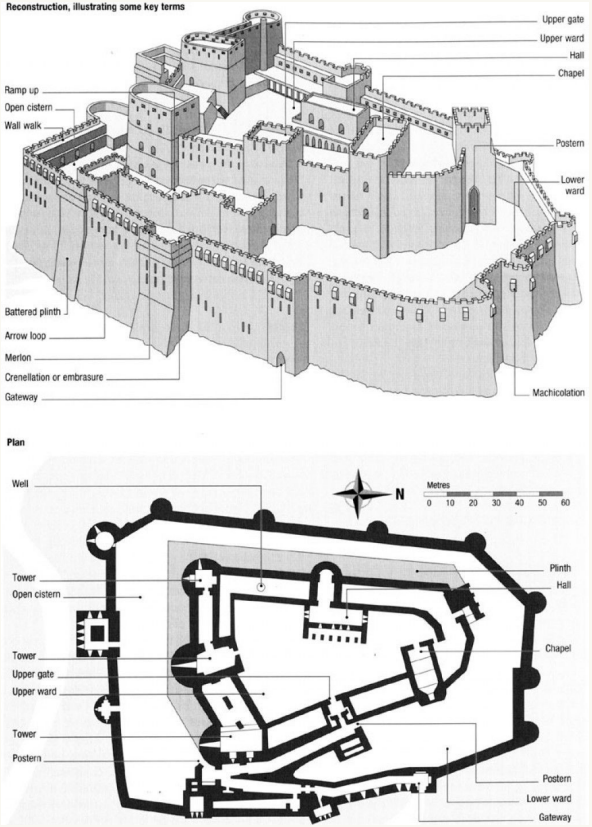

Painting the threat surface

Similar to development we can create “**Malicious** User Stories” in the form of Threat Scenarios.

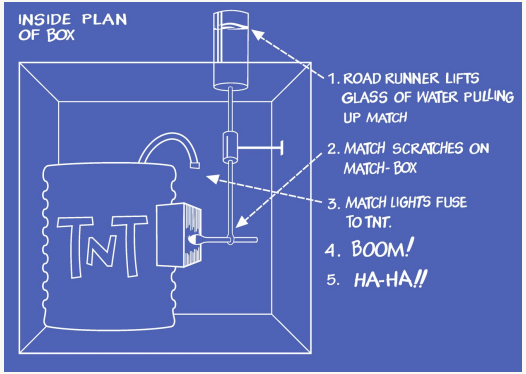

Work with engineers to decompose their systems end-to-end in creating atomic **Threat Scenarios** that describe what things can go wrong at each stage. This would include:

- What are the ways adversaries can obtain access to this system
- What can they do with pre-existing access? How can they expand it?
- Can they laterally move to systems in proximity to the crown jewels?
- How can they bypass controls and evade detections?

What’s next? Prioritise these by their overarching Threat Score based on the Threat Scenario’s **Severity**, **Impact**, **Complexity** and **Likelihood** of happening. Offer a “discount” based on the degree of **Controls** that are present.

LOW	MODERATE	HIGH	CRITICAL
1-39	40-69	70-89	90-100

$$Threat\ Score = (Threat\ Rating - Threat\ Controls) \times \frac{100}{Max(Threat\ Rating)}$$

$$Threat\ Rating = Severity + Impact - Complexity + Likelihood$$

$$Threat\ Controls = Threat\ Rating \times \frac{Controls - 1}{Max(Controls)}$$

Assess the Threat Scenarios using **Mitre ATT&CK** to create a technique map, capable of offering insights on what to prioritize for **Control Recommendations**.

Threat Modelling

Sample Threat Scenarios

Threat - Initial Access through API service RCE
Threat Scenario - A threat actor enumerates and exploits API service endpoints in an attempt to obtain a foothold into ACME Corp's environment
Severity - 4 (Significant CIA impact, direct access to privileged systems)
Impact - 3 (Confirmed business impact, extended actions required to contain/resolve the incident potentially including service downtime)
Complexity - 3 (More tailored exploitation and service enumeration observed, good recon, some defense evasion)
Likelihood - 4 (High value asset or stepping stone to such)
Controls - 2 (AWS WAF rules + rate limiting)
Threat Score - 44
Threat Category - Moderate
ATT&CK Techniques - T1595, T1190, T1027, T1203, T1059
Control Recommendations - Container hardening, code auditing/external pentest, command execution monitoring

Threat - Detonation of customer devices
Threat Scenario - A threat actor with pre-existing access to ACME Corp's API service may issue a scheduled detonation for all customer accounts
Severity - 5 (Max severity, destructive nature, potentially harming human lives)
Impact - 5 (Max business impact, catastrophic monetary and reputational loss)
Complexity - 1 (Low complexity, pre-existing access and API secrets already present in environment variables)
Likelihood - 5 (Max likelihood, high reward for low effort)
Controls - 1 (No controls in place)
Threat Score - 100
Threat Category - Critical
ATT&CK Techniques - T1078, T1053, T1528, T1552
Control Recommendations - Configure 2-step verification for privileged customer actions, create per-customer PKI detonation validation, monitor privileged calls

Threat Hunting

Building on top of Threat Models

With the contextual information about the service threat landscape ACME Corp's Threat Hunters can go on the prowl.

The Threat Scenarios can be expanded into **hypothesis-driven hunting engagements.**

Successful hunts are one which are structured, goal-oriented and scoped well.

Palantir's **Alerting and Detection Strategy** (ADS) framework represents an excellent method of doing this, as it formalises the hunting objective and turns it into a monitoring solution that is structured.

Threat Hunting

Hunting for the big Boom

Goal	Detection of malicious scheduled remote detonation jobs for a large volume of customers
Categorisation	T1078, T1053, T1528, T1552.
Strategy Abstract	The ADS looks for abnormal volumes of RabbitMQ jobs being scheduled for detonation, which go above a 2 week moving average threshold.
Technical Context	<Technical Context section for responders elaborating the architecture and context of the data source in a SIEM with steps on how to confirm if this is a FP or not>
Blind Spots and Assumptions	RabbitMQ logs are continuously fed into the SIEM and no form of data manipulation is happening
False Positives	False positives may occur on Mondays when construction work begins and detonations are scheduled at larger volumes.
Validation	A Lambda function can be invoked to populate test scheduled jobs
Priority	Critical
Response	The API service team is automatically paged out and the job queue is temporarily halted
Additional Details	Links to architecture diagrams, repositories and other resources

But wait, how is this different from regular alerts?

Trying harder

Targeted hunting engagements can produce alerts, but way too often the tuning and allowlisting creates gaps a sophisticated adversary can exploit. We should not replace them fully, but we should stay conscious that if a condition fails for the tuned alert - we may be one step closer to kaboom (in our case)**.** It's **better to catch everyone eventually, than to catch only some instantly**.

We should try harder to look for outliers more frequently and not await the alert's review cycle to discover deficiencies in an alert **and what the sliced up data can offers us in addition to it.**

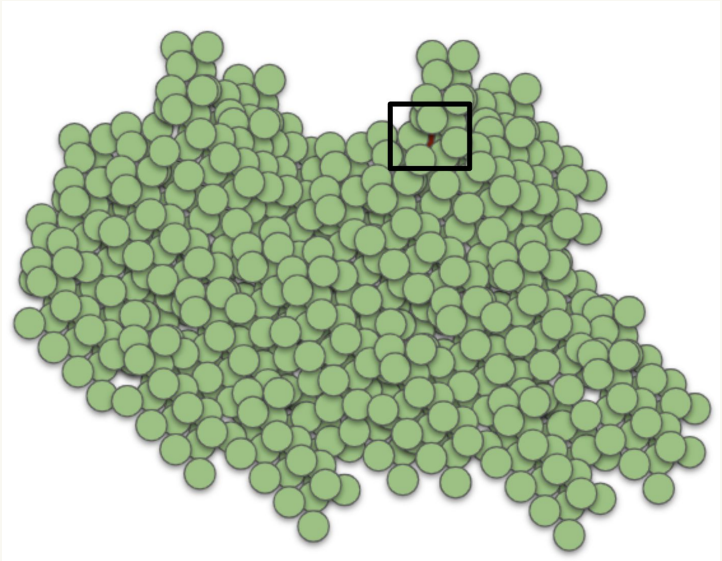

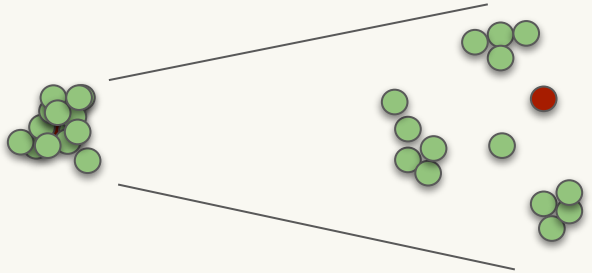

Digging deeper

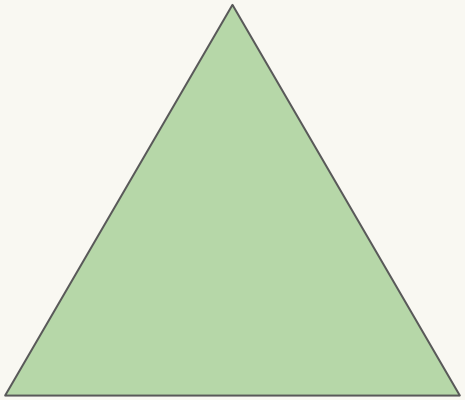

This is where we should be spending most of our time

Refine

Transform
Multiply
Transition

Data “mine”
“Millions” of events
Good for blue sky research and play
Bad for actual discovery
Example: All Powershell

Haystack
“Thousands” of events
Good for discovery
Bad for constant review
Example: Powershell with command lines longer than 200 characters

Signature
“Ones” of events
Good for starting IR
Bad for humans
Example: Powershell with Mimikatz keywords

The Enabler

None of the before mentioned topics would be possible without the power of Automation and data enrichment.

Whether it's :

- Slackbot notifications that directly allow engineers to confirm security events
- Splunk IaC to manage all alerts
- Sigma to abstract down complex SPL and to make it more maintainable
- Playbooks for Cloud containment, response, data enrichment or many many other

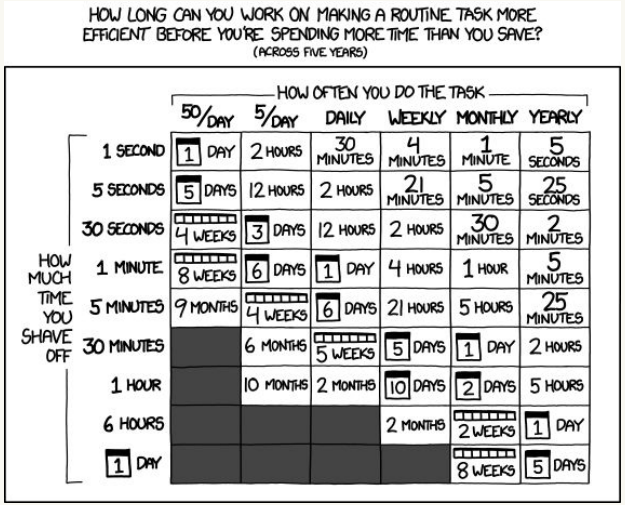

Automation saves time, which can be used towards hunting for adversaries.

Q&A

Jobs: expressvpn.com/jobs