Introduction to Threat Hunting

Carnegie Mellon University
Software Engineering Institute

Notices

This material is based upon work funded and supported by the Department of State under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center sponsored by the United States Department of Defense.

DM23-0953

Modules

1. What Is Threat Hunting – Definition & Types
2. Threat Hunting Cycle
3. Adversary Frameworks for Threat Hunting
4. Reading the Threat Landscape
5. Hunt Teams, Roles, Skills
6. Implications of AI on Threat Hunting

What Is Threat Hunting – Definition & Types

Carnegie Mellon University
Software Engineering Institute

What is threat hunting?

Threat hunting is proactively searching for evidence that you are about to be targeted, being targeted, or have already been compromised.

Quotes

“[Hunting is] the ability to proactively search through network and configuration data with the goal of identifying events or misconfigurations that would be indicative of malicious activity... A prerequisite for performing a quality “hunt” is having a high degree of visibility and introspection into your network and endpoints.”

- Chris Lee, Palantir

“Hunting almost always requires investigators to pull data from multiple systems and make sense of it, needing to fetch, join, and normalize disparate data in order to answer specific questions.”

- Ely Kahn SQRRL, AWS, (now at SentinalOne)

The Name Is New – But the Steps Are Not

Some Professionals call it *"IOC Free analysis"*

Others say it is "*Proactively Looking for Incidents instead of being Reactive*"

Sample definitions from others

Definition	Source
"Threat hunting is the practice of proactively searching for cyber threats that are lurking undetected in a network. Cyber threat hunting digs deep to find malicious actors in your environment that have slipped past your initial endpoint security defenses."	What is Cyber Threat Hunting? By Scott Taschler at CrowdStrike (April 17, 2023)
"Threat hunting [] moves the bar for network defense beyond looking at the known threats and allows a team to pursue adversaries that are attacking in novel ways that have not previously been seen."	Maurice, Chad, et al. The Foundations of Threat Hunting: Organize and Design Effective Cyber Threat Hunts to Meet Business Needs. N.p., Packt Publishing, Limited, 2022.
"Threat hunting assumes that compromise has already happened in some way, shape or form. The process involves proactively searching for cyber threats, vulnerabilities, and malicious actors hiding in your environment that have somehow escaped detection by the rest of the security toolset."	A Threat Hunting Primer by Innovate Cybersecurity (November 22, 2021)

There are many more definitions

Some Say Threat Hunting Is NOT

1. Generating Alerts or Responding to Alerts
2. Receiving a list of Indicators of Compromise (IOCs) and running scans for matches to the list in your own environment
3. “A passive activity”
4. Incident Response …
5. A product, it is not automated, it is not something you can put in a script or flow chart – But Automated processes and tools can help with efficiency during hunts

There is also debate on all of the above items.

Some Types of Hunting

Internal

- Search for evidence of compromise in our internal environment
- Search for insider threat activity

External

- Is there evidence of our compromise outside out environment? Is our customer data being sold?
- Are people talking about attacking us? Or attacking software that we use?

Hybrid – both Internal and External

Laying Traps – Proactive activity to improve future detections

- If we aren't compromised using a specific attack yet, can we make a HoneyPot or HoneyToken or HoneyData to alert us when we are?

Threat Hunting Cycle

Threat Hunting Cycle -1

1. Understand threat environment
2. Understand organization's environment and establish scope
3. Create hypothesis – develop some goals for the hunt
4. Collect data
5. Perform analysis, develop results
6. Communicate results, distribution
7. After actions

Threat Hunting Cycle -2

Carnegie Mellon University Software Engineering Institute

Gather Intelligence on Threat activity

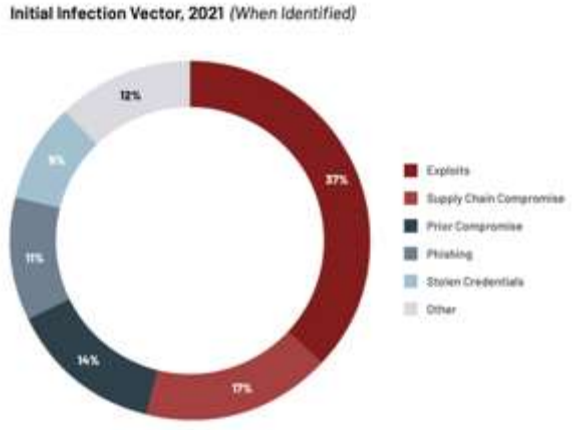

Understand Internal Priorities & Scope

Develop Hunting Hypothesis

Threat Hunting Cycle -3

Carnegie Mellon University Software Engineering Institute

Perform Data Collection → Perform Analysis & Develop Results → Communicate Results


DRAGOS

A simple example can be observed in Figure 2 in the form of an Excel spreadsheet filled out at an electric grid control center and its transmission level substation. Here, the question type is structured against the intrusion kill chain phases, noting that the analyst can answer the types of questions related to that phase of the kill chain, such as exploitation of the system.

	CONTROL CENTER	CONTROL CENTER	CONTROL CENTER	TRANSMISSION SUBSTATION	TRANSMISSION SUBSTATION
ASSET TYPE	Windows Human Machine Interface	Data Historian	Network Monitoring Appliance	Windows Human Machine Interface	Remote Terminal Units
DATA TYPE	Windows Event Logs	Alarms	Alerts	Windows Event Logs	Syslog
QUESTION TYPE (KILL CHAIN PHASES)	Exploration, Installation, Actions on Objectives	Actions on Objectives	Internal Reconnaissance, Command and Control, Delivery, Actions on Objectives	Exploitation, Installation, Actions on Objectives	Installation, Actions on Objectives
FOLLOW-ON COLLECTION	Registry Keys	Set Points and Tags	Packet Capture	Registry Keys	Controller Logic
DATA STORAGE LOCATION	Enterprise SIEM	Local	Enterprise SIEM	Local	Local
DATA STORAGE TIME	60 Days	120 Days	30 Days	30 Days	7 Days

Figure 2: Sample CMF of a Hypothetical Electric Company

Collection Management Frameworks often go beyond a typical IT asset inventory to include logs, data sources, duration and more

Gather Intelligence on **Adversary Behavior** and **Recent Activity**

- The team collects information to decide on the most probable threats
- What attackers are actively doing
- Attacks against multiple organizations
- Attacks against similar organizations

Example Threat Activity Report

Example Items of Interest

- 71% of incidents were “malware free”
- Ransomware
- Access Brokers are using specific vulnerabilities, then reselling the access.
- Active Threat Groups and their tactics
- 80% increase in attacks on Financial Service
- Cross Platform Attacks

We cover ***Attacker Behavior Frameworks*** and ***Reading the Threat Landscape*** in more detail in another Sub-Module

Understand Internal Priorities and Scope

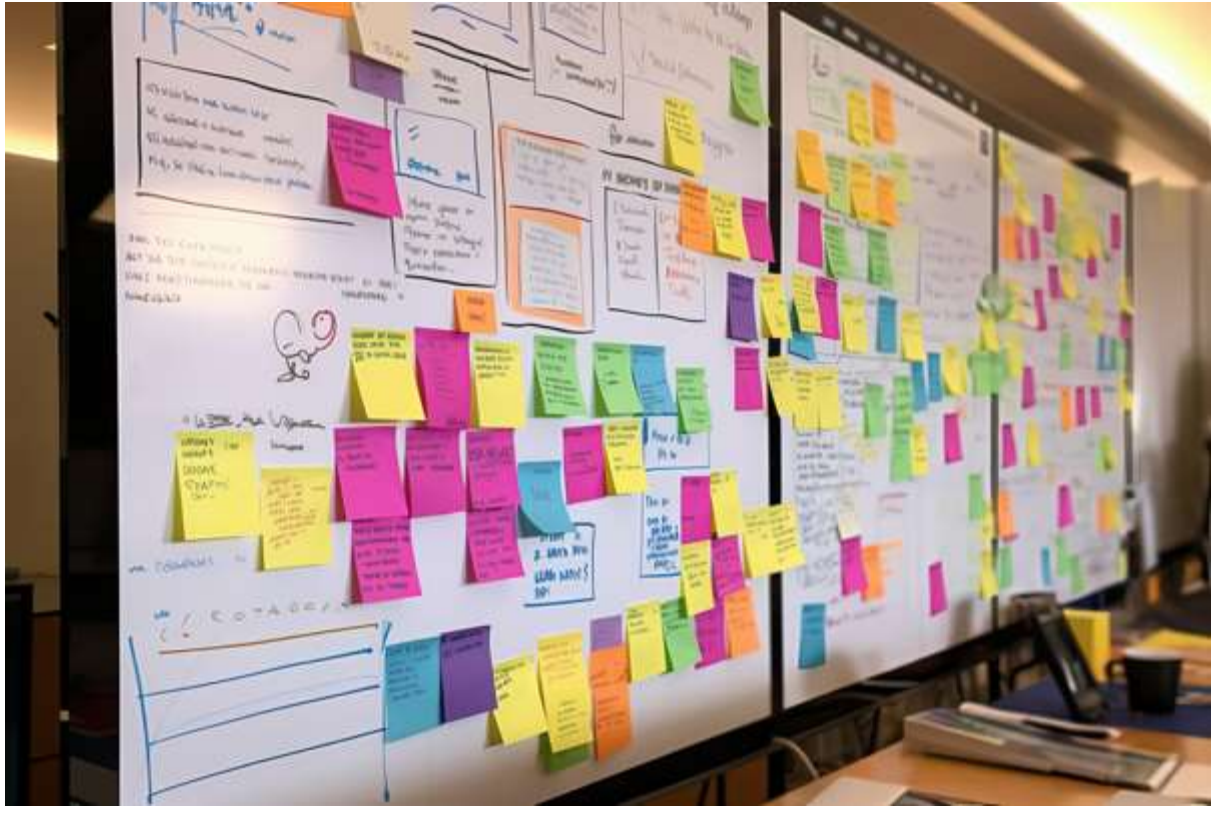

- Align hunts with organization's key assets and internal priorities
- Identify key activities occurring soon
  - Critical Business Processes and Special Events, Acquisitions/ Launches, R&D, etc.
  - What assets and information support these activities?
  - Under what adverse IT conditions would these activities fail?

Develop Hunting Hypothesis – Questions

Combine information from intelligence sources with internal priorities.

Can be an artform. Usually not easily automated.

Threat Hunting often starts with the assumption that the organization has been compromised, and initiates activities to look for evidence to see if it is true. The activity is almost always informative.

Developing Hunting Hypothesis

Carnegie Mellon University Software Engineering Institute

Question: How is it possible to respond to an incident if you do not know it exists?

Answer: Assume that you have been compromised, ask yourself or your team:

1. How might it have happened?
2. What is a list of weaknesses or tools that adversaries might try to use?
3. What are the most likely paths they would choose to compromise us?
4. What are common techniques attackers are using to compromise other organizations and How might they use those to compromise us?

Example Data Driven Hypothesis Questions

Carnegie Mellon University Software Engineering Institute

From Intelligence Reports:

- **CrowdStrike reports the Kerberoasting technique was used in 583% more incidents in 2022 than in 2021.**
- Attackers were stealing tickets that are associated with Service Principle Names (SPNs).
  - SPNs are often tied to service accounts,
  - Service accounts often have higher administrative privileges.
  - After identifying service accounts, attackers would also use HashCat to brute force PW hashes

Sample Hunt Questions

- What environments are we using service accounts in?
- Do any of our service accounts have unnecessarily high privileges?
- Have attackers already been able to Kerberoast us?
  - Have they already stolen any accounts or hashes this way?
- If not, how easy would it be for them to do so?
- Would we be able to detect it if they did?
- What processes would be affects?
- Are we logging and alerting on those?

Understand How Threats Might Attack You

Carnegie Mellon University Software Engineering Institute

Kerberoast Attack Diagram by RedSiege

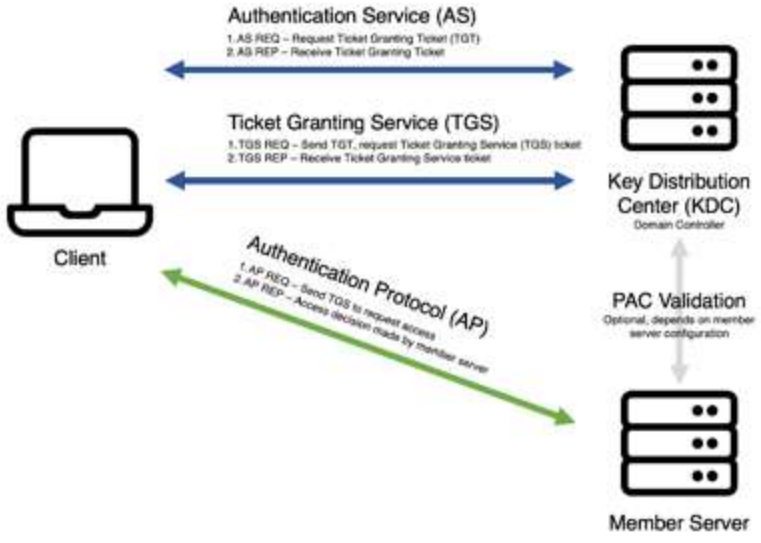

https://redsiege.com/tools-techniques/2020/10/detecting-kerberoasting/

Become familiar with the Technical Details of the Kerberoasting and Detection techniques.

After understanding the attack, ask yourself and your team members

- What tools were used?
- What traces do these specific tools leave on my network?
- What logs are they left in?
- What other behaviors would need to occur?

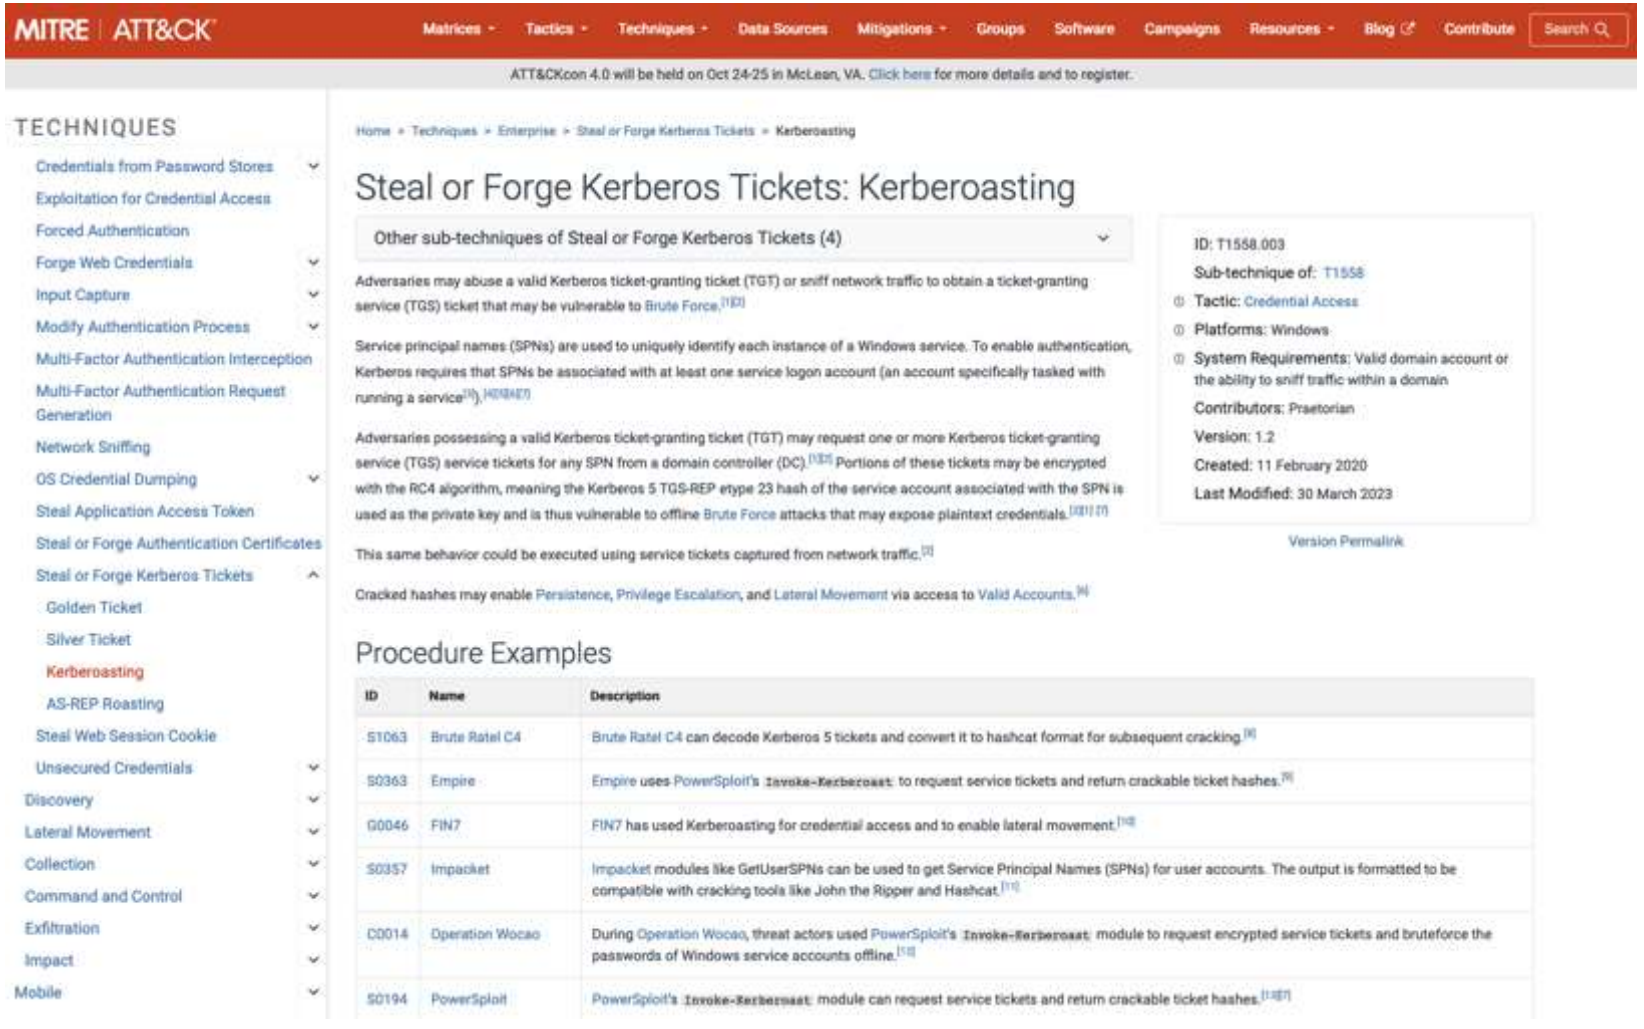

We cover ***Attacker Behavior Frameworks*** and ***MITRE ATT&CK*** *in more detail in the next sub-module*

Data – Collection Management Framework

When answering the hypothesis question, a CMF will contain information about data sources that might be relevant to the hunt.

Hypothesis 1 - here's where the educated guess goes that determines the required data sets					
Data Type	Data Source	Retention Length	Data Owners	Data Collection Method	Data Value
logs	webserver	24 hours	NOC	Manual	medium
logs	dns	7 days	NOC	Manual	high
logs	proxy	72 hours	NOC	Manual	medium
binary logs	exchange	45 min	NOC	Manual	high
binary logs	endpoint	variable	Security team	Manual	critical
logs	antivirus console/endpoint	90 days	Security team	Manual	critical
logs	firewall	24 hours	NOC	Manual	high
packet capture	IDS appliance	48 hours	Security team	Manual	critical
NetFlow	IDS appliance	14 days	Security team	Manual	critical

Figure 5.1 – Collection Management Framework

Image Source: Maurice, Chad, et al. The Foundations of Threat Hunting: Organize and Design Effective Cyber Threat Hunts to Meet Business Needs. N.p., Packt Publishing, Limited, 2022.

Kerberoast Example Hunting – Data Collection

Develop a plan to collect the relevant data, logs, and configurations to perform the hunt.

There may be a lot of data and it may also be sensitive.

- Credentials
- Tickets
- Sensitive Communications
- Security configurations
- Permission logs

The hunt team will need procedures to secure the data.

Kerberoast Hunt Example

- Windows Event Log
  - Look for ticket granting service requests and approvals (e.g., 4769, 4770)
- List of Service Accounts, the IT systems they are on, and the criticality of those systems
- Endpoint logs from the systems with service accounts.
- And more

Perform Analysis & Develop Results

The team works to identify patterns, build connections, and attempt to answer the hypothesis question.

Questions to consider

- What was executed on the machine?
- Are any service account actions over the past few days/weeks suspicious?
- Were any new accounts created recently? Are they all accounted for?
- Have any suspicious login attempts (successful or failure) been logged?

Use tools to help with analysis.

Example: Kerberoast Windows Event Log Data Analysis

Windows Event Log entries… Are these bad?

EventID	Date	AccountName	ServiceName
4769	1/25/2017 9:36:07 PM	JoeUser@LAB.ADSECURITY.ORG	svc-VDIPVS01
4769	1/25/2017 9:36:07 PM	JoeUser@LAB.ADSECURITY.ORG	Svc-BizTalk01
4769	1/25/2017 9:36:07 PM	JoeUser@LAB.ADSECURITY.ORG	SVC-BOADS-01
4769	1/25/2017 9:36:07 PM	JoeUser@LAB.ADSECURITY.ORG	SVC-AGPM-01
4769	1/25/2017 9:36:07 PM	JoeUser@LAB.ADSECURITY.ORG	svc-adsMSSQL10
4769	1/25/2017 9:36:07 PM	JoeUser@LAB.ADSECURITY.ORG	svc-adsSQLSA
4769	1/25/2017 9:36:07 PM	JoeUser@LAB.ADSECURITY.ORG	svc-adsMSSQL11
4769	1/25/2017 9:36:06 PM	JoeUser@LAB.ADSECURITY.ORG	SQL-ADSDB317-SVC

Example logs and images from https://adsecurity.org/?p=3513

Findings

- Event 4769 : a Kerberos service ticket was requested
- Ticket Encryption Type = 0x17 : uh oh, this indicates a weak encryption algorithm (RC4) was chosen during configuration
- Tickets Requested for Several different Service Names (BizTalk, Microsoft SQL Service, and more) by the same account within microseconds of each other.

Analysis: What Is the Impact?

Where did we find this activity?

- *We found this on a domain controller (DC)*

What departments or network activities does this particular DC provide service to?

What service accounts are potentially a problem?

- *In our example, look at the Service Names to help answer this question!*

What other assets and applications does this service account have access to?

How recently has this service account accessed other IT assets or applications?

What business activities do those other IT assets support?

What other logs can be collection from their machine.

Should the hunt team pass this over to the incident response for additional investigation and forensic analysis?

- *In our example, yes, asap. And probably assume that the service account password is cracked*

Communicate Results

- What do stakeholders need to know about what you found / find?
- How will you communicate it to them so that it remains relevant ?
- Try to collect feedback on the usefulness of the results to the audience?
- Surveys, testimonials, delivery platforms.
- If activity is found, get the incident response team involved.
- Communicate Technical information at the executive level to stakeholders.

Report Templates Can Help with Communication

Carnegie Mellon University Software Engineering Institute

Report templates can help ensure teams perform the right activities during hunts.

Templates help audiences organize and understand complex information.

Some threat hunting teams share their templates.

Templates can be coded into tools or platforms.

Sample Report Template from CyborgSecurity

1. **Executive Summary**: A concise overview of the threat hunting operation and its results.
2. **Abstract & Hypothesis**: Describes the focus of the hunt and proposes a theory about the potential threat.
3. **Technical Summary**: Provides in-depth insights about the hunt, the technologies used, and the results obtained.
4. **Mitigation Recommendations**: Outlines proactive steps to mitigate identified threats.
5. **Analysis**: Transforms raw data into actionable intelligence, highlighting significant patterns or indicators of compromise.
6. **Conclusion**: Wraps up the findings, summarizing what was achieved during the threat hunt.

Source: https://www.cyborgsecurity.com/hunter-platform/building-an-effective-threat-hunting-report-template/

After Actions

Perform an Internal Review of the Hunt

What challenges were faced?

How could those have been overcome faster?

What went well? What did not go well? What improvements can be made?

What metrics were collected?

What new metrics would be useful to start tracking weighed against how much will these cost to track.

Example After Action observations:

- Infrastructure Changes
  - We need to switch from RC4 to AES encryption for tickets
  - Ensure all DC are logging Kerberos related Windows event IDs
  - Make sure all service accounts have very strong passwords to prevent fast cracking
  - Filter the logs in the SIEM for requests, weak encryption, etc.
- Our IT team does not have a central list of all service accounts, the Hunt team had to contact business units on separate networks to collect a full-service account list.
- The transportation team does not maintain a mapping of IT server to critical processes: we had to construct it with them during the hunt

Other Hunt Planning Tips

Collaboration with other business groups is often needed during a hunt.

Consensus is often needed between the hunting team and other teams, such as IT, strategy, operations, finance, and others.

Methods of hunting may need to be evaluated and weighed for cost, effectiveness, time to implement, and impact on current operations.

A combination of different experts and perspectives is needed during hunting activities.

There ARE other approaches. You could start with a broad question or large pile of data and filter it down iteratively.

Also: Parts of the Hunting Cycle Can Apply to Specific Technologies or Areas

Example: ICS environments have unique attack vectors and attack surfaces

CISA slide on the top attack vectors in ICS in 2021

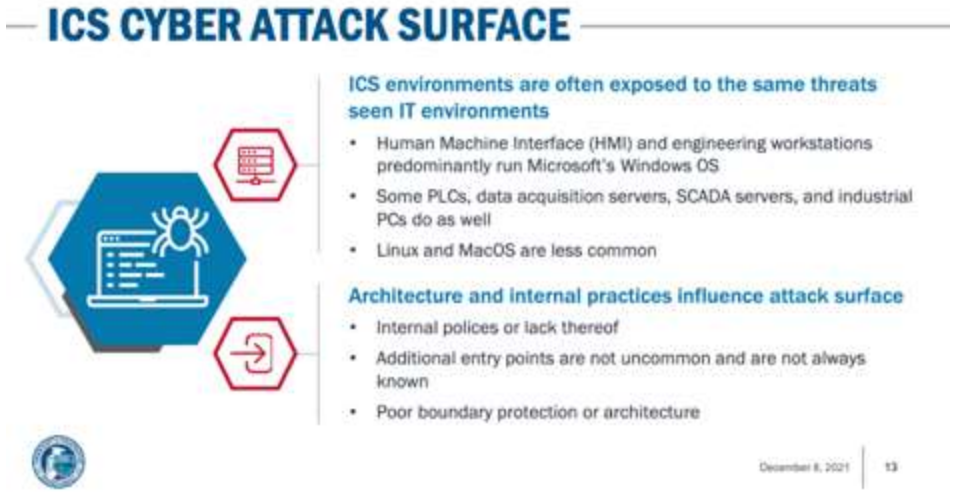

CISA slide on unique ICS attack surface in ICS in 2021

Summary

Threat Hunting Lifecycle

1. Understand Threat Environment
2. Create Hypothesis
3. Collect Data
4. Perform Analysis, Develop Results
5. Communicate Results
6. After Actions

Adversary Frameworks for Threat Hunting

Carnegie Mellon University
Software Engineering Institute

Lockheed Martin: Cyber Kill Chain®

LHMC Cyber Kill Chain® Framework

The term kill chain originates from the military to specify offensive actions.

Lockheed adapted the term to Model common cyber attacker behavior.

Use the chain to identify where you are in an attack.

Disrupting any one link in the chain can prevent attacker's reaching objectives.

LHMC Kill Chain – Example of How to Apply It for Threat Hunting

Carnegie Mellon University Software Engineering Institute

Example of what Adversaries are attempting during the **Reconnaissance** phase.

Adversary Activities

- Harvest email addresses
- Identify employees on social media networks
- Collect press releases, contract awards, conference attendee lists
- Discover internet-facing servers

Recommended Defender Activities and key Data Sources

- Collect website visitor logs for alerting and historical searching.
- Collaborate with web administrators to utilize their existing browser analytics.
- Build detections for browsing behaviors unique to reconnaissance.
- Prioritize defenses around particular technologies or people based on recon activity.

Source: https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/Gaining_the_Advantage_Cyber_Kill_Chain.pdf

MITRE Attack

Carnegie Mellon University Software Engineering Institute

The MITRE ATT&CK® framework is a knowledge base of attacker tactics and techniques.

Designed for cyber operators (including threat hunters) to help create a common vocabulary of all known attacks.

Created in 2013 to enable testing for a research project called FMX

The objective of FMX was to investigate how endpoint telemetry data and analytics could help improve post-intrusion detection of attackers operating within enterprise networks.

The ATT&CK framework was used as the basis for testing the efficacy of the sensors and analytics under FMX

Served as the common language that both offense and defense could use to improve over time.

Lists out the details for the “next level down” in each stage of the cyber kill chain…

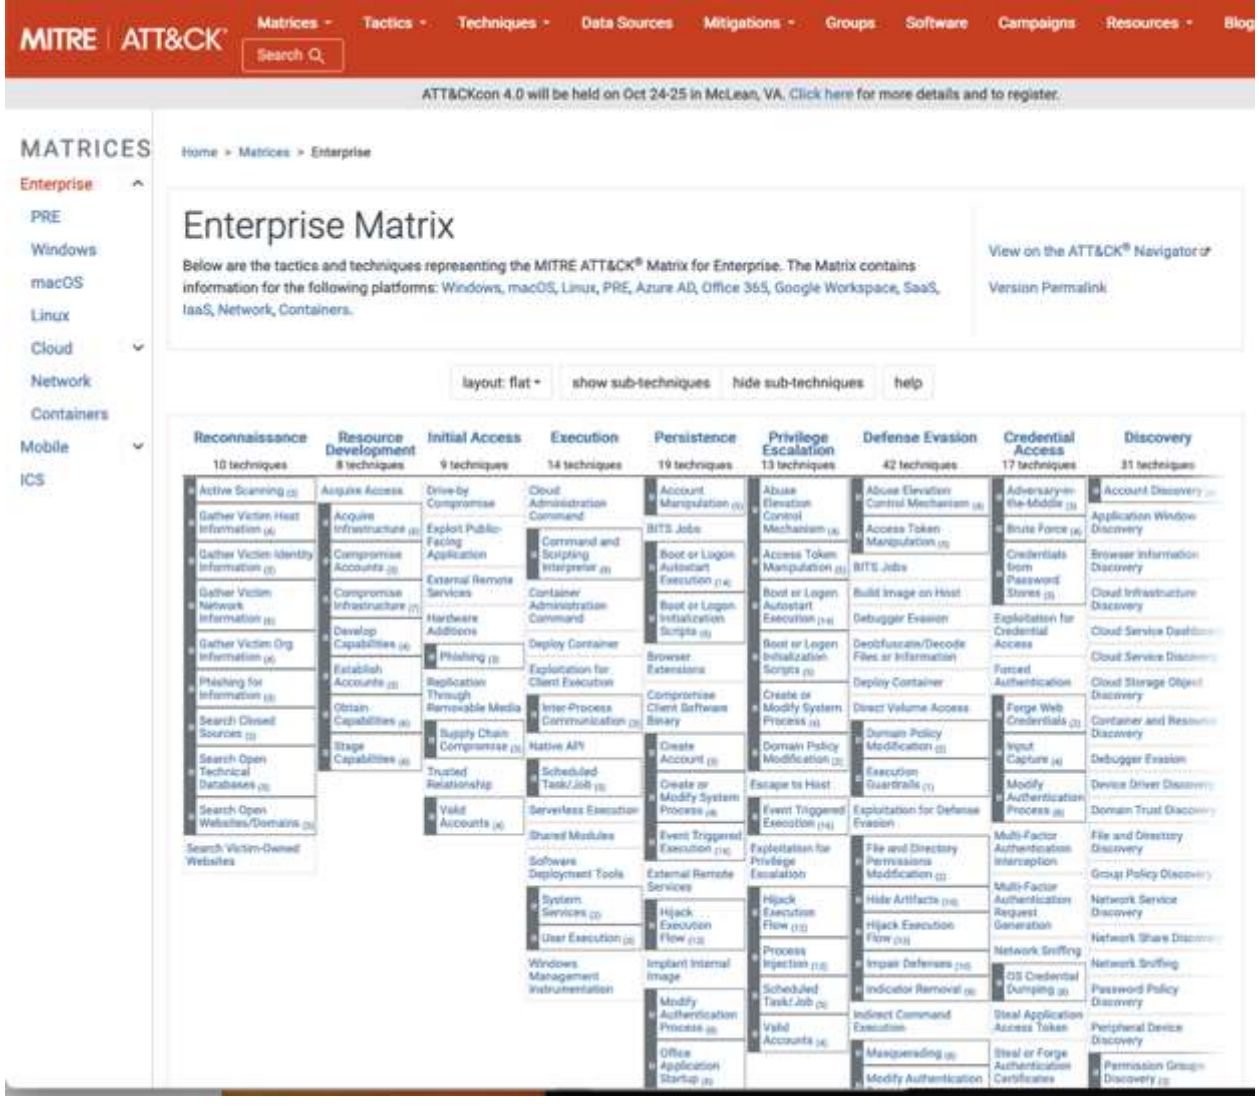

Example from **Reconnaissance** Phase

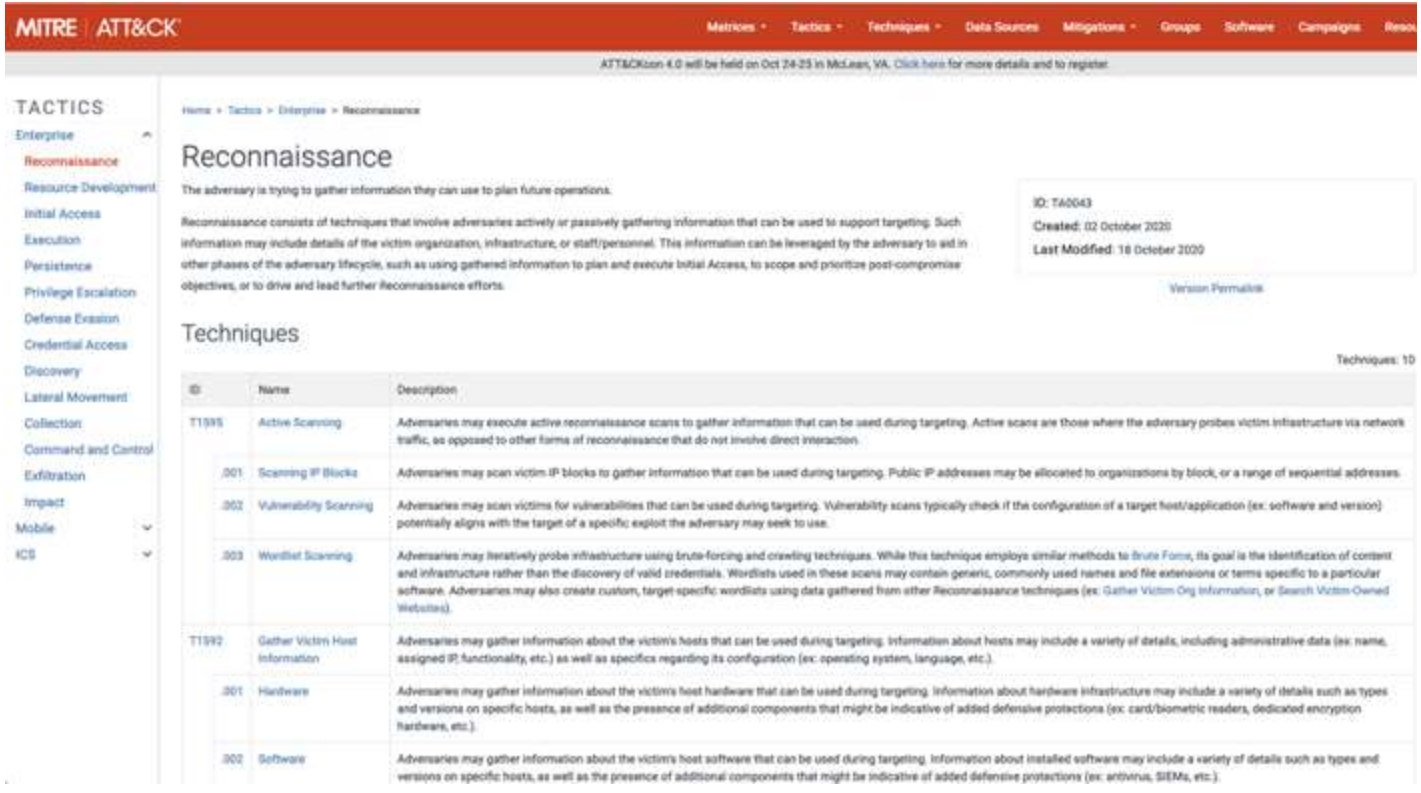

https://attack.mitre.org/tactics/TA0043/

10 specific techniques adversaries use for reconnaissance

Example: Reconnaissance: Active Scanning

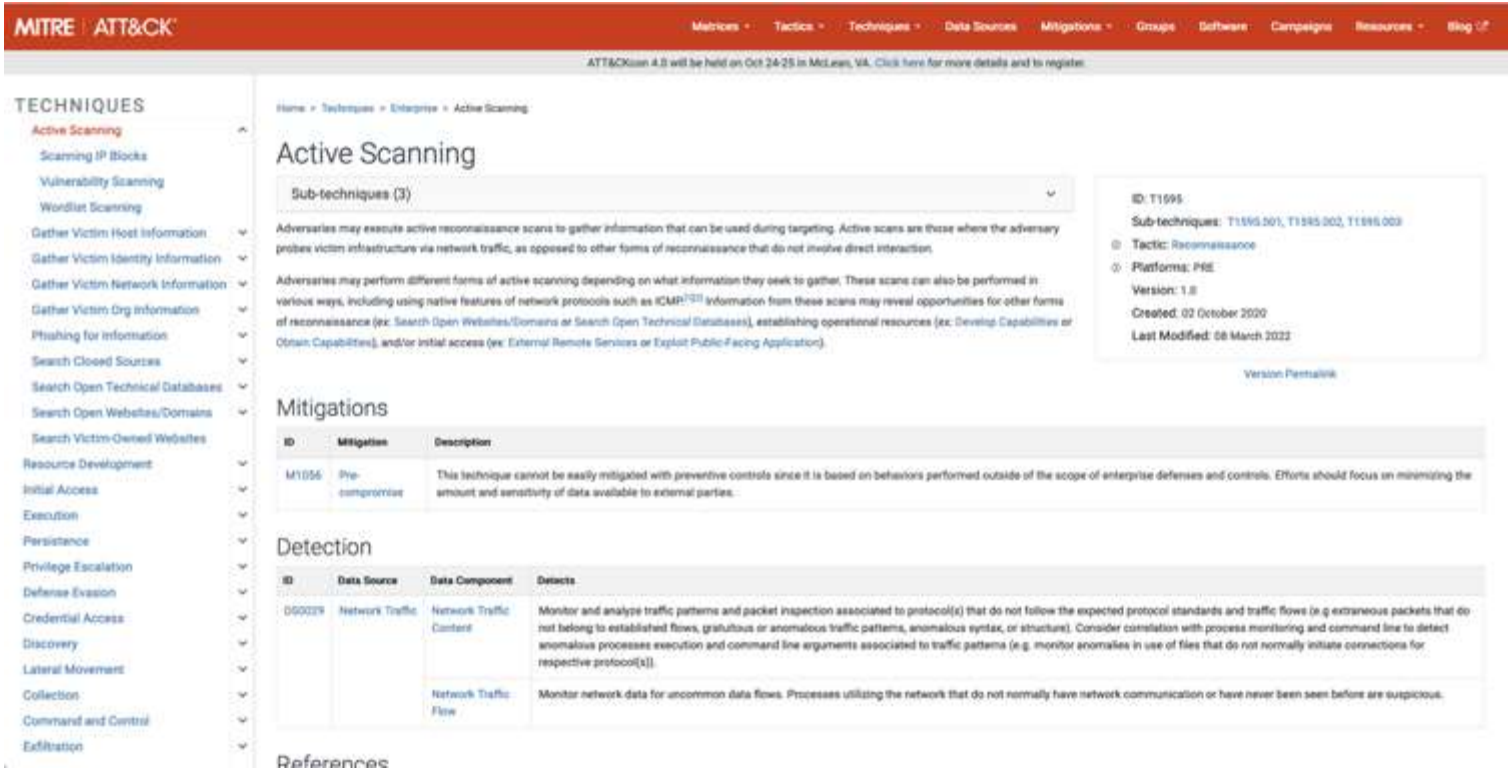

https://attack.mitre.org/tactics/TA0043/

Each technique has

- Description
- Mitigations
- Detection
- References
- Linkages to other techniques
- And more

Other Examples of How Teams Use MITRE Attack

Top 10 observed techniques in customer environments

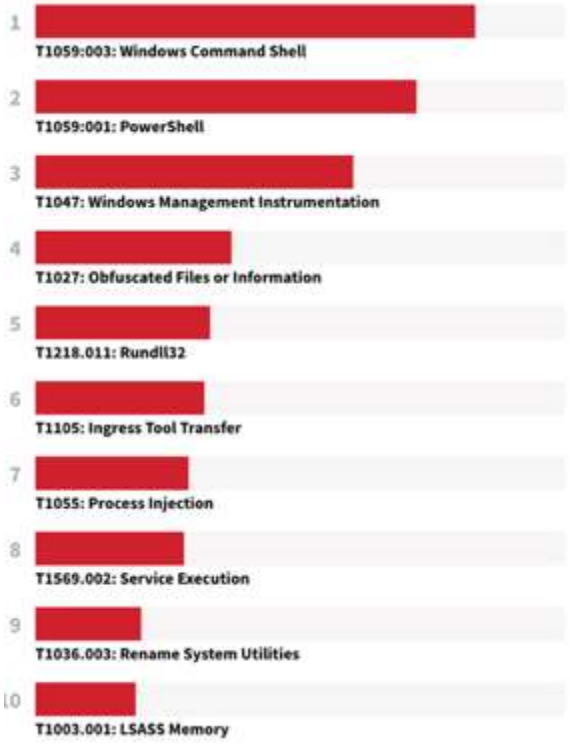

Red Canary 2023 Threat Detection Report

Red Canary uses it to report the *Most Common* threat activity seen in customer environments

How it is used:

"We have a library of roughly 3,500 detection analytics []. These are mapped to corresponding MITRE ATT&CK techniques whenever possible, allowing us to associate the behaviors that comprise a confirmed threat detection with the industry standard for classifying adversary activity.

More Examples Operationalizing MITRE ATTACK

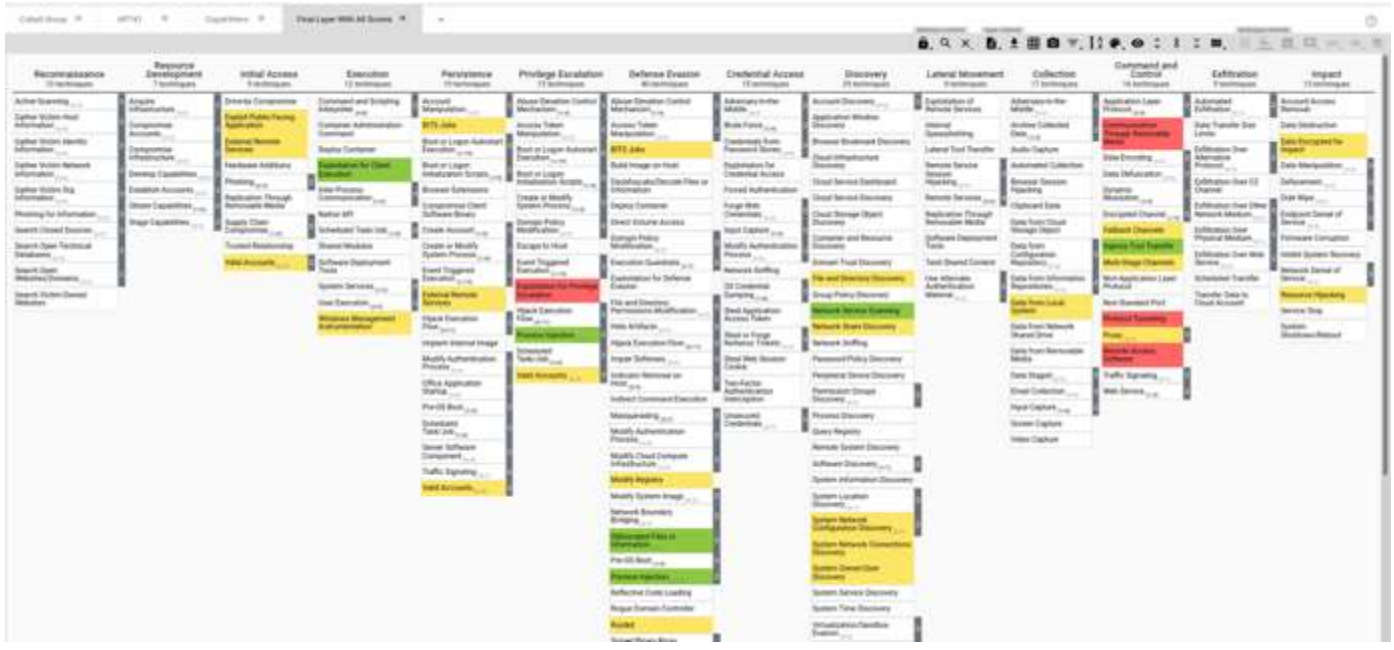

One team shows how they use Mitre Attack to display the results of their control testing and use scores as an overlay

https://www.signalblur.io/getting-started-with-mitres-att-ck-navigator

Heatmaps can show

- What controls you have / don't have
- What data you collect / are not yet collecting
- What areas attackers are observed to be in (vendors often report these)

Many Vendor Tools Integrate or Map to MITRE Attack

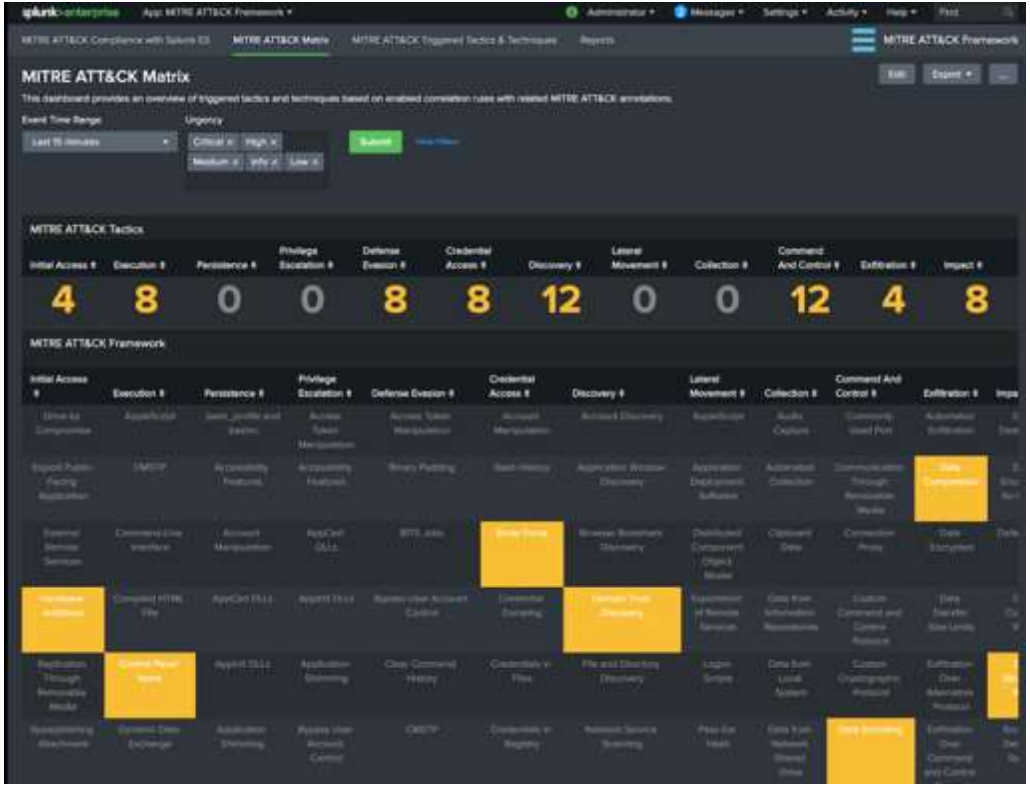

Example: Splunk Enterprise adopts MITRE ATT&CK - https://apps.splunk.com/app/4617/

MITRE ATTACK Framework

Carnegie Mellon University Software Engineering Institute

Advantages

- Common language / format
- Contains many adversary attacks and defenses against them
- Contains links to other resources for more detail
- Incorporated into many tools

Important to Remember

- Not quite exhaustive – has nearly all attacks but not fully 100%
- Learning curve for terminology

Pyramid of Pain

How expensive is it for attackers to change their approach if you block it?

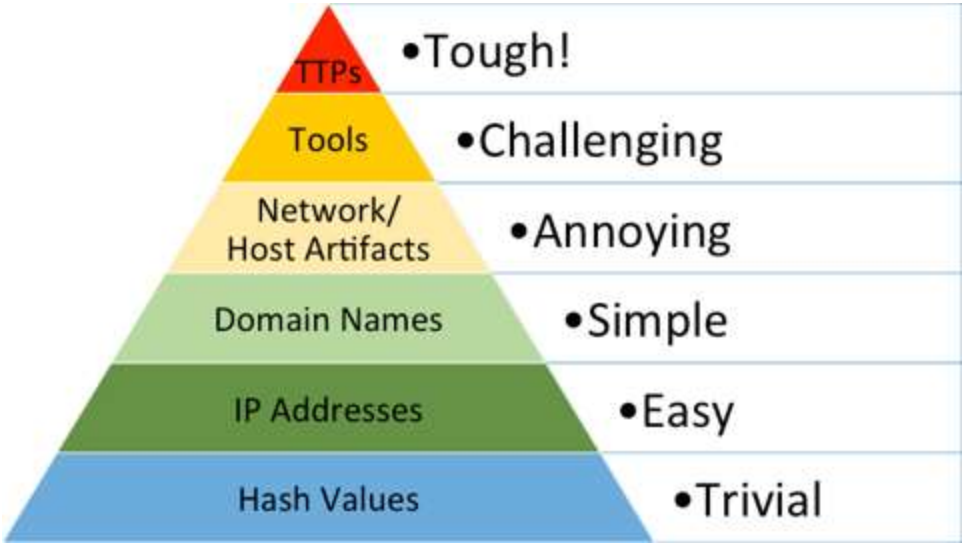

A cybersecurity professional named David Bianco created the pyramid of pain
http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html

Some types of intel lead to longer term defenses than others.

It is less useful to build a defensive capability around it.

So, things that are hard for attackers to change are better for defenders to use.

Also Useful: Sandia Generic Threat Matrix

Threat Level	THREAT PROFILE						
	Commitment			Resources			
					Knowledge		
	Intensity	Stealth	Time	Technical personnel	Cyber	Kinetic	Access
1	H	H	Years to decades	Hundreds	H	H	H
2	H	H	Years to decades	Tens of tens	M	H	M
3	H	H	Months to years	Tens of tens	H	M	M
4	M	H	Weeks to months	Tens	H	M	M
5	H	M	Weeks to months	Tens	M	M	M
6	M	M	Weeks to months	Ones	M	M	L
7	M	M	Months to years	Tens	L	L	L
8	L	L	Days to weeks	Ones	L	L	L

Reproduced from Duggan et al. [8].

Trevino, Cassandra M., et al. *Cyber threat metrics*. Sandia National Laboratories, 2012.

Summary

Threat hunting Frameworks are useful at standardizing terminology and ideas and for sharing a common picture with each other.

Popular frameworks include:

- Lockheed Martin Kill Chain
- MITRE ATT&CK
- Pyramid of Pain
- Sandia generic threat matrix
- And many more

Frameworks are not standalone; they are frequently combined with each other.

Many vendor tools incorporate these frameworks to improve internal reporting and sharing of information with other teams.

Reading the Threat Landscape

Carnegie Mellon University
Software Engineering Institute

Threat Landscape and Sources

Threats can be

- external
- internal

Need to learn trends and patterns for both types

- current
  - passive
  - active
- general TTPs
- specific TTPs – reported by other teams
- sources of threat behavior
- identifying ‘shifts’

Data Gathering – Monitor Reported Threat Activity

Read what others are saying.

Create trends from your own internal (incident/alert) data.

Gather your own data on external threats (see Honeypots).

Read carefully – be wary of data collection, methodology, sample sizes, false positives, date ranges, and resulting claims.

Data Gathering – Sources of Threat Information

Stories

Vendor threat trend reports and data feeds

- events from their sensors
- events from clients they are servicing

Internet databases (may be sponsored by vendors)

Academic studies

Information Security Analysis Center (ISAC) reports

Sharing networks

And more

For all sources, carefully read descriptions of methodology, sample size, and resulting claims.

Not All Sources of Data for Threat Hunting Are Equal

Different sources of threat intelligence will give you different levels of information and different levels of confidence.

Some are more useful than others.

Context matters.

Anecdotes can still be very useful.

Be aware of sample sizes, the origin of results, the methods of collection, statistical analysis methods, and more.

What Are Sharing Networks?

A community of defensive organizations sharing information about what they are seeing on their networks

Often use sharing rules such as “traffic light protocol”

Networks of all types exist: Open to all vs Invite only, Sector specific, Attributed vs. non-attributed, and more…

Tools now exist to set up your own private threat sharing network and invite other organizations to join.

Threat hunters can send and answer queries and mine networks for trends.

Government Reports

CYBERSECURITY & INFRASTRUCTURE SECURITY AGENCY — AMERICA'S CYBER DEFENSE AGENCY

Search

Topics | Spotlight | Resources & Tools | News & Events | Careers | About

Home / News & Events / Cybersecurity Advisories / Cybersecurity Advisory

CYBERSECURITY ADVISORY

Iranian Government-Sponsored APT Actors Compromise Federal Network, Deploy Crypto Miner, Credential Harvester

Last Revised: November 25, 2022 **Alert Code:** AA22-320A

Summary

From mid-June through mid-July 2022, CISA conducted an incident response engagement at a Federal Civilian Executive Branch (FCEB) organization where CISA observed suspected advanced persistent threat (APT) activity. In the course of incident response activities, CISA determined that cyber threat actors exploited the Log4Shell vulnerability in an unpatched VMware Horizon server, installed XMRig crypto mining software, moved laterally to the domain controller (DC), compromised credentials, and then implanted Ngrok reverse proxies on several hosts to maintain persistence. CISA and the Federal Bureau of Investigation (FBI) assess that the FCEB network was compromised by Iranian government-sponsored APT actors.

https://www.cisa.gov/news-events/cybersecurity-advisories/aa22-320a

Advisories contain

- PDF
- Mappings to MITRE ATT&CK
- Context
- Separate files for IOCs and STIX

Government Automated Sharing: CISA AIS Program

Automated Indicator Sharing

OVERVIEW

The Department of Homeland Security (DHS) Cybersecurity and Infrastructure Security Agency's (CISA's) free Automated Indicator Sharing (AIS) program enables organizations to share and receive machine-readable cyber threat indicators (CTIs) and defensive measures (DMs) in real time to monitor and defend their networks against known threats that are relevant to AIS participants.

WHY PARTICIPATE IN AIS?

By participating in AIS, organizations can send and receive CTIs/DMs with other organizations and can be on the lookout for similar activity to proactively defend their network. This allows organizations to benefit from the collective knowledge of participant organizations. AIS also offers anonymity, as well as liability, and privacy protections to encourage the submission of CTIs/DMs related to successful or attempted compromises.

THE CYBERSECURITY INFORMATION SHARING ACT OF 2015

AIS is available through CISA's Cybersecurity Division and CISA Central which are designated as the hub for the sharing of CTIs/DMs between the federal government and private sector by the Cybersecurity Information Sharing Act of 2015. This law grants liability protection, privacy protections, and other protections to organizations that share CTIs/DMs through AIS in accordance with the Act's requirements. As mandated by the Cybersecurity Information Sharing Act of 2015, DHS certified the operation of AIS in March 2016. The goal is to share tactical CTIs/DMs through AIS broadly among the public and private sector, enabling everyone to be better protected against cyberattacks.

LIABILITY PROTECTION

Liability protection is granted to organizations for sharing through AIS if the sharing of CTIs/DMs is done in accordance with the Cybersecurity Information Sharing Act of 2015. Liability protection applies to:

- Non-federal entities sharing with other non-federal entities;
- Non-federal organizations sharing with information sharing and analysis centers (ISACs) and information sharing and analysis organizations (ISAOs);
- Non-federal entities sharing with CISA and other federal agencies through AIS.[1]

Federal organizations do not receive liability protection when sharing with one another, but some aspects of the Cybersecurity Information Sharing Act of 2015 apply (e.g. privacy requirements when sharing CTIs).

PRIVACY PROTECTIONS

CISA has taken careful measures to ensure appropriate privacy and civil liberties protections are fully implemented in AIS. CISA has published a privacy impact assessment of AIS found on https://www.cisa.gov/automated-indicator-sharing-ais.

The Department of Homeland Security (DHS) Cybersecurity and Infrastructure Security Agency's (CISA's) free Automated Indicator Sharing (AIS) program enables

- organizations to share and receive machine-readable cyber threat indicators (CTIs) and defensive measures (DMs)
- real time to monitor and defend their networks against known threats that are relevant to AIS participants

Example: Recorded Future Global Trends

List of Vendor Threat Report Examples

- Verizon: Data Breach Investigations Report
- Websense Threat Report
- Symantec: Internet Security Threat Report
- Sophos: Security Threat Report, Cisco's Annual Security Report
- Hewlett Packard: Cyber risk report
- EY: Under Cyber Attack – EY's Global Information Security Survey
- Booz Allen: Cyber Power Index
- Office of Management and BudArbor Annual Threat Report focused on Denial of Service
- Ponemon Institute Exposing the Cybersecurity Cracks: A Global Perspective
- CSRIC IV WG5 "Remediation of Server-Based DDoS Attacks" Final Report
- Guide to Cyber Threat Information Sharing (Draft), NIST Special Publication 800-150 (Draft)
- Annual Report to Congress – Federal Information Security Management

Source: CYBERSECURITY RISK MANAGEMENT AND BEST PRACTICES WORKING GROUP 4: Final Report

Research Studies

Academic and private sector research offers longer term studies, publish their collection methods, and sometimes their data.

Examples:

- *Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem* by Kyle Soska and Nicolas Christin
  - In Proceedings of the 24th USENIX Security Symposium (USENIX Security'15), pages 33-48. Washington, DC. August 2015.
- *Framing Dependencies Introduced by Underground Commoditization* by Kurt Thomas et al
  - Workshop on the Economics of Information Security, 2015

Example lesson learned: **Attackers can outsource parts of their attack against a target for relatively low cost and very high amounts of specialization.**

Platform Versus Protocol

In the next few sections, we are going to discuss various platforms and protocols for information sharing.

Platforms are software applications built to facilitate communication of a variety of data types.

- Examples for general platforms would be Twitter, Facebook, MISP, etc.

Protocols in Information Sharing have the same meaning as protocols in networking.

- Protocols are the rules for how to send data to another organization, person, or application.
- For example, TAXII is a protocol for sending Cyber Threat Intelligence data to another application, STIX, using the HTTPS protocol.

Sometimes an information sharing platform has its own protocol; so MISP is a platform and a protocol.

Information Sharing Platforms

A platform may perform a variety of security tasks to help defenders such as:

- Receive data about threats from other network participants or publicly available data sources
- Help analysts perform correlation analysis between events such as ‘linking’ them together
- Allow analysts to add metadata to values such as URLs, Domain names, Filenames, Hash values, and much more
- Integrate with other services such as malware sandbox analysis and importing results (enrichment)
- Integrate with defensive tools including Firewalls, Intrusion Detection Systems (IDS), SIEM, or other programmable network/host event sensors such as Zeek (formerly Bro)

Sharing Platforms

Only as valuable as the data they contain

- All of the platform use cases require the receipt of useful, timely, and ‘actionable’ security data.
- Data that is out of date or incorrect can cause unnecessary outages rather than prevent against attacks.
- Using incorrect and inaccurate cyber intelligence can lead to many other security failures.
- There are many teams that provide data for others to use (often called Feeds), but careful examination of each dataset is recommended.

Example Platform: MISP

Carnegie Mellon University Software Engineering Institute

Key Features

- Store, share, collaborate on cyber security indicators, malware analysis, and use to detect and prevent attacks or threats
- Support for Events to have tags, to apply different taxonomies
- Multi-layered Sharing groups for multiple organizations with permissions and protocols (including TLP)
- Import/Export events in various formats including indicator extraction via Regex
- Linking of attributes (observables and IOCs) between events

MISP – Event List

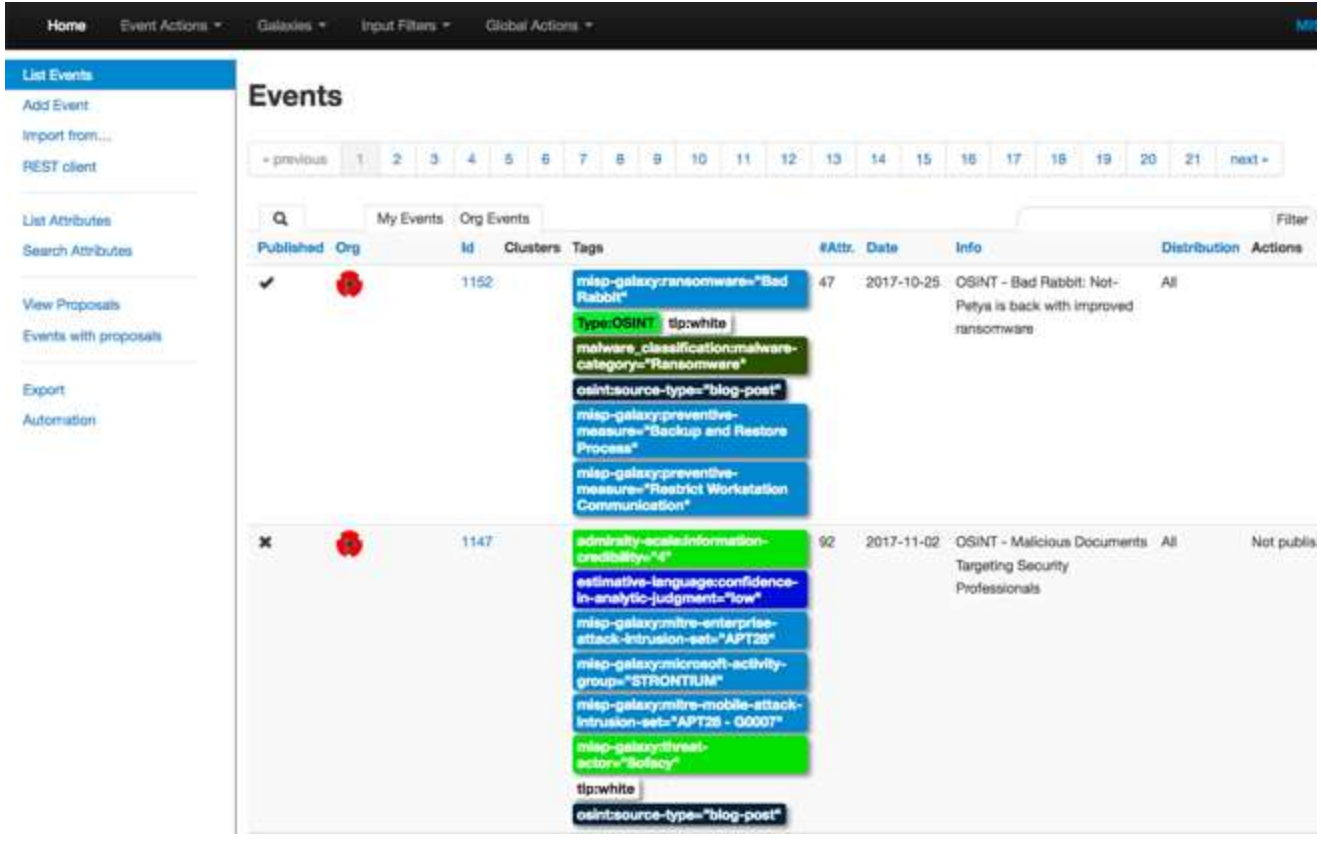

Events from the the CIRCL MISP dataset imported into a MISP server. tlp:white

Sample Design Plan

open-src tool integration and data flow

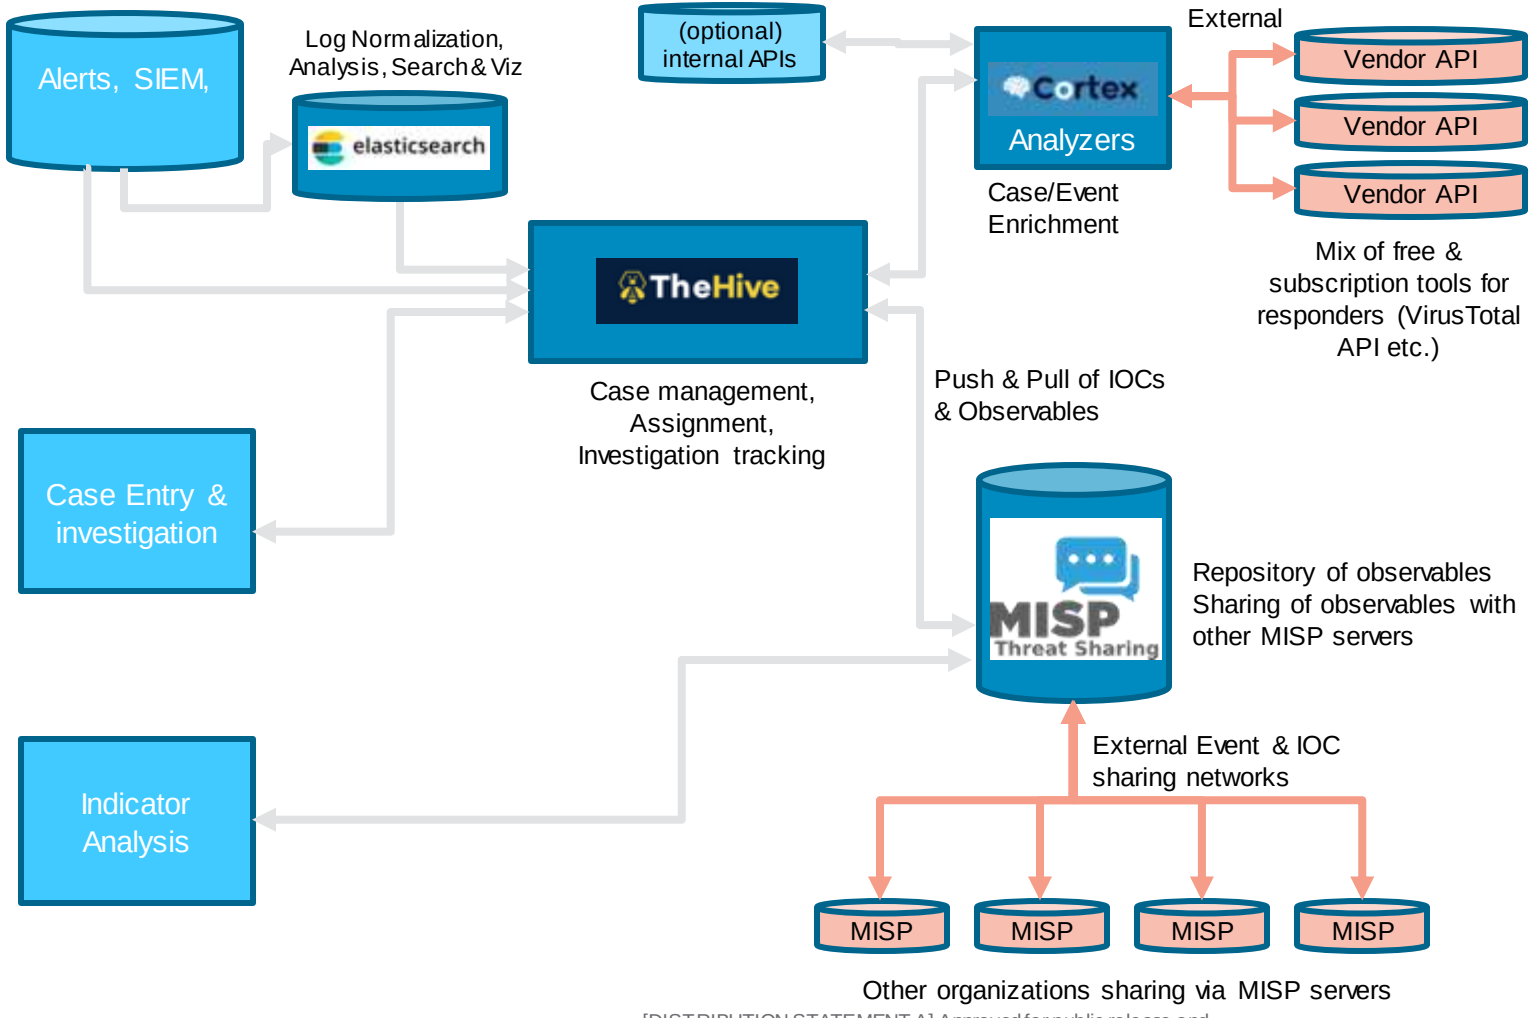

Structured Threat Information Expression (STIX™)

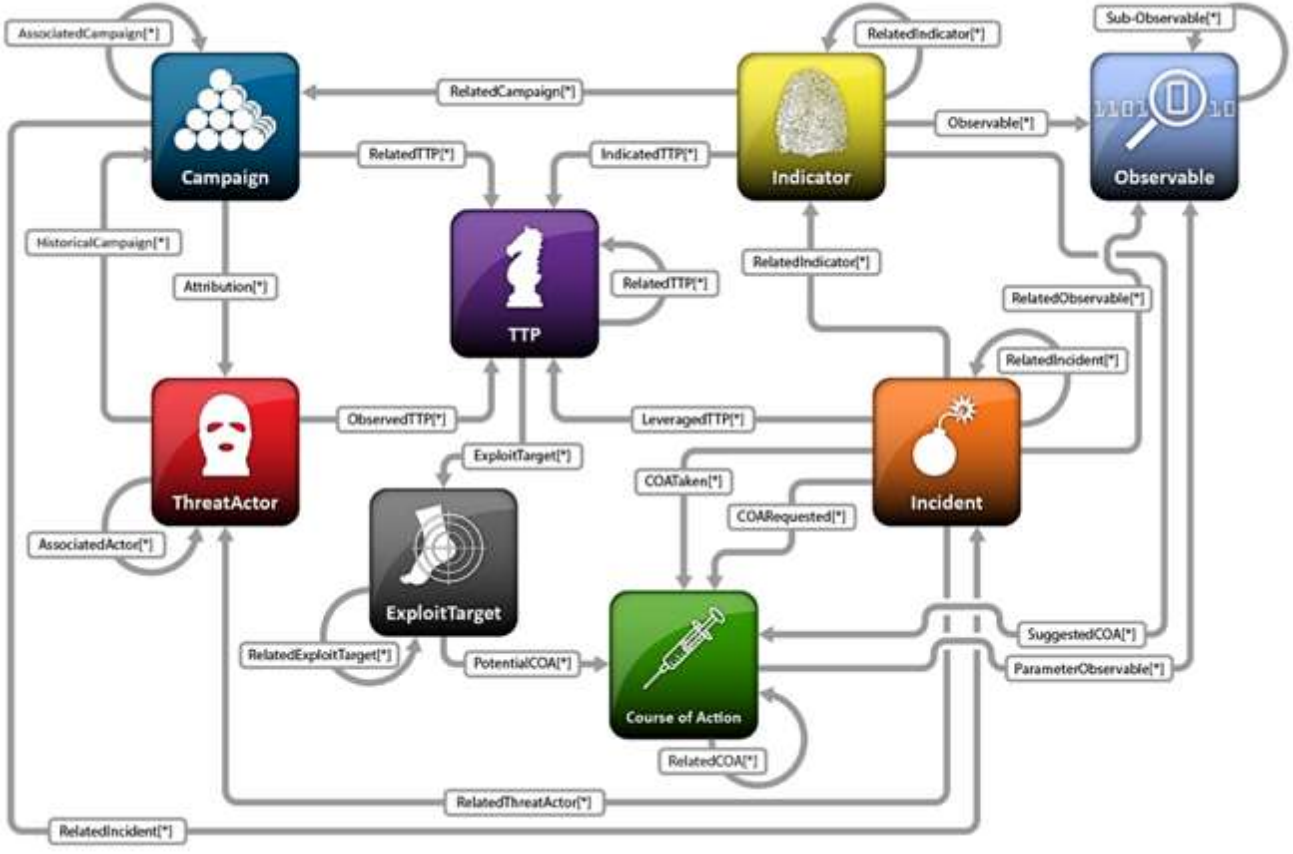

Image Source: http://stixproject.github.io/about/

- Structured language for describing cyber threat information
- Easier to share, store, analyze in a consistent manner.
- STIX – Maintained by OASIS
- TAXII – A transport protocol for STIX

Example STIX Data Model of a TTP

Scenario represents 3 IP addresses that are 'known' C2 for an adversary

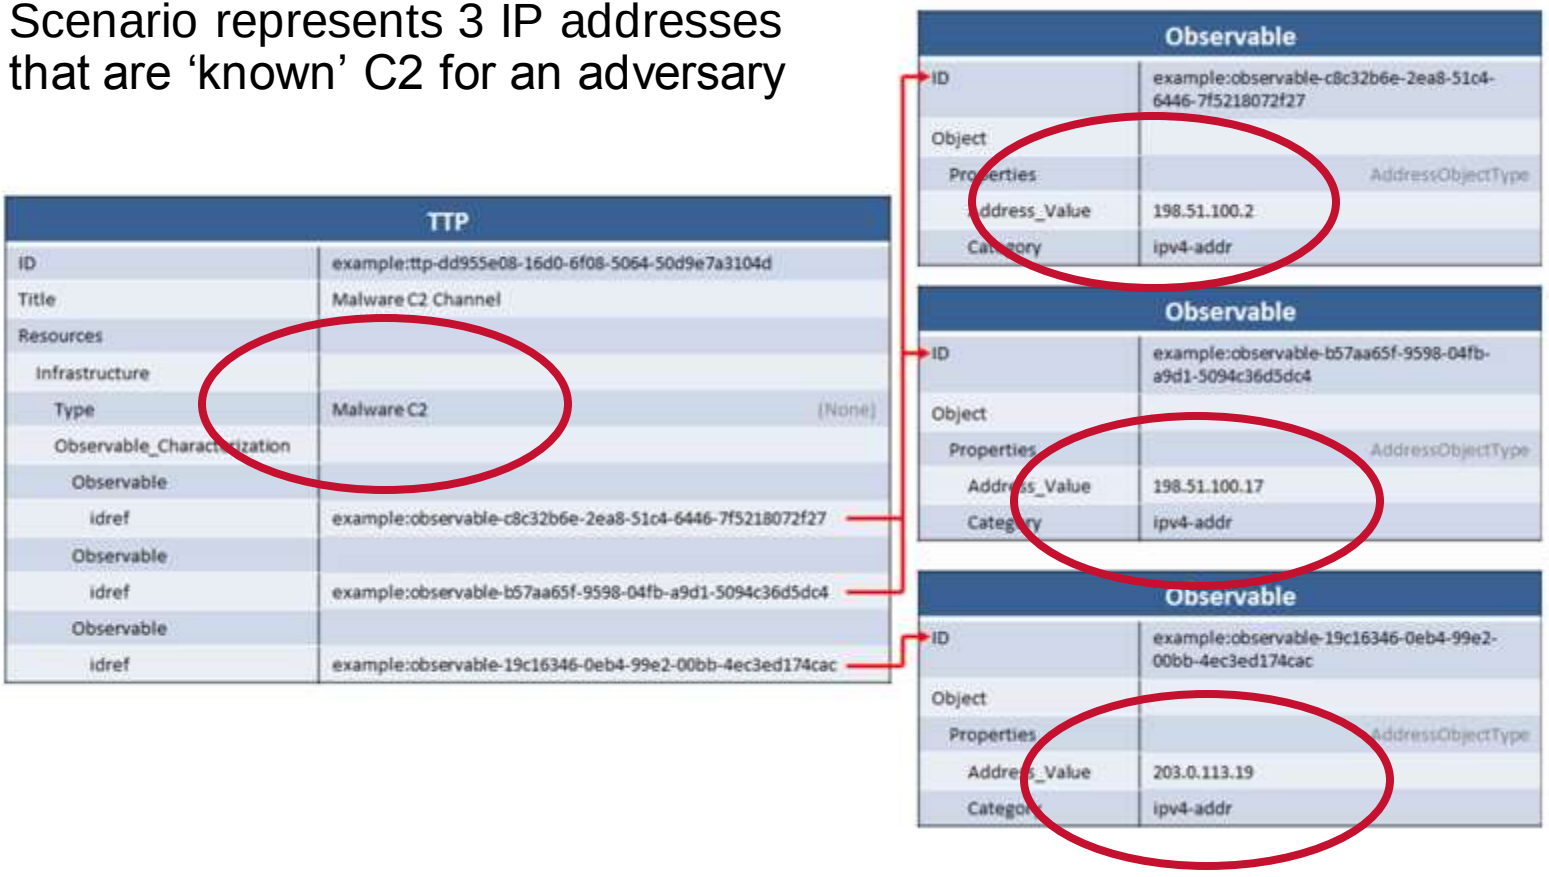

Anecdotal Sources Can Be Useful

Sources to learn about threats may be anecdotal, such as news stories:

- CNN – “5-year-old boy hacks dad's Xbox account” (http://www.cnn.com/2014/04/04/tech/gaming-gadgets/5-year-old-xbox-hack/)
- Zdnet – “Teenager hacks Google Chrome with three 0day vulnerabilities” (http://arstechnica.com/security/2012/10/google-chrome-exploit-fetches-pinkie-pie-60000-hacking-prize/ and http://www.zdnet.com/article/teenager-hacks-google-chrome-with-three-0day-vulnerabilities/)
  - It took about one-and-a-half weeks to find the vulnerabilities and write a reliable exploit.
  - The exploit worked on a fully patched Windows 7 machine (64-bit) and did not require user action beyond normal web browsing.

Summary

Threat Hunting Analysis is a process performed by analysis with loops, changes, and exploration to discover a reasonable and repeatable answer.

Trends and patterns can help uncover unusual circumstances and highlight the need for action.

Many sources and networks exist to discuss, compare, and contrast threat actor and group behavior.

Data Sources include news, reports, databases, and research papers.

For vendor reports, read and fully understand the methodologies and sample sizes before you make judgements about the claims and findings in the reports; be a skeptic.

Academic and private sector research offers longer term studies with published collection methods.

There is growth in threat information sharing using machine readable formats.

Hunt Teams, Analysts, Maturity Models

Who Should Be on the Hunt Team?

What skills would help hunters answer each question?

- Where are all the places we should be hunting?
- How can we automate the hunt?
- How prevalent is this problem for us?
- How do we explain the problem?
- What are all the likely ways attackers might harm us?

Who Should Be on the Hunt Team?

Role	***Sample Activities***
Network & IT Infrastructure	• Collect logs from IT assets • Provide knowledge and background on “normal” activity vs. anomaly
Security Experts	• Understand security architecture and controls • Understand weaknesses • Identify security significant impacts in data
Programmers/Developers	• Help automate collection and analysis tasks • Improve tools and platforms
Data Scientists	• Work with large datasets • Apply Machine Learning & Deep Learning to data • Interpret results
Visualization & Communications	• Develop effective charts and graphs • Improve communication and actionability of technical findings
Business	• As needed, provide expertise on business processes
Management	• Plan and track tasks, report Metric • Align hunting with priorities of organization • Financial management, approve tools acquisition, hiring etc.

The Analysts Mind

Example slide from "The Mind of a Hunter: A Cognitive, Data-Driven Approach - SANS Threat Hunting Summit 2017" by Chris Sanders

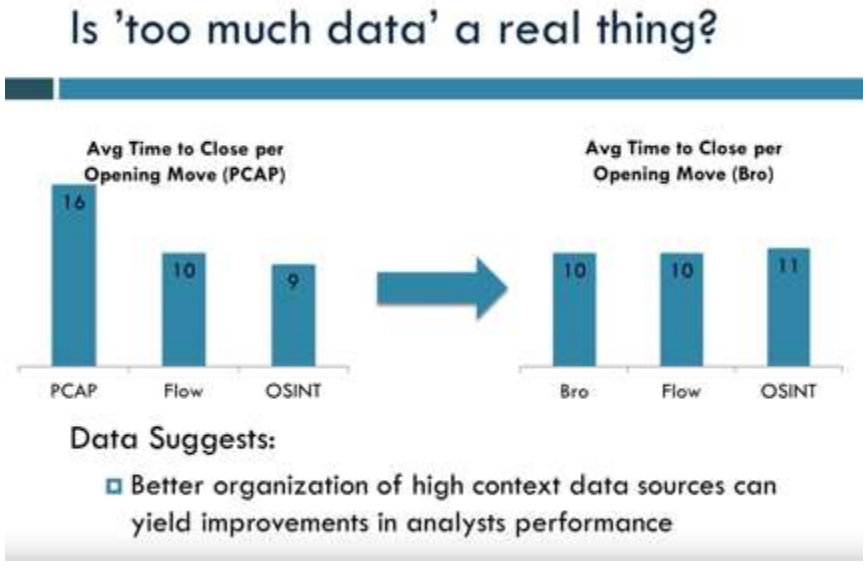

How should analysts perform threat hunting?

Chris Sanders studied the time it took for analysts to achieve a desired outcome depending upon which cybersecurity data source they started with.

For more, see the video at *The Mind of a Hunter: A Cognitive, Data-Driven Approach - SANS Threat Hunting Summit 2017* and Chris Sander's doctoral thesis: *The Analyst Mindset: A Cognitive Skills Assessment of Digital Forensic Analysts.* He also offers a course on **Investigation Theory**

Analyst Techniques

Establishing Baselines and searching for Anomalies in different domains.

Attempting to scale searching, expansion, pattern matching, and correlation activities.

May need to apply techniques from other fields including statistical models, social sciences (language), human computer interaction, and more.

Awareness of Business Events, Calendar, Cyclical Behaviors, volumes, and more.

See Additional Resources handout and slide Notes for links to more specific Data Analysis techniques

- Depending upon the data type: Network, Server, Endpoint, Appliance, etc.
- Or for Statistical Analysis / Machine Learning including Deep Learning.

SEI Survey of Non-signature-based Hunting Techniques

By George Jones, John Stogoski

- Analysts focus on a few common protocols
- They apply understanding of related business processes
- They look for expected behaviors.
- Excluding certain data highlights interesting things for additional examination.
- Hunting results in the creation of new repeatable processes to look for suspicious artifacts (including IPs, domain names, certificates, and others).
- New signatures minimizes analysts' future workloads
- Note this was in 2014….

SQRRL Hunting Maturity Model (HMM)

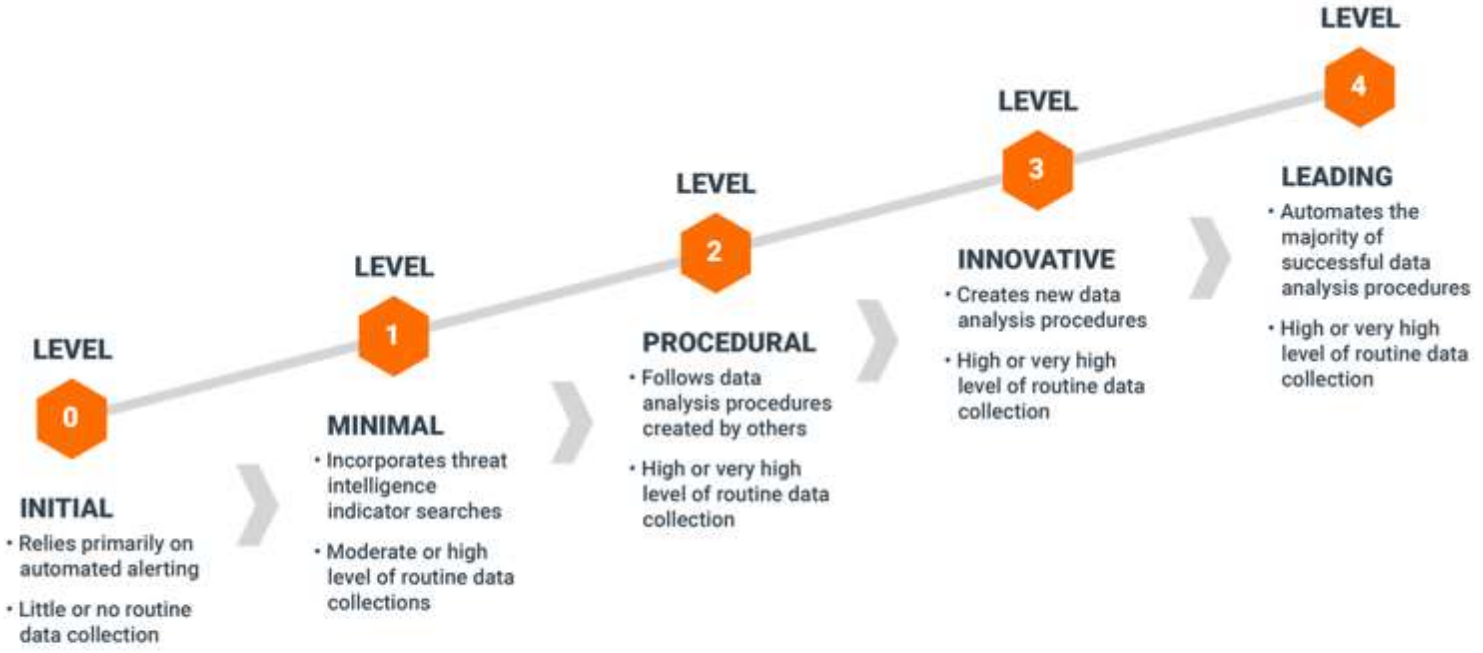

The Hunting Maturity Model (HMM)

Author: David J. Bianco SQRRL

Implications of AI Technologies on Threat Hunting

Carnegie Mellon University
Software Engineering Institute

Markets
The Big Take

Deepfake Imposter Scams Are Driving a New Wave of Fraud

AI could turbocharge the cybertheft economy. The world's banking industry is scrambling to contain the risk.

Illustration: Jinhwa Jang for Bloomberg Markets

By Nabila Ahmed, Adam Haigh, Ainsley Thomson, and Ellie Harmsworth
August 21, 2023 at 7:00 PM EDT

Computer-generated children's voices so realistic they fool their own parents. Masks created with photos from social media that can

3 ways ChatGPT can help criminals take advantage of you

By Chris Smith Published Mar 28th, 2023 10:18AM EDT

Image: phonlamaiphoto/Adobe

European Union's police force warned us that malicious individuals can use ChatGPT to assist with various criminal activities.

1. Generate text that reads just like a regular message from one of those companies
2. Generate a 'specific narrative' with little effort – helps fraudsters
3. Produce malicious code

Chat Bots for Attackers?

WormGPT Is a ChatGPT Alternative With 'No Ethical Boundaries or Limitations'

The developer of WormGPT is selling access to the chatbot, which can help hackers create malware and phishing attacks, according to email security provider SlashNext.

By Michael Kan July 14, 2023

(Credit: Hacking forum)

Used in Business Email Compromise (BEC) attacks to generate messages

Trained on Malicious code

Guardrails removed – will generate malicious executable code if asked

> *"The results were unsettling. WormGPT produced an email that was not only* ***remarkably persuasive*** *but also strategically cunning, showcasing its potential for sophisticated phishing and BEC attacks,"* SlashNext said.

Source: https://www.pcmag.com/news/wormgpt-is-a-chatgpt-alternative-with-no-ethical-boundaries-or-limitations

Other Attacks on AI technology

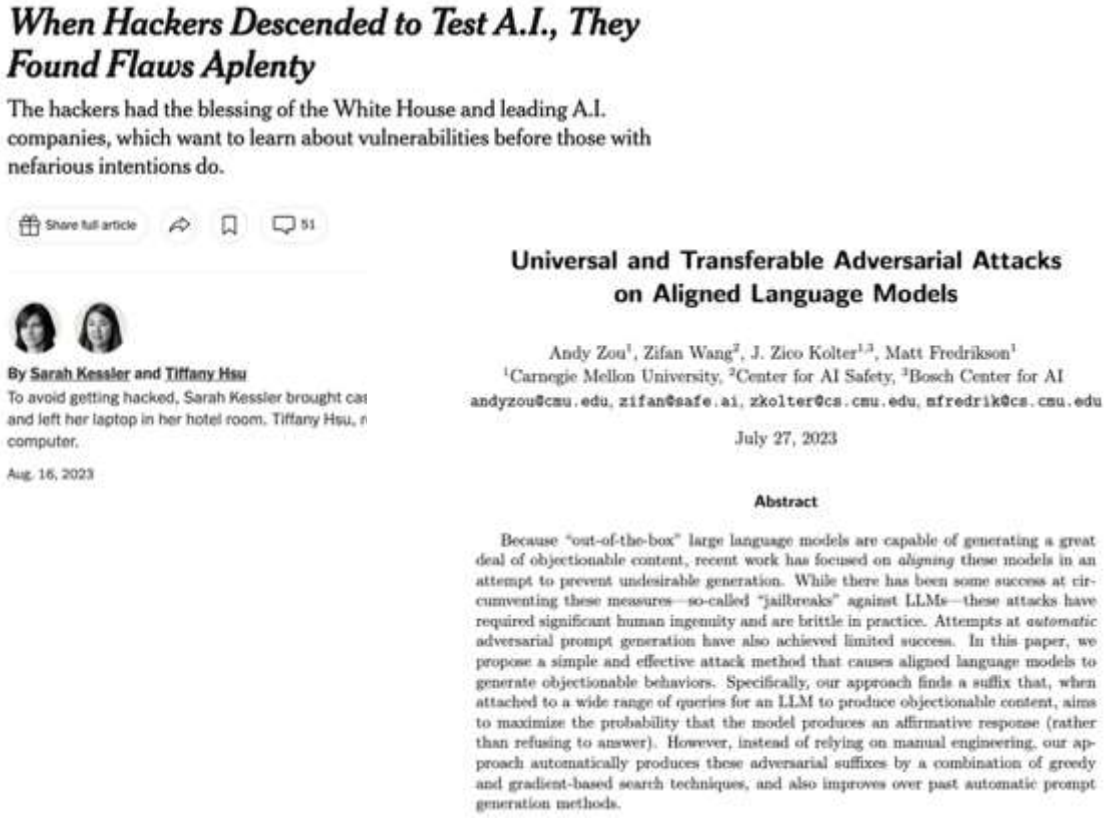

When Hackers Descended to Test A.I., They Found Flaws Aplenty

The hackers had the blessing of the White House and leading A.I. companies, which want to learn about vulnerabilities before those with nefarious intentions do.

Share full article 51

By Sarah Kessler and Tiffany Hsu

To avoid getting hacked, Sarah Kessler brought cas and left her laptop in her hotel room. Tiffany Hsu, r computer.

Aug. 16, 2023

Universal and Transferable Adversarial Attacks on Aligned Language Models

Andy Zou[1], Zifan Wang[2], J. Zico Kolter[1,3], Matt Fredrikson[1]
[1]Carnegie Mellon University, [2]Center for AI Safety, [3]Bosch Center for AI
andyzou@cmu.edu, zifan@safe.ai, zkolter@cs.cmu.edu, mfredrik@cs.cmu.edu

July 27, 2023

Abstract

Because "out-of-the-box" large language models are capable of generating a great deal of objectionable content, recent work has focused on *aligning* these models in an attempt to prevent undesirable generation. While there has been some success at circumventing these measures—so-called "jailbreaks" against LLMs—these attacks have required significant human ingenuity and are brittle in practice. Attempts at *automatic* adversarial prompt generation have also achieved limited success. In this paper, we propose a simple and effective attack method that causes aligned language models to generate objectionable behaviors. Specifically, our approach finds a suffix that, when attached to a wide range of queries for an LLM to produce objectionable content, aims to maximize the probability that the model produces an affirmative response (rather than refusing to answer). However, instead of relying on manual engineering, our approach automatically produces these adversarial suffixes by a combination of greedy and gradient-based search techniques, and also improves over past automatic prompt generation methods.

Techniques Include

- Model poisoning
- Prompt Injections
- Prompt Engineering

"attack suffix is able to induce objection-able content in the public interfaces to ChatGPT, Bard, and Claude, as well as open source LLMs such as LLaMA-2-Chat, Pythia, Falcon, and others. "

Increased Use of AI Tools in Cyber Defense

Automating Attack and Defense – see Darpa Cyber Grand Challenge

Create and manage IT Infrastructure

Pair Programming via AI assistant – Ex. Co-Pilot

Automated Patch generation and deployment

And more

Increased Use of AI Tools in Threat Hunting

SentinalOne incorporating Generative AI into its platform allowing for

- Natural language translated into syntax queries or commands
- Result summarization and ‘storyline generation’
- Classification of technical event activity into industry standard terminology “This is lateral movement”

Microsoft Security Copilot

- GPT-4 Assistant for cybersecurity analysts
- Summarize incidents, events, and reporting
- Accepts mixed data types: natural language, URLs, code/log snippets
- Collaboration tools (pinning prompts, reuse, team tools)

Copilot | Dynamics 365 and Power Platform

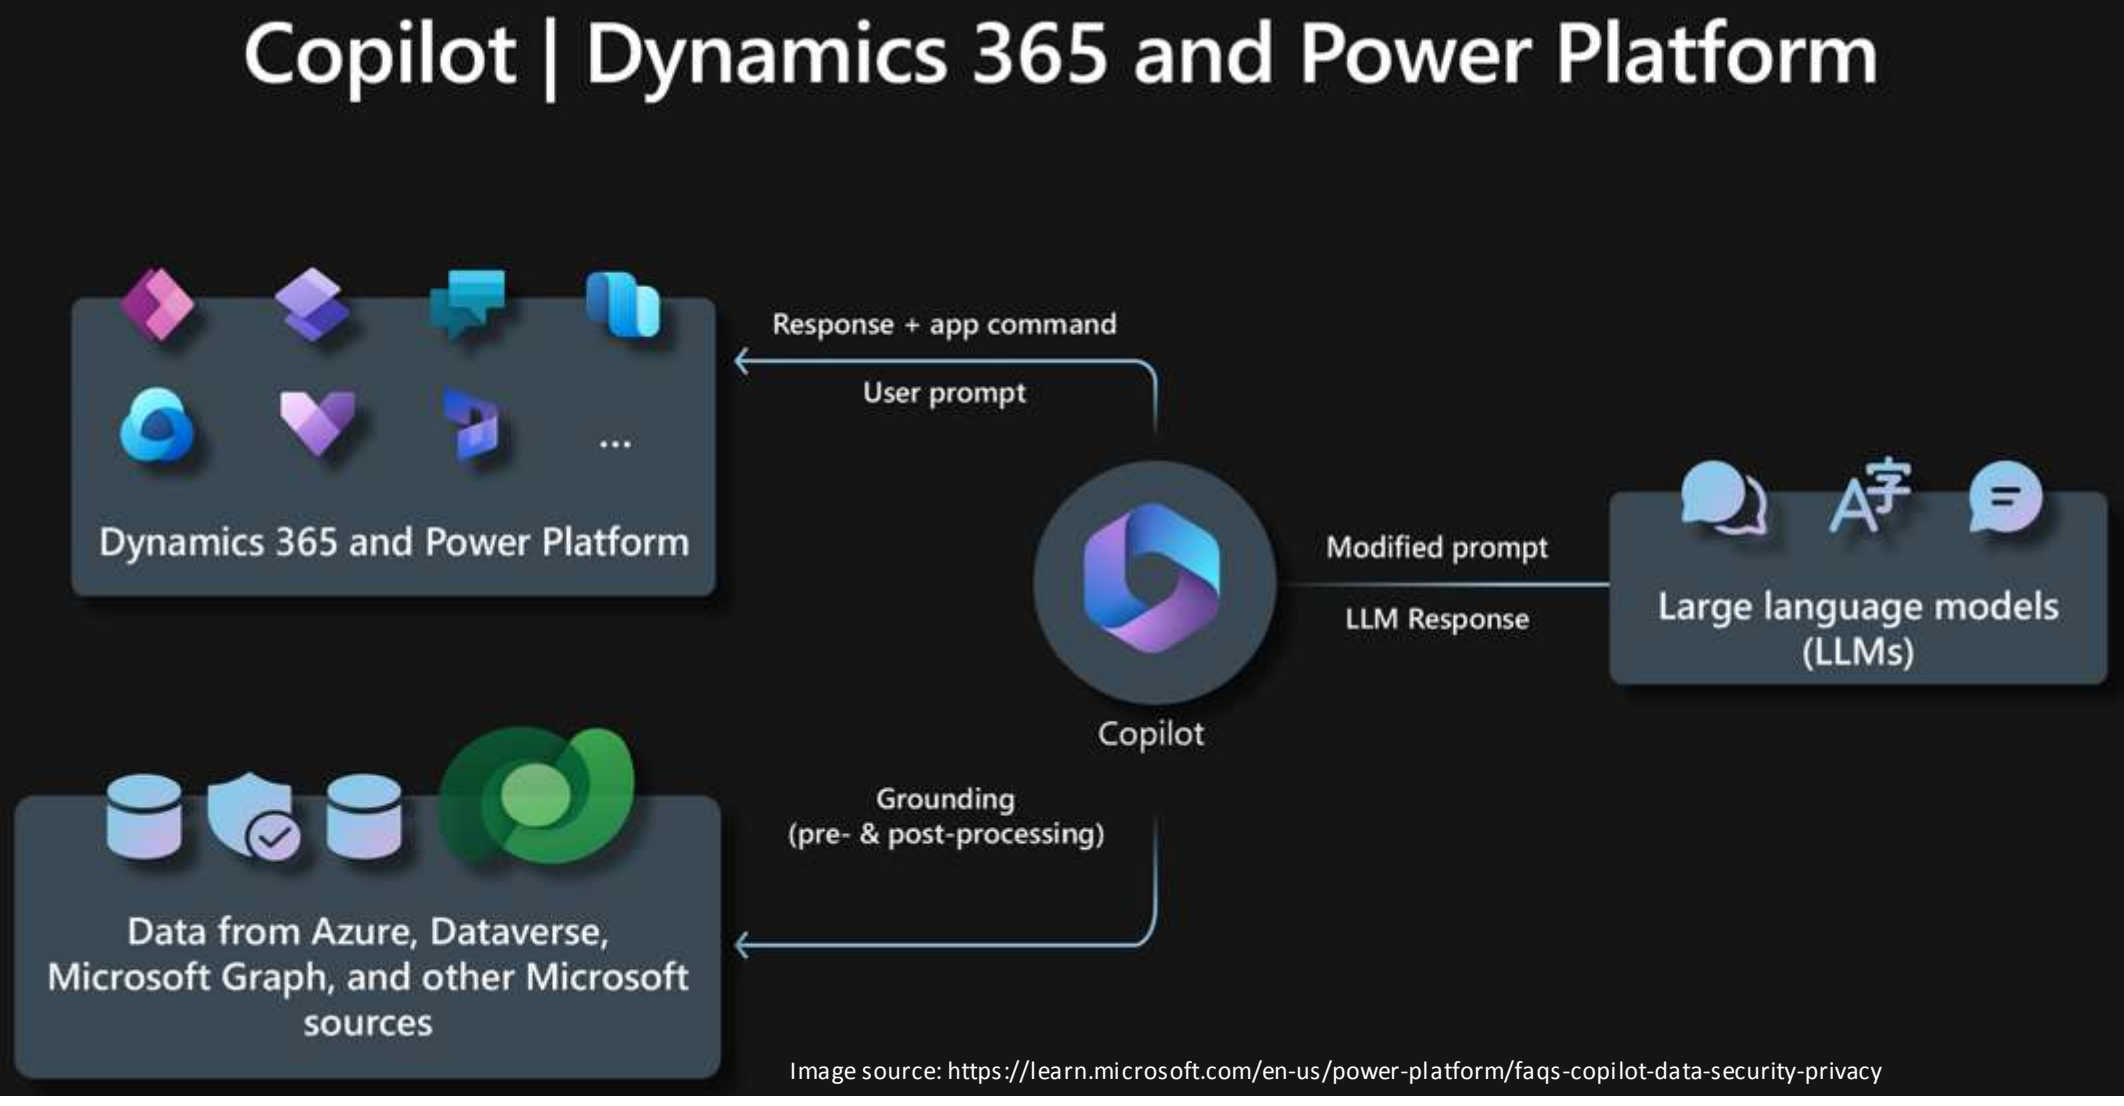

Image source: https://learn.microsoft.com/en-us/power-platform/faqs-copilot-data-security-privacy

Example Incorporating LLMs into Tools for Threat Hunters

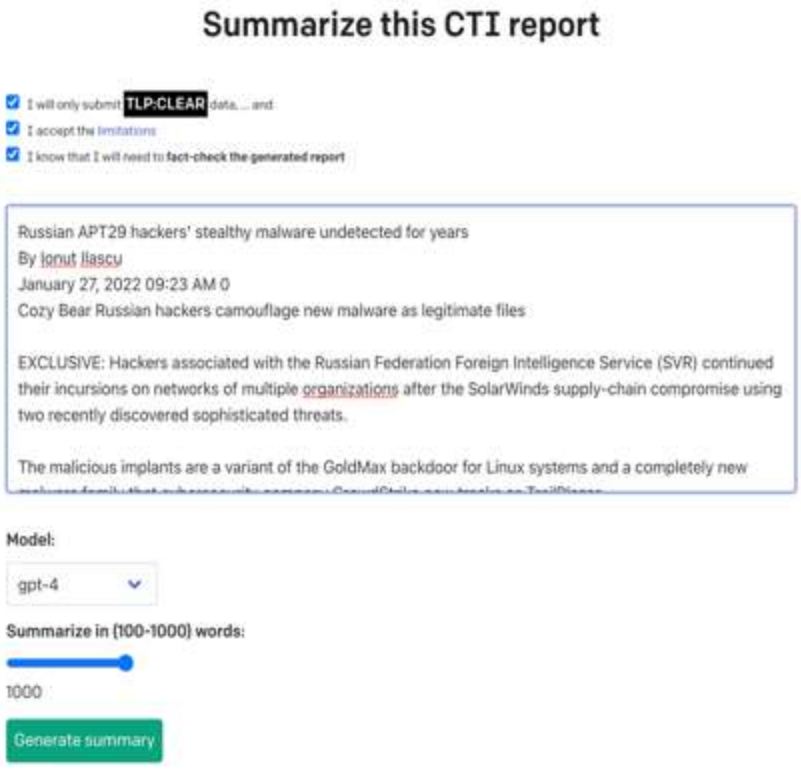

Output

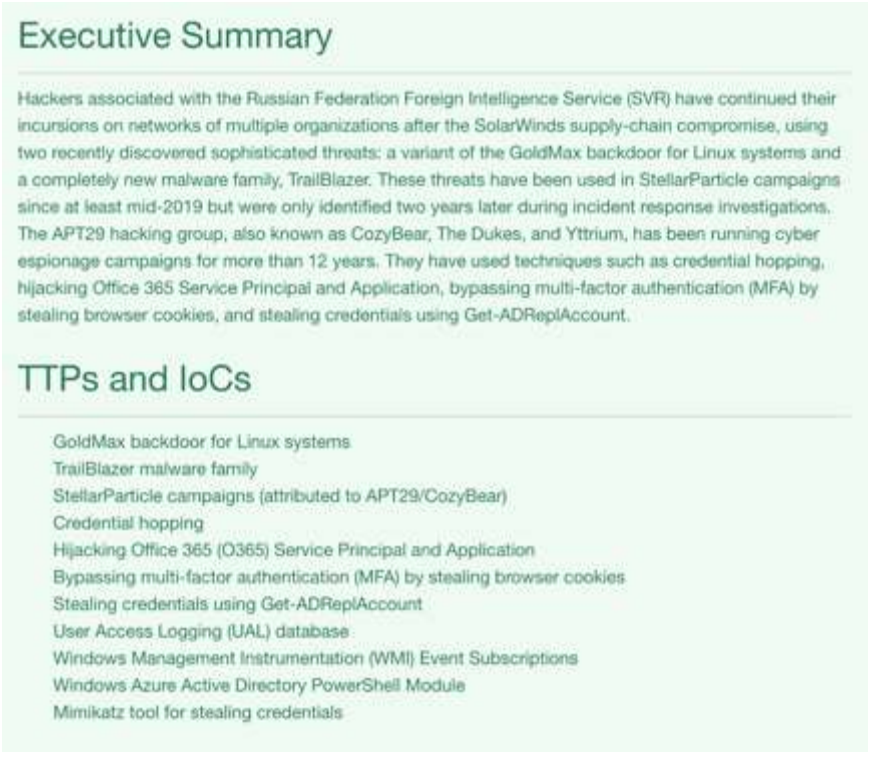

https://github.com/EC-DIGIT-CSIRC/openai-cti-summarizer by Aaron Kaplan

Trend Micro White Paper

Malicious Abuses of AI

- Deepfakes and Human Impersonation
- AI-Supported Password Guessing
  - Some attackers trained a model on how humans change their passwords
  - Example: April to Apr1l to A9r!l over time
- AI enabling new Hacking tools
  - DeepExploit, WormGPT, Pwnagotchi (WiFi)
- AI written malware
  - A CyberArk research team got ChatGPT to write polymorphic malware

Summary

AI Technology is here and is already being deployed by threat actors for malicious activity.

There is growing research in attacks on AI Technology, how will we defend AI systems?

Increased use of AI Tools in IT infrastructure and in Cybersecurity Defense

Increased use of AI Tools in Threat Hunting to automate manual tasks

AI based tooling is enabling both attackers and defenders

Scenario: Exfil via Powershell

Sample Method Walkthrough -1

1. You read about an attacker technique using PowerShell to exfiltrate data in a recent public incident report.
2. You create a hunt and decide to look for: automated data exfiltration via PowerShell in your own environment.
3. You further investigate the specifics of the technique on a threat analysis blog and discover: use of PowerShell may sometimes alter `User-Agent` strings.
4. Based on this, you narrow your focus to look for:
   - anomalies in HTTP user agent strings
   - consistent and reoccurring `HTTP PUT` methods
   - possibly `HTTP POST` methods (beware false positives here)

Sample Method Walkthrough -2

6. You look in your internal datasets for the activity.
   - Netflow ("flow" data in general)
   - Packet Captures
   - Proxy Logs
   - Firewall Logs (if logging HTTP headers)
   - *Others (not configured in your environment but very helpful if you had them…):* Sysmon, Windows Event Logs for PowerShell, PowerShell Transcript Logs
7. in Proxy Logs, you find a User-Agent string which contains *"Mozilla/5.0 (Windows NT; Windows NT 6.1; en-US) WindowsPowerShell/3.0".*

```
$ rwfilter --start-date=2004/10/04:20 --end-date=2005/01/08:05  \
       --sensor=S0,S1 --type=all --proto=1,6,17 --print-volume \
       --threads=4 --pass-destination=stdout                   \
  | rwuniq --fields=proto --sort-output                        \
       --values=records,bytes,packets,stime,etime

     |     Recs|    Packets|        Bytes|  Files|
Total|  5866314|  155520999|  88858102591|    452|
 Pass|  5851584|  155228649|  88779771406|       |
 Fail|    14730|     292350|     78331185|       |
pro|Records|      Bytes|  Packets|     sTime-Earliest|       eTime-Latest|
  1| 321678|   58471992|   865991|2004/10/04T20:03:44|2005/01/08T05:28:34|
  6|1935300|75022603954|127277668|2004/10/04T20:03:41|2005/01/08T05:28:37|
 17|3594606|13698695460| 27084990|2004/10/04T20:03:41|2005/01/08T05:28:37|
```

Sample Netflow SiLK query and its output

https://tools.netsa.cert.org/silk/referencedata.html

Sample Method Walkthrough -3

9. You then look in Netflow and use the machine name and IP address of the host from the proxy logs to determine any other connections that internal machine is making internally and externally.
   - *You discover a consistent pulse every 2 hours on the internal machine during working hours for a period of 3 days.*
   - *You find the machine is connecting each pulse to a different domain, but you use Passive DNS data and you see that each resolves to one of 3 IP addresses.*
10. You decide to investigate the host for further infection and pass it to the incident response team…

Automation to Ease Hunting in the Environment

Based upon this discovery and the information from other threat reports, you decide to automate the hunting of more types of PowerShell command activity.

One approach (there are many others)

1. Install MSFT Sysmon. Configure it to record windows process execution, network utilization, etc. on each machine via policy.
2. Connect Sysmon to Splunk Tech Add-on (TA) (or connect it to other tools).
3. Write a Splunk alert for given processes and events from the trends such as: powershell.exe, cmd.exe, or net.exe.
   - Note: You will need to tune a lot here – some processes generate a lot more 'noise' than others.
   - Powershell.exe can be obfuscated via cmd.exe (Metasploit does this).
   - You read more articles on the use of powershell by attackers.
4. You begin to receive alerts and investigate suspicious commands and flags. You tweak the alerts based upon false positives in your environment.

Tools that May Help Hunters (a sample)

Carnegie Mellon University Software Engineering Institute

Hunting is a human's investigation of their environment for malicious activity – in particular, activity designed to evade traditional tools. New tools may be needed to discover activity. Here are a few tools used by threat hunting teams:

- Sysmon
- OSQuery, Kolide Fleet, Graylog
- Powershell Empire (also a post-attack tool)
- Caldera
- Graylog
- Ansible Playbooks
- Factor
- Sysdig

There are many other tools available. For example, see "threat hunting solution providers" on Wikipedia.

Example: Using OSQUERY to Find (vulnerable) Browser Extensions -1

You read about CVE-2017-6753 “Cisco WebEx Browser Extension Remote Code Execution Vulnerability” being used for attacks.

You suspect that some of your employees have WebEx installed for working with customers but are unsure which users on which machines have done so. Nor do you know what version they are using.

The following versions are affected

- Versions prior to 1.0.12 of the Cisco WebEx extension on Google Chrome
- Versions prior to 1.0.12 of the Cisco WebEx extension on Mozilla Firefox

Your end users should not have this version anymore, but what if some do?

Example: Using OSQUERY to Find (vulnerable) Browser Extensions -2

Doug Wilson from Uptycs shows how to use OSQUERY to solve this

Chrome

```
osquery> select name, version from chrome_extensions
where name like "%Cisco%";
| name                    | version |
| Cisco WebEx Extension   | 1.0.12 |
```

Firefox

```
select name, version from firefox_addons
where name like "%Cisco%";
```

This example is local, but osquery can run a query across all browser extensions on all machines in the entire org - in minutes. (The syntax is slightly different for that query).

Using a daemon, OSQuery can aggregate host and user information across all machines into tables that can be queried in near-real time.

Questions and Discussion

End of Module

Resources

ALTernatives to Signatures (ALTS) - George M. Jones John Stogoski April 2014 - WHITE PAPER CERT-CC-2014-35 (http://blog.sqrrl.com/the-cyber-hunting-maturity-model)

HP (https://www.rsaconference.com/writable/presentations/file_upload/anf-w04_hunting-the-undefined-threat-advanced-analytics-_-visualization.pdf)

Automatically Detecting Vulnerable Websites Before They Turn Malicious (https://www.usenix.org/system/files/conference/usenixsecurity14/sec14-paper-soska.pdf)

An Anthropological Approach to Studying CSIRTs (http://www.arguslab.org/documents/spsi_csirts_preprint.pdf)

Prioritizing Information Security Risks with Threat Agent Risk Assessment (http://www.intel.com/Assets/en_US/PDF/whitepaper/wp_IT_Security_RiskAssessment.pdf)

The Diamond Model of Intrusion Analysis (http://www.dtic.mil/get-tr-doc/pdf?AD=ADA586960)

Appendix

Hunting Template

Background (Problem)

Hypothesis Description

Where will you search?

How will you search?

What is a summary of your reasoning?

What do you expect to find?

(Bonus) How will you measure the hunt?

Hunting Template – Example

Background	We found some weird port behavior from a device that occurred during non-working hours.
Hypothesis Description	Determine if the behavior is malicious, if the team authorized the activity. If authorized, determine if any company-wide policies should be created for teams exploring IoT devices.
Where will you search?	Identify all machines with the behavior, and on the network for other similar behavior. Specifically, what process is causing the weird port behavior activity, and why is it occurring after hours?
How will you search?	Talk to the user and the machine's team owners. Investigate & Understand what the device is intended to be doing, why the change in behavior might have occurred, and what if any explanation there might be for it occurring after normal hours.
What is a summary of your reasoning?	This machine is probably transmitting IoT data to a cloud server (the machine appears to be on the network segment for a team that might be experimenting with new applications.) We don't want to disrupt their work, but we need to know if this is legitimate behavior.
What do you expect to find?	We expect to find a machine transmitting data to a cloud server – but we don't currently recognize the destination and there is a chance it would be malicious. We also would like to know if the team is testing a new IoT device.
How will you measure the hunt?	Combination of hours spent vs. importance of findings vs. impact of threat (if any). Outcomes might be: discovery of malicious activity, understanding of new cyber activity a group is exploring using in future, thinking of ideas for a new company policy, and more.

Reporting – Sample Report Contents

Threat Hunting Goal

- What you were hunting and why.

Data

- What data did you gather and where was it from?

How You Hunted and What Analysis You Did

- Description of your Method. Commands you ran on the environment. Results.

Evidence of Findings

- The evidence found, possibly displayed in a chart or graph.

Risk Implications for the Organization

- Include remediation performed (if any) and any remaining risk.

Future Recommendations to the Organization based upon findings

Reporting – Example Post Hunt Options

Sample optional security recommendations following a VPN credential hunt:

- Help IT implement a process / tool to ensure that certificates can only be distributed and stored in encrypted form.
- Move from single factor to multi-factor.
- Move from certificate-based multi-factor to token based.
- Make end user machines more resilient to widely available certificate extraction tools, or have hosts to detect their presence/usage.
- Determine when users on the VPN are acting out of character (Hard).

Reporting – Lessons Learned

It is important to identify your audience and communicate effectively with them.

Who is your audience – what ‘action’ do you want from them?

What is their criteria for making a judgement?

Principles for presenting data findings to audiences in general

- data visualization
- communication of technical results

Example from “Quality Over Quantity: Determining Your CTI Detection Efficacy” by David Bianco

Plot of collected IOCs against the Pyramid of Pain (Y axis) and the Adversary Lifecycle in which that IOC occurred (X axis). This can tell you:

- what do I have?
- where is it most valuable?
- what am I missing?

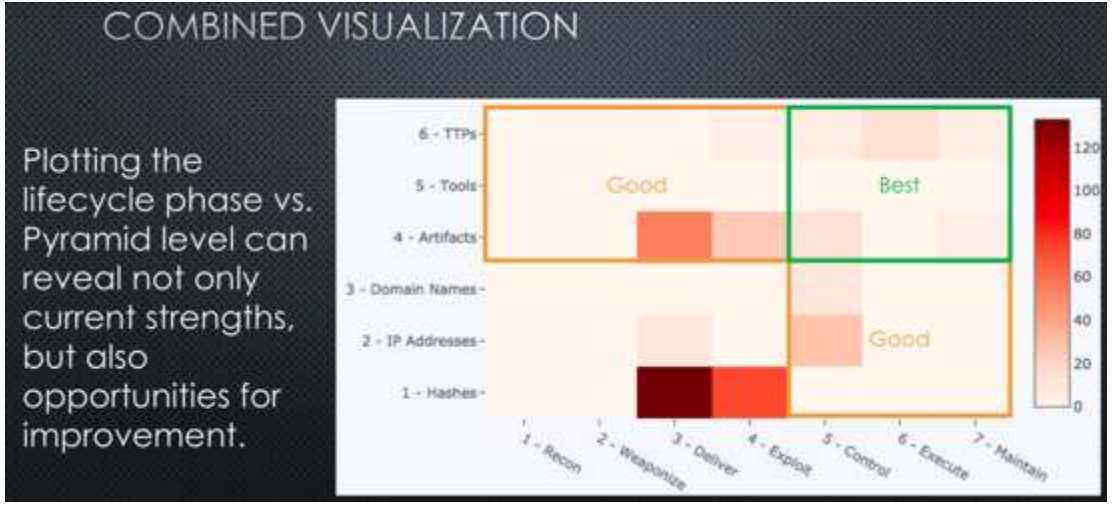

Source: Bianco, David, Quality Over Quantity: Determining Your CTI Detection Efficacy, SANS CTI Summit 2019, Arlington, VA

Other Hunting Process Recommendations – Results from an SEI Study

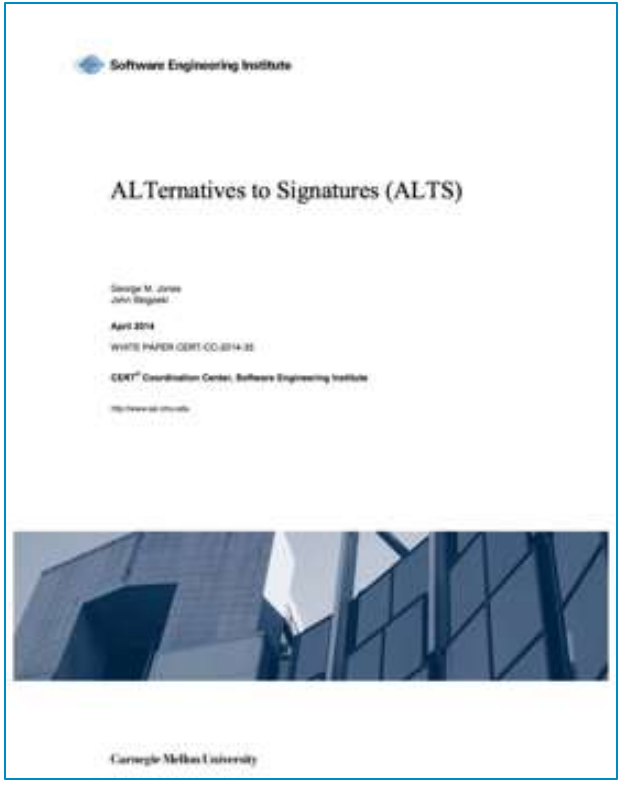

By George Jones, John Stogoski

- Promote adoption of hunting operations, sandboxing, DNS analytics, and network profiling with both policy and funding, including R&D.
- Develop an HR and staffing strategy to support hunting.
- Integrate research roles into operational environments for joint learning.
- Augment DNS designs to improve DNS analytics and the collection of passive DNS.
- Clarify policies regarding deception.
- Determine policies to address security and privacy concerns related to large collection and storage.

Differences Between Hunting and Other Teams

Hunting Teams compared to Incident Management Teams

- Tend to focus on threats that are pervasive or that are against the most high-profile targets
- More proactive toward detection as well (run traps)
- May be temporary (e.g., future acquisition date is planned)

Hunting Teams compared to Pen Testing Teams

- More specialized toward a range of current and near-term capabilities of most likely threats – not searching for all possible vulnerabilities
- Likely to work with other teams to
  - exchange data
  - design tests and ensure that highest risk areas are being tested
  - consult with experts, borrow skills, etc.

FYI Definition of Threats

Carnegie Mellon University Software Engineering Institute

RFC 4949 – Internet Security Glossary

- Threat - 1a. (I) A potential for violation of security, which exists when there is an entity, circumstance, capability, action, or event that could cause harm. (See: dangling threat, INFOCON level, threat action, threat agent, threat consequence. Compare: attack, vulnerability.)

NIST 800-53 rev. 4 references the definition proposed by the Committee on National Security Systems (CNSS)

- Threat [CNSSI 4009, Adapted] – Any circumstance or event with the potential to adversely impact organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, or the Nation through an information system via unauthorized access, destruction, disclosure, modification of information, and/or denial of service.

Intel Corp – Threat Agent Risk Assessment Resources

Carnegie Mellon University Software Engineering Institute

Threat Agent Risk Assessment (TARA)

- concentrates on agents, motives, methods, and objectives
- maps to controls, not weak points
- **attempts to determine most likely attack**

Component: threat agent library (TAL)

- defines eight common threat agent attributes, such as intent
- identifies 22 unique threat agent archetypes, such as disgruntled employee, competitor, and organized crime

Component: common exposure library (CEL)

- maps vulnerabilities against existing controls to show residual exposures

Component: methods and objectives library (MOL)

- lists known threat agent objectives—what they want to accomplish
- lists the most likely methods they will use to reach these objectives

http://www.intel.com/Assets/en_US/PDF/whitepaper/wp_IT_Security_RiskAssessment.pdf

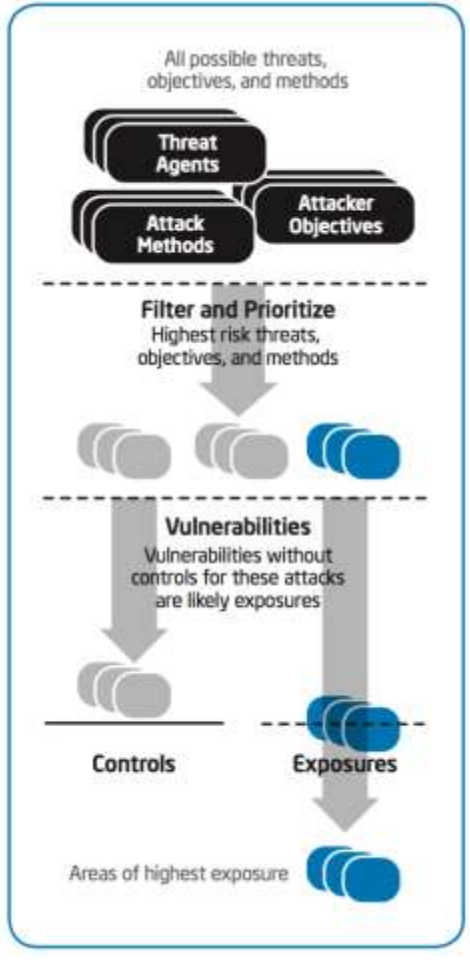

Other Threat Modeling Resources

The Diamond Model of Intrusion Analysis
(http://www.dtic.mil/get-tr-doc/pdf?AD=ADA586960)

MITRE ATT&CK
(https://attack.mitre.org/index.php/Main_Page)

Threat Modeling – Designing for Security by Adam Shostack
(http://threatmodelingbook.com/)

Comparison of Threat Modeling Methods by Mead & Shull
(publication pending)

Teaming for Automation

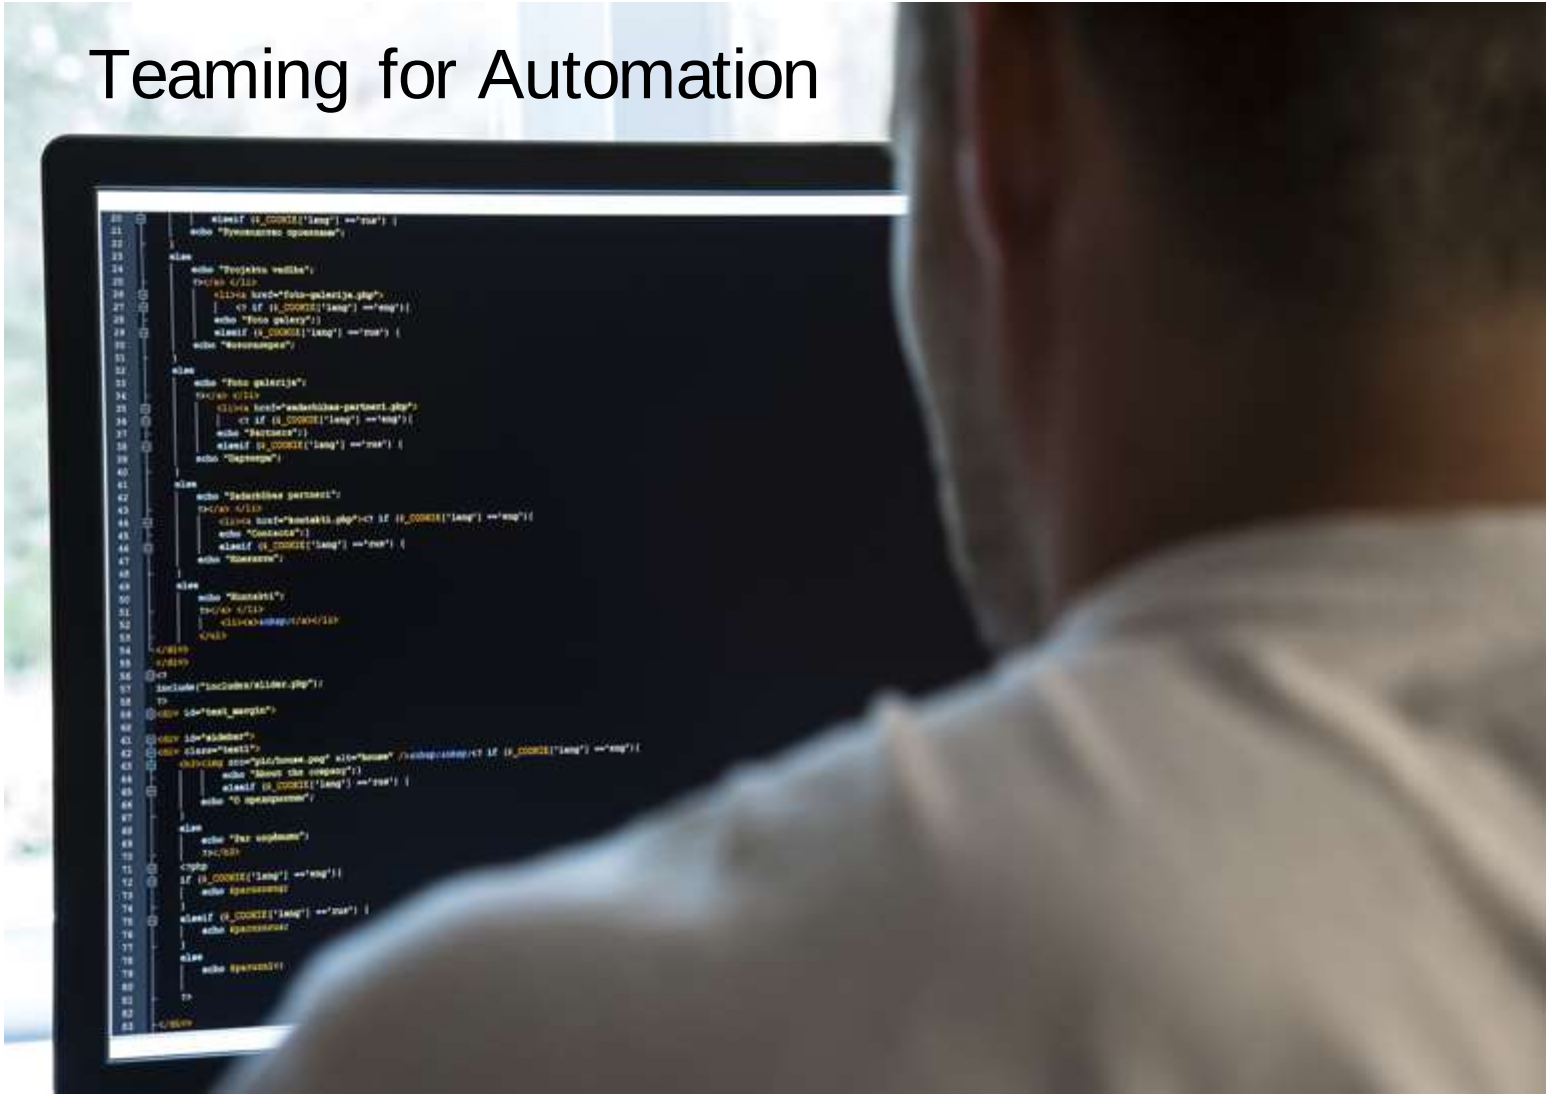

A Story on Teaming

From *An Anthropological Approach to Studying CSIRTs*

- A Research Analyst was sitting with a SOC and wanted to help them by building something.
- They wouldn't assign him to sensitive incidents, and rarely had time to talk with him.
- He felt that all his SOC time was spent on carrying out repetitive operational tasks and was frustrated because he did not feel he was gaining any insight at all.
- The SOC was focused on getting incidents processed quickly.
- It did not have the time for contemplating a long-term vision of improved efficiency.

Sundaramurthy, Sathya Chandran, et al. "An anthropological approach to studying CSIRTs." *IEEE Security & Privacy* 5 (2014): 52-60.

Observed Repeated Events

From *An Anthropological Approach to Studying CSIRTs*

- The SOC receives alerts on malicious network traffic from a number of trusted sources as well as from their own intrusion detection system (IDS).
- The alerts contain the IP address (usually that of the NATing firewall) of the infected host and the external IP address with which it was communicating.
- The real internal IP address has to be extracted from the firewall logs; the MAC address identifying the infected host from DHCP logs.
- Finding the log entry for a given event and looking up the associated information to resolve the ticket takes about 5 minutes.
- This repeats and repeats.

Sundaramurthy, Sathya Chandran, et al. "An anthropological approach to studying CSIRTs." *IEEE Security & Privacy* 5 (2014): 52-60.

Engineering for Scale

From *An Anthropological Approach to Studying CSIRTs*

Set Small Goals

- He thought about ways to speed up the ticket handling by building a database of connections and an IP address to MAC address mapping.
- Noting that most active alerts are a week or less old, he decided to build a caching database retaining seven days of mapping information.

Try new things and fail fast

- Initially tried using MySQL but it didn't index inputs in real time.
- Tried MongoDB which stores data as JSON and has a sufficiently high ingest capability.

Recognize Success

- Asked the incident response analyst to use the database.
- The analyst was extremely happy with the performance improvement which reduced handling time from from five minutes to two seconds.

Iterate

- Led to discussion of new tool extensions and new data types to expand
- Eventually, the two arrived at a "Threat Intelligence Framework" that added information sources and relationships among them to the database, allowing a variety of incidents to be handled.

Sundaramurthy, Sathya Chandran, et al. "An anthropological approach to studying CSIRTs." *IEEE Security & Privacy* 5 (2014): 52-60.

2014 Survey & Interviews of Methods

In 2014, the SEI performed a survey and assessment of non-signature-based approaches, tools, and techniques:

By George Jones, John Stogoski
http://resources.sei.cmu.edu/library/asset-view.cfm?assetid=296146

Study Goals

- Collect non-signature-based approaches, tools, and techniques.
- Observe maturity and adoption of methods.
- Identify promising emerging methods.
- Focus on network-based detection of malicious activity.

Interviews with Participants Revealed

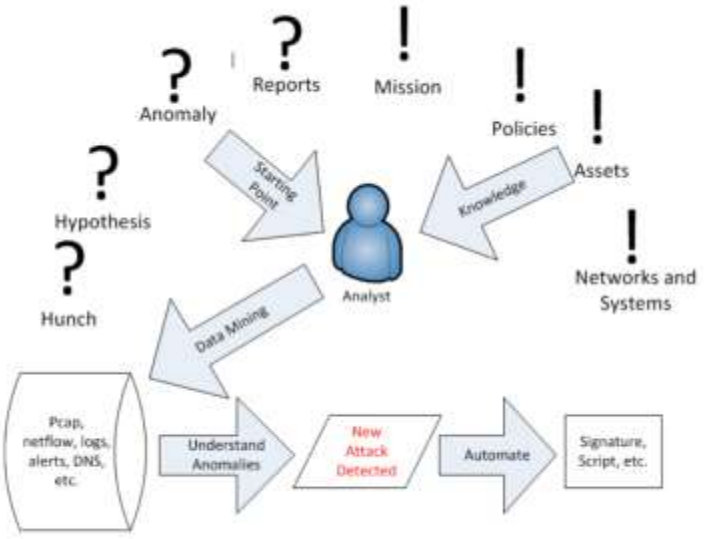

Figure 2: Hunting Operations Process

Hunting Process

1. Mine data.
2. Identify suspicious activity.
3. Investigate.
4. Codify a repeatable analytic.

Analysts perform heuristic queries informed by expert knowledge of the allowed and expected behavior, the controls implemented, and other situational knowledge.

Other Attacks on AI Technology

Model poisoning

Prompt injections

Adversarial AI

Researchers tricked self driving car algorithms using stickers on stop signs

Source: https://www.autoblog.com/2017/08/04/self-driving-car-sign-hack-stickers/

Example: Hunting for Credential Theft (from 2015)

Walk Through Example of Threat Hunting

You read a report from Mandiant which says that the most commonly observed attack across all Mandiant engagements was:

- VPN compromise methods
  - This particular method gives attackers two huge advantages.
    1. They can persist in an environment without having to deploy backdoors.
    2. They can blend in by imitating authorized users.

Example of How Threat Hunting works

The report further lists the main ways that attackers were observed compromising VPN credentials of other victims. The ways differ depending upon the type of environment the victim has:

- If Single factor – Attackers re-used credentials stolen from compromised end-user systems or the Active Directory domain.
- If Certificate-based multi factor – Attackers used available tools (such as Mimikatz) to extract certificates from compromised end-user systems or found certificates that had been distributed in an insecure manner.
- Also, Attackers stole credentials via direct compromise – this was less common than the others!

You Think of a Few Hunting Goals and Pick One

You develop the following goal.

Prevent VPN compromises by looking for insecure certificates and insecure distribution of certificates.

For this goal, you plan to look in the following places:

- attached emails in unencrypted form
- available on open network file-shares
- posted in SharePoint systems

Skills Needed to Achieve Hunting Goal

You think about the skills team members will need to accomplish all the items in this goal.

Necessary skills for this Hunt include

- network & infrastructure: Where are all the places in our infrastructure we should be hunting?
- security SMEs: What are tell-tale signs of insecure certificates?
- programming: How can we automate the hunt?
- data science: How prevalent is this problem for us?
- visualization: How do we explain the problem and report the results?
- IT/process: What are all the ways you distribute certificates?

You assign or pull those individuals into the team.

You Think of Another Goal

Goal: **Find attacks on our network that are trying to use stolen certificates or attempting to steal certificates.**

List of hunting methods

- Collect source IPs and geolocation for connections. Alert on large location changes (country/state) which is similar to methods used by enterprise cloud apps.
- Work with departments to reduce false positives for certain staff members.
- Alert on presence of tools like Mimikatz (and others) in traffic and on hosts such as via Yara Rules.
- Create fake certificates and watch them (aka “Honey Hashes”).
  - Consider staging them in risky machines/areas, in DMZ, and/or randomly.
  - Set up alerts for attempted use of the fake accounts.
  - Other design considerations: name schema for usernames, high privileges on your domain, proper metadata (last login, etc.).

Insight From the Hunt

These hunting methods can take a lot of time and may miss some activity.

You look into advanced methods and more resources on ways to perform credential theft.

User Behavior Activity Monitoring

- Enterprise tools for comparing user behavior against itself such as
  - Rapid7 InsightUBA (user behavior analytics)
  - Microsoft Advanced Threat Analytics

Resources for Hunting Credential Thefts

- https://dfir-blog.com/2015/11/24/protecting-windows-networks-dealing-with-credential-theft/
- https://isc.sans.edu/forums/diary/Detecting+Mimikatz+Use+On+Your+Network/19311/
- https://adsecurity.org/?tag=yara

After Action

Many other ideas could have been developed from the VPN observation:

- Move from single factor to multi-factor.
- Move from certificate-based multi-factor to token based (or other).
- Make end user machines more resilient to widely available certificate extraction tools or have hosts to detect their presence/usage.
- Help IT implement a process / tool to ensure that certificates can only be distributed and stored in encrypted form.
- A harder goal is to determine when users on the VPN are acting out of character.

Threat Landscape Example – Microsoft SIR

Figure 35. Encounter rates for different types of exploit attempts in 2015

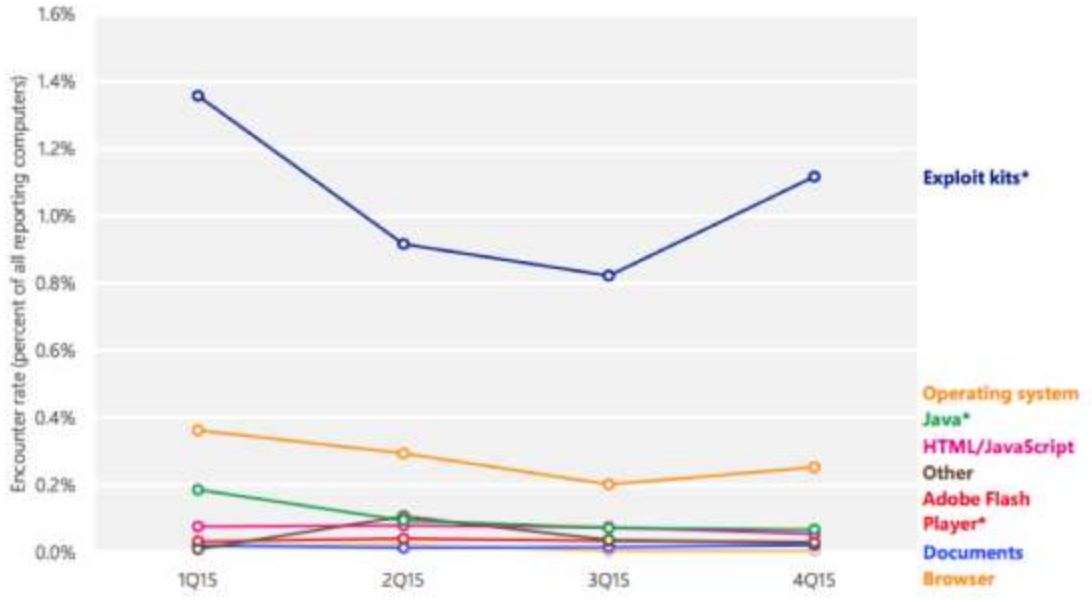

* Figures for exploit kits, Java, and Adobe Flash Player exploits are affected by IExtensionValidation in Internet Explorer, which blocks many threats before they are encountered. See page 76 for more information.

Chart from Page 64 of Microsoft SIR #20 (https://www.microsoft.com/security/sir/)

Encounter rate is the percentage of computers running Microsoft real-time security products that report a malware encounter. … Only computers whose users have opted in to provide data to Microsoft are considered when calculating encounter rates. See page 157 for Data Sources

Example: Kerberoast Analysis

Windows Event Log entries

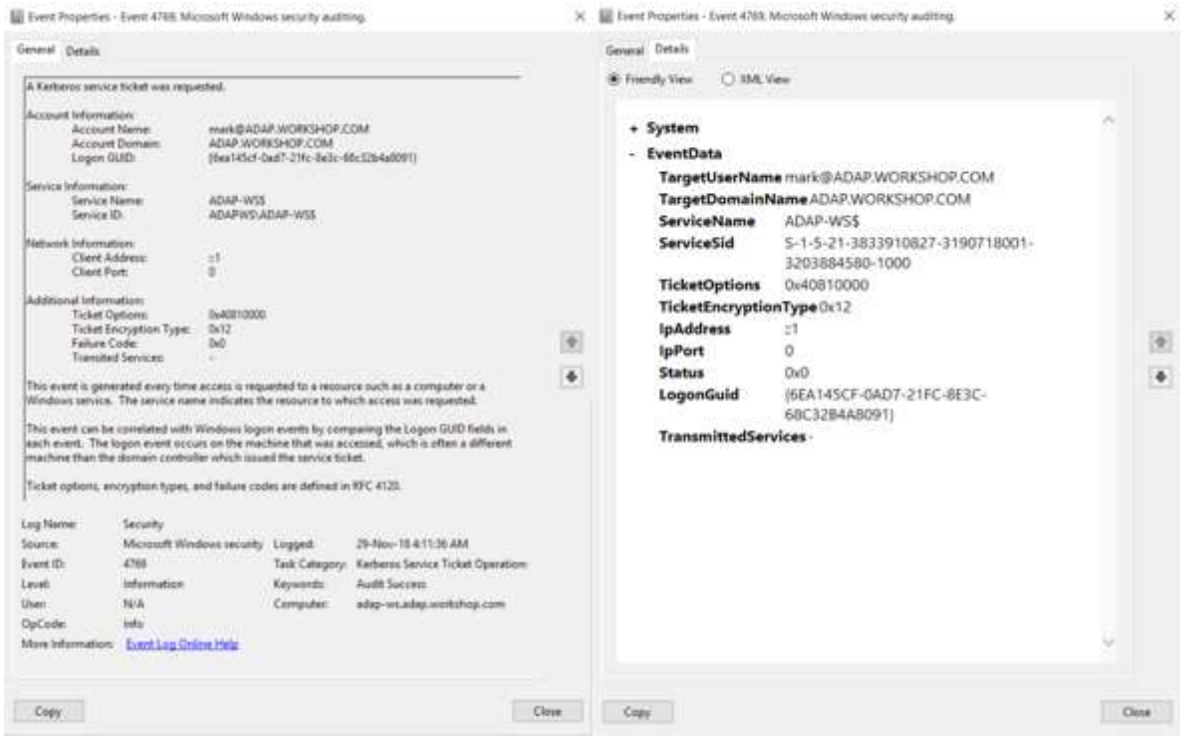

Example taken from https://www.manageengine.com/products/active-directory-audit/kb/windows-security-log-event-id-4769.html

Findings

- The windows event log indicates Event 4769 a Kerberos service ticket was requested.
- The Ticket encryption type is 0x12.

Hunting Looks for Rare Events

Carnegie Mellon University Software Engineering Institute

Hunting is the exploration of guesses that you think are likely, but you might be wrong.

For example:

- Question: What is the danger of cryptomining malware if it doesn't affect system performance?
- Answer: Low?
  - Unless the attackers behind it decide to expand to more lucrative system/data targets when they see that you did not remove their malware.
- Answer: On second thought maybe Med or High.

"Attack likelihood" is difficult to estimate and is also a moving target.