578.4
Analysis and Production of Intelligence

SANS

THE MOST TRUSTED SOURCE FOR INFORMATION SECURITY TRAINING, CERTIFICATION, AND RESEARCH | sans.org

FOR578.4 | Cyber Threat Intelligence

Analysis and Production of Intelligence

Robert M. Lee (Lead Author)

Robert M. Lee is the CEO and Co-Founder of the industrial (ICS/OT/IIoT) cybersecurity company Dragos, Inc. which creates and delivers technology, services, and cyber threat intelligence to the industrial community. He is a SANS Senior Instructor and the course author of SANS ICS515: ICS Active Defense and Incident Response and the lead author of SANS FOR578, Cyber Threat Intelligence. Robert is also a Department of Energy employee serving on the Electric Advisory Committee and the Vice-Chair of the Grid Security Committee. He also serves on the World Economic Forum's Oil and Gas and Electricity Subcommittees focusing on the cybersecurity of global infrastructure.

Robert obtained his start in cybersecurity in the U.S. Air Force, where he served as a Cyber Warfare Operations Officer tasked to the National Security Agency. He has performed defense, intelligence, and attack missions in various government organizations, including the establishment of a first-of-its-kind ICS/SCADA cyber threat intelligence and intrusion analysis mission. Routinely sought for his expertise, he has keynoted and spoken at major conferences such as RSA, BlackHat, and DEFCON and has testified to the U.S. Senate Energy and Natural Resources Committee. Robert is also the author of the book *SCADA and Me*, *Cyber Threat Intelligence and Me*, and *Santa and Me: The SCADA before Christmas* as well as the weekly webcomic LittleBobbyComic.com. Robert may be found on Twitter @RobertMLee or contacted via email at RLee@Dragos.com.

Rebekah Brown (co-author)

Rebekah is a cybersecurity and intelligence analysis professional specializing in threat intelligence, network warfare analysis, systems analysis, and threat modeling. Rebekah spent over a decade on active duty as a cryptologic linguist, network warfare analyst, and cyber operations chief in the U.S. Marine Corps before moving to the private sector, where she has developed threat intelligence programs at multiple Fortune 500 companies. She received degrees in International Relations from Hawaii Pacific University, and Homeland Security with a cybersecurity focus, and a graduate certificate in Intelligence Analysis from American Military University. She is a published author, instructor, and public speaker on intelligence-driven incident response and adversary tactics.

Course Agenda

- Cyber Threat Intelligence and Requirements
- The Fundamental Skill Set: Intrusion Analysis
- Collection Sources
- Analysis and Production of Intelligence
- Dissemination and Attribution
- Capstone

This page intentionally left blank.

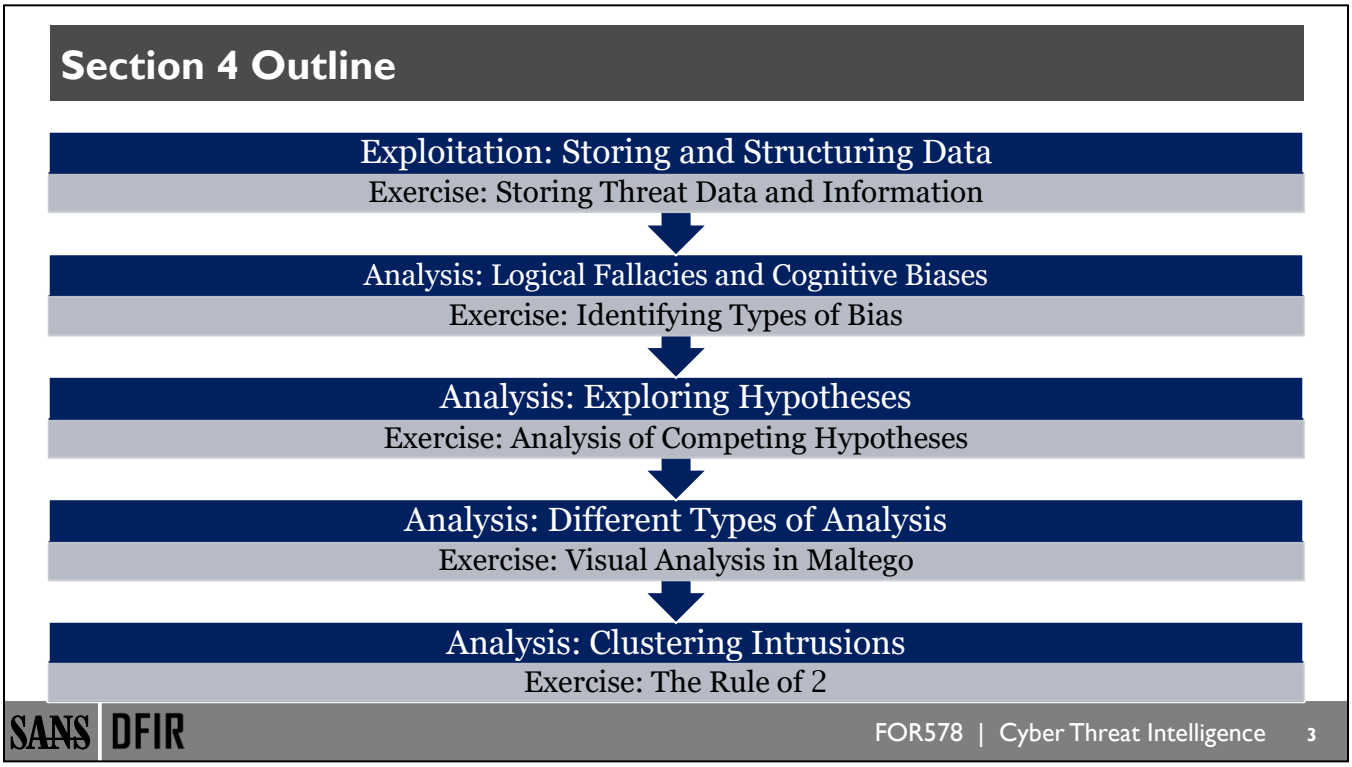

This page intentionally left blank.

Case Study: Human Operated Ransomware

This page intentionally left blank.

Human Operated Ransomware Operations

- Ransomware has been around for decades, but a few factors mixed in recent years to see a significant increase in the number of cases:
  - More initial access vulnerabilities and tradecraft proliferation
  - Easy access to open-source tools such as PowerShell and Cobalt Strike
  - Willingness of victims to pay the ransom
- Auto-spreading ransomware such as WannaCry poses an untargeted risk that can be relatively easily mitigated
- Human operated ransomware is a class of ransomware attacks where the adversary penetrates the network and tailors their operation to the victim
  - Much more impactful and time to remediate is often in the minutes or hours

Human Operated Ransomware Operations

Ransomware is hardly a new topic in information security, but the prevalence of it in recent years is noteworthy. The leak of the ETERNALBLUE vulnerabilities definitely enabled some ransomware operations such as WannaCry as a class of auto-spreading malware, but the cases that were most damaging were those that human adversaries leveraged directly. Because the adversary compromises the organization and takes time to learn the environment and where to place the malware for most impact, such as propagating it through the domain controller, these styles of operations have maximum impact. Many folks inappropriately linked some of the most major cases to ETERNALBLUE when the biggest impacted victims were actually impacted by the adversary stealing credentials inside the network and leveraging them. ETERNALBLUE fell off as organizations patched it and yet the human operated ransomware cases only went up. There are many factors for this, but the proliferation of criminal networks and tools combined with a willingness of victims to pay definitely contributed.

Reference:

- https://www.microsoft.com/security/blog/2020/03/05/human-operated-ransomware-attacks-a-preventable-disaster/

Wadhrama Attack Chain by PARINACOTA

- Microsoft's intel team did a great job of abstracting out the technical details to provide a TTP and kill chain view of different ransomware attacks
- Note the overlay of MITRE ATT&CK IDs per adversary step in an easy to consume method

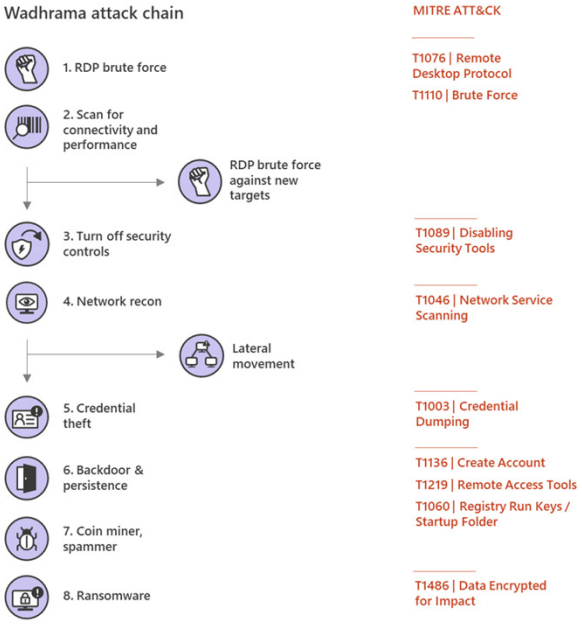

SANS DFIR

Wadhrama Attack Chain by PARINACOTA

In this case study, we're going to heavily reference Microsoft's excellent blog post, Human Operated Ransomware Attacks: A Preventable Disaster. Too often, analysts associate cyber threat intel with indicators of compromise or just clusters of intrusions. Threat feeds and pew pew maps instantly come to mind when people think about CTI. But some of the most effective CTI is the intel that's easily consumed by defenders. Here, we have beautiful and easy to understand visualizations for defenders that combine TTPs and a kill chain view of the attack to educate defenders on what they need to be prepared to deal with.

Reference:

- https://www.microsoft.com/security/blog/2020/03/05/human-operated-ransomware-attacks-a-preventable-disaster/

Doppelpaymer Ransomware

- The Doppelpaymer ransomware is a good example where the nuance between TTPs and tradecraft can matter
  - There might be RDP access like Wadhrama, but the use of Dridex or other malware would early on signify an entirely different Activity Group and thus a different investigation strategy
- The defensive strategies ahead of time may overlap but where you investigate and how, especially in a time sensitive scenario, can massively depend on deep understanding of the adversary

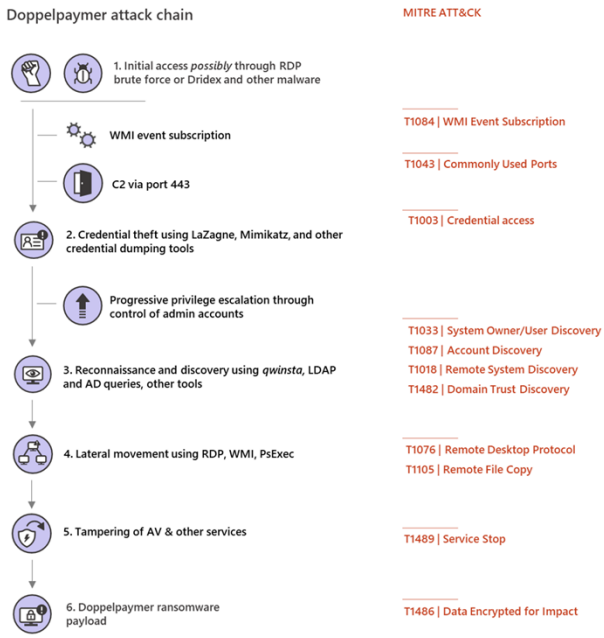

SANS DFIR

Doppelpaymer Ransomware

TTPs and tradecraft are different. This isn't nuance that's normally worth debating, but in this case, we see the nuance matters. Tradecraft is a combination of tactics, techniques, and adversary capabilities. Tradecraft is not necessarily specific to any one adversary, but it forms a very tangible insight to build defensive playbooks off of. Further, if you are compromised and see RDP abused, you could easily associate this activity accidentally with Wadhrama, which is going to give you very different later stage kill chain views. If you see that Dridex or other malware is involved with the RDP abuse, though, you could quickly ask your defenders and investigators to look for different tradecraft internal to the network to include looking for the Doppelpaymer capability. When these attacks take place over the course of minutes or a few hours, the difference in how you guide and train your incident responders can be the difference in a costly breach or being able to enjoy your afternoon.

Reference:

- https://www.microsoft.com/security/blog/2020/03/05/human-operated-ransomware-attacks-a-preventable-disaster/

Ryuk from TrickBot Infections

- As another example of tradecraft, we see not only the TTP which is the Spearphishing Attachment but the addition of the adversary's capability, which is TrickBot
- Cobalt Strike, PowerShell, and eventually the Ryuk ransomware give a combination of tools and TTPs that can help build much more robust defensive strategies without ever focusing on the indicators

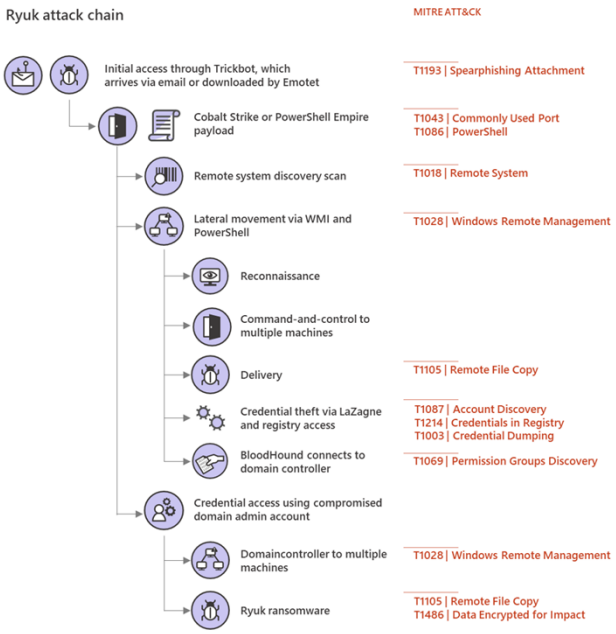

SANS DFIR

Ryuk from TrickBot Infections

One particularly common form of human operated ransomware took advantage of TrickBot or Emotet as a loader, first stage malware, or first stage implant, to then allow the adversary to leverage open-source tools and native tools such as Cobalt Strike and PowerShell to expertly move around the network to ultimately position the Ryuk ransomware. The combination of Bloodhound to compromise and leverage the Domain Controller inside the network makes this combination of TTPs and capabilities particularly difficult to defend against. Looking simply for indicators in these types of scenarios are unlikely to be highly effective.

Reference:

- https://www.microsoft.com/security/blog/2020/03/05/human-operated-ransomware-attacks-a-preventable-disaster/

Make It Easy for Defenders

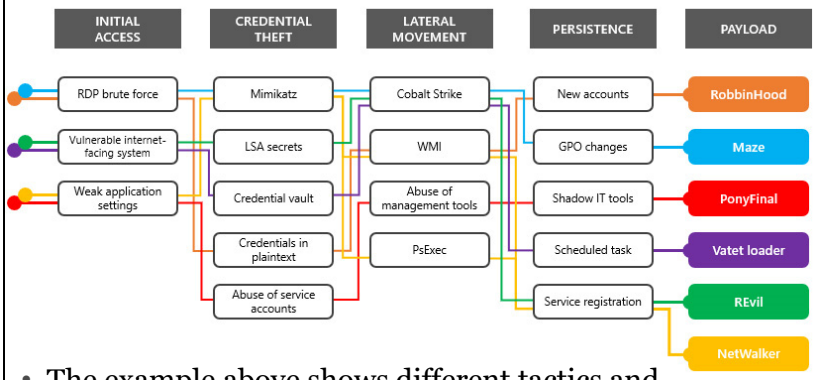

- The example above shows different tactics and techniques with easy-to-follow visual cues, which can map to malware families that defenders might encounter

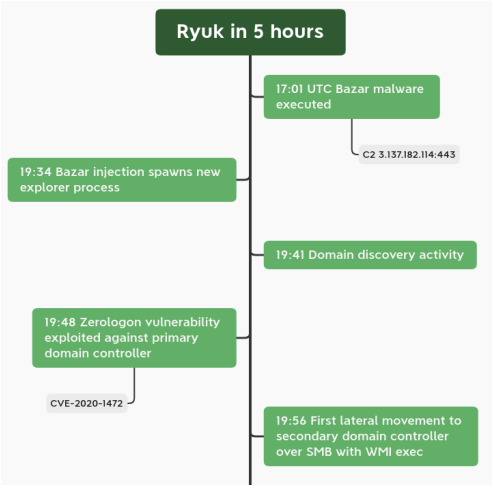

- The above is a great example of part of a timeline to give the defender a "kill chain" like view with an understanding of timeline

Make It Easy for Defenders

Remember, the intelligence output that benefits the intel analyst more than the defender is an embarrassment on our profession. Creating intelligence that's easily consumable to help satisfy requirements, even broad requirements such as education of large groups of people, is a great example of intelligence. I understand that the books aren't in color, but as presented in the class, the colors will come through and can also be seen in the second link below. The colors on the graph tying each family of malware to the different objectives and tactics above is a great way to communicate to defenders what they need to be ready to deal with.

The other graphic showing Ryuk operations and how they can be done in five hours or less with what you might expect to see is a really great way to educate people; these types of intelligence products are great to use not only in broad education but also in tabletop exercises ahead of incidents.

References:

- https://thedfirreport.com/2020/10/18/ryuk-in-5-hours/
- https://www.helpnetsecurity.com/2020/04/30/ransomware-campaigns/
- https://redcanary.com/blog/how-one-hospital-thwarted-a-ryuk-ransomware-outbreak/

Example of Effective Visual Communication of TTPs

- Huge credit to Microsoft and their threat intelligence team for their public education but also their expert-level visual communication of TTPs and defenses to a wide audience
- Great example of a non-traditional and effective intelligence product

SANS DFIR

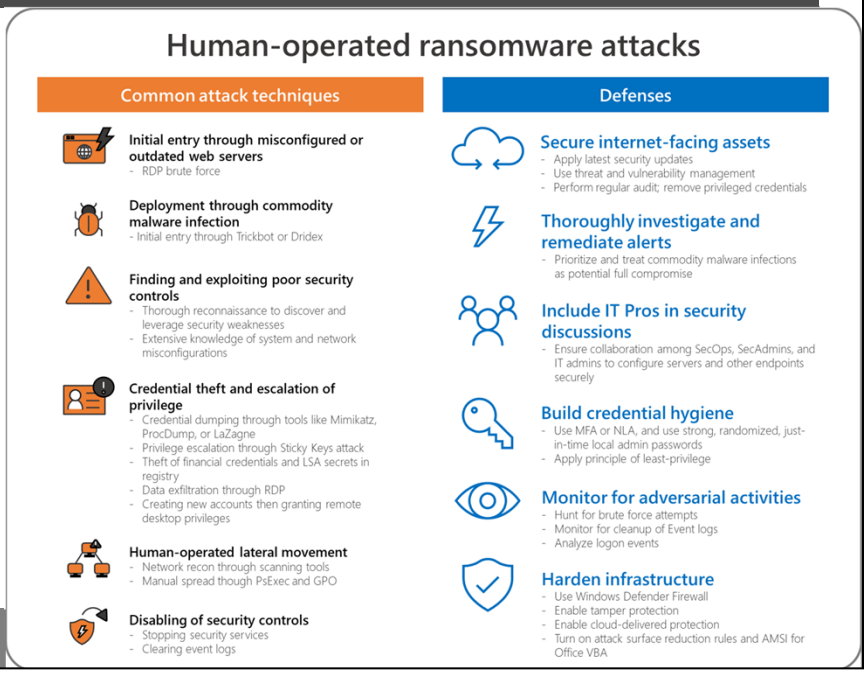

Example of Effective Visual Communication of TTPs

My favorite type of intelligence is the intelligence that focuses heavily on defender outcomes. Microsoft really did the community a service with their fantastic blog post and in a very easy to consume understanding of common attack techniques paired up with the most realistic defenses against them. Too often, intel analysts want to suggest all the types of defense possible for every attack technique. The recommendations then become too numerous to act upon and the reader loses focus. Using your expert understanding of the adversary in combination with your understanding of defense allows you to pair the 1-2 most critical defenses for everything you observe and paint a more tailored picture for the defender.

Reference:

- https://www.microsoft.com/security/blog/2020/03/05/human-operated-ransomware-attacks-a-preventable-disaster/

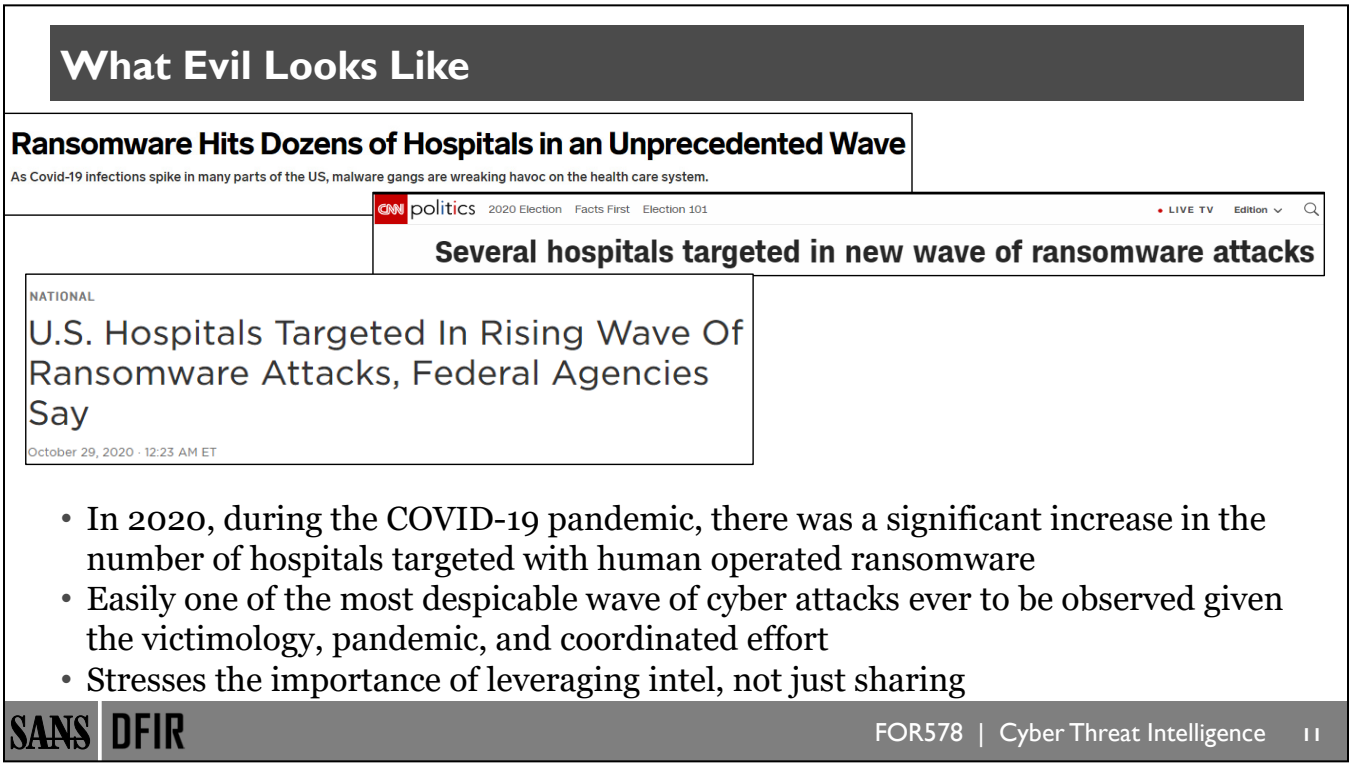

What Evil Looks Like

SANS would censor the words I'd like to use here, so I'll just say professionally: Those that orchestrated and coordinated the explicit targeting of hospitals especially during a pandemic in such a way that drove fear and discouraged people from seeking medical help—are evil. I don't use that word lightly.

But what's the lesson here? I hope this doesn't come off too critical, but I saw so many intel analysts start passing around indicators and sharing IOCs related to these series of attacks and patting themselves on the back. They deserve credit for doing something. I'm not putting that down. However, the hospitals that were getting hit were in no way at scale able to use, consume, or otherwise leverage that indicator sharing that was largely taking place on social media. Even at the ISAC level, there were few that were able to take advantage of it. Without being overly critical at good people trying to help, that is not my intention here. I simply want to call attention to the fact that in cases like this, the most useful intelligence was that which was either already in the hands of the incident responders responding to these attacks or to the service and product vendors already in place at those hospitals. The healthcare community, widely, was not well positioned to use indicators and that type of sharing to achieve any meaningful help outside the top hospitals in the world.

So again, thank you to those that tried to help, but we have to get creative as a community on how to leverage intel appropriately especially to teams that don't have SOCs or intel teams.

Exploitation: Storing and Structuring Data

This page intentionally left blank.

Storing Collected Intelligence

- Often discussed in the context of threat intelligence platforms (TIPs)
- The focus is on storing information in a quickly accessible and useful format
- Should be available to internal security personnel as well as analysts who will productize the information
- Some common tools include MISP, Threat_Note, and CRITs

Storing Collected Intelligence

It is important to be able to store collected information and intelligence in your environment. Storing the information in a usable and quickly accessible format allows it to be made available to those who need it, such as security personnel, as well as made available for intelligence analysts producing assessments. Storing intelligence is often fairly unique to the companies that do it, although there are some out-of-the-box threat intelligence storing/sharing platforms that can serve as a starting place.

Reference:

- https://www.enisa.europa.eu/publications/exploring-the-opportunities-and-limitations-of-current-threat-intelligence-platforms/at_download/fullReport

Storing Platforms

Open Source
- MISP
- Threat_Note
- CRITs

Pros: Free, ample storage, open-source sharing communities

Cons: Difficult to implement and maintain

Commercial
- ThreatConnect
- ThreatQuotient
- Anomali
- EclecticIQ

Pros: Fully supported, ease of installation, integration with other tools, data analytics

Cons: Can be pricey, may not fit established workflows

Storing Platforms

There is a wide variety of storing platforms out there. One of the biggest complaints usually is that it's not the storage that people struggle with but the access to good data. Storage platforms are something that should be considered later in the stages of doing internal analysis so that good requirements can be levied. I would highly recommend using open source for a while, determining requirements, and then moving to professional tools with support as the CTI team scales in size and responsibility.

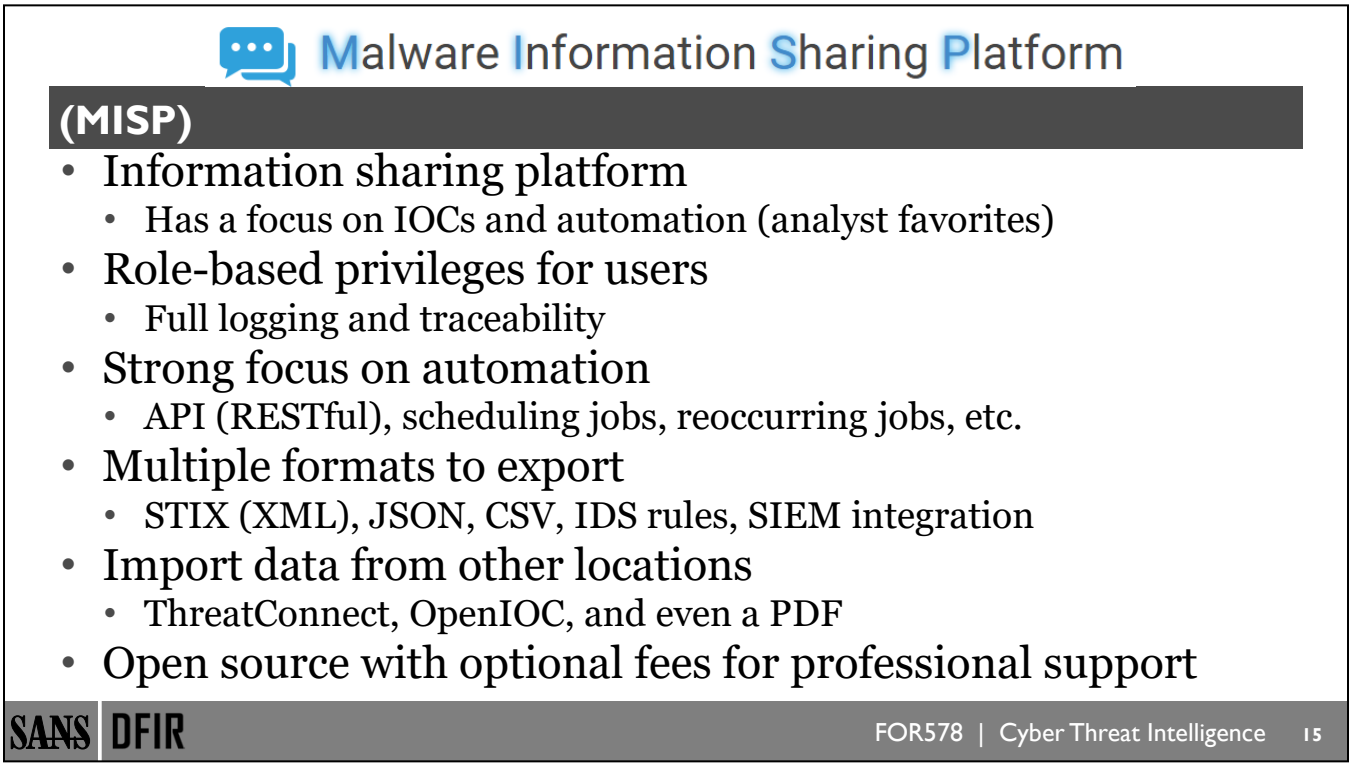

Malware Information Sharing Platform (MISP)

MISP is a sharing methodology/platform similar in nature to STIX/TAXII. MISP is unique in a number of ways, though, and can integrate with STIX or other IOC standards such as OpenIOC instead of needing to outright replace it. In other words, the efforts can be complementary if desired.

MISP has a strong usage in Europe due to a number of the CERT's active involvement in the development of the code base and community. The North Atlantic Treaty Organization (NATO) also sponsored the project and put an emphasis on it for the fusion of threat information between different NATO countries. This helps showcase the international flavor of MISP.

There are a lot of great features in MISP from user management (such as role-based privileges for users to ensure analysts only get access to the data they need, as well as logging and full traceability of user actions) to automation through a RESTful API and an ability to schedule jobs and reoccurring jobs. Additionally, a big benefit of the platform is the ability to import data from other locations, such as existing IOCs in formats like OpenIOC, as well as the ability to export the information in a wide variety of formats from STIX (XML), JSON, and CSV to IDS rules and connectors to popular SIEM systems.

References:

- https://www.misp-project.org/

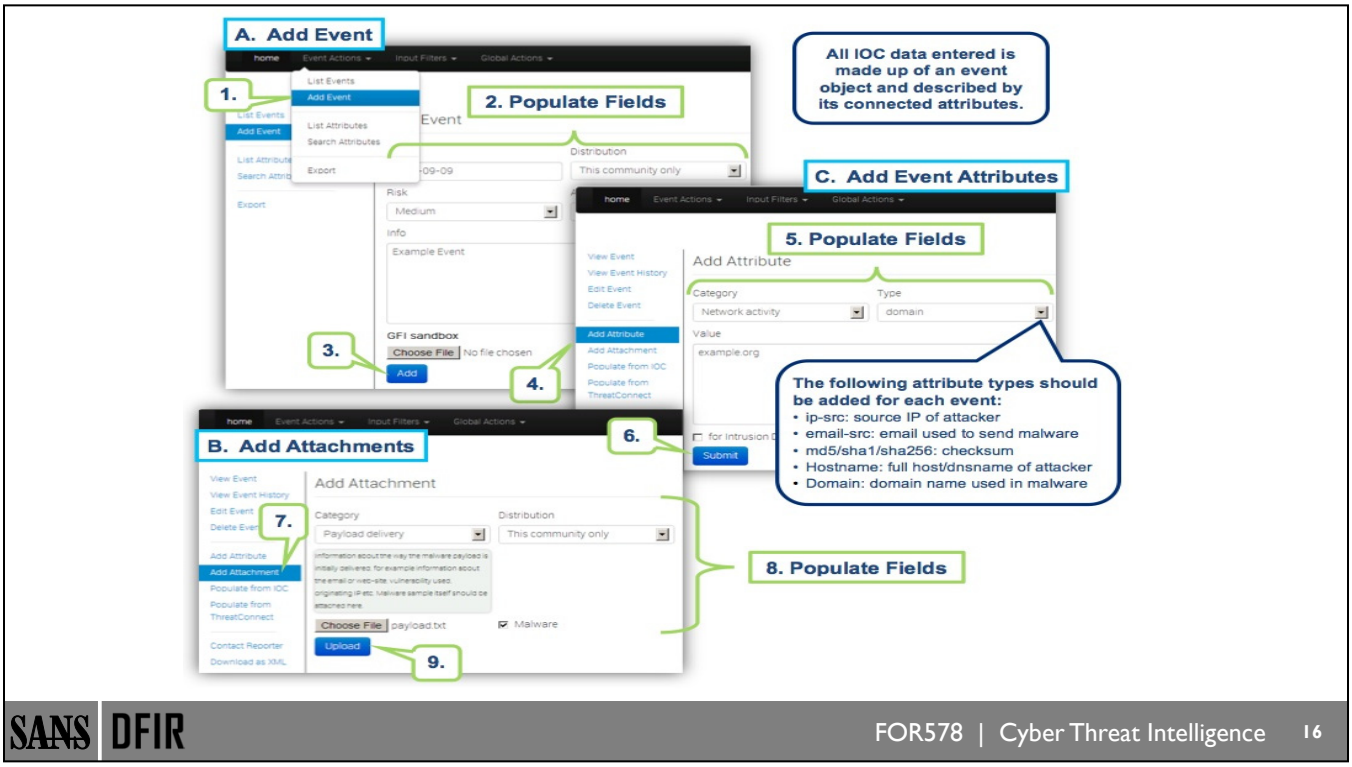

Creating an MISP Event

The MISP quick-start guide demonstrates an easy nine-step process across three phases (Add Event, Add Attachments, and Add Event Attributes) to adding an event into MISP. First of all, notice the focus on IOCs, which stresses the tactical and operational level value of the tool.

In the first phase (Add Event), the user selects the Event Actions dropdown menu and chooses Add Event. Then the analyst would populate the fields with information such as the date, the perceived risk level of the event being inserted, information about it, analysis level, and the desired distribution level of the event to determine who the information is shared with. The process ends with the analyst choosing Add. It is important to stress the analysis level. By being able to note if the analysis is in its initial phase or final phases, it is extremely helpful to categorize confidence in the IOCs and its static or dynamic nature. Early analysis of indicators often means that much of the information will change over time.

In the second phase (Add Attachments), MISP gives the option to add files such as the malware itself, accompanying files, phishing emails, etc. Once the information is populated and uploaded, the file is now available to others. One of the issues with many IOCs is the difficulty in verifying that they work correctly. By quickly being able to share samples with the IOCs themselves, it minimizes that problem if users appropriately take advantage of this feature.

In the third phase (Add Event Attributes), the analyst can enter attributes as identified by the popup in the graphic; this is an area to include indicators associated with the adversary such as their IP address, email address, or hash of the file. This is extremely important to be able to link indicators together across events and visualize a pattern (as shown in the next slide).

References:

- https://github.com/MISP/misp-book/tree/master/quick-start
- https://github.com/MISP/misp-book/tree/master/using-the-system

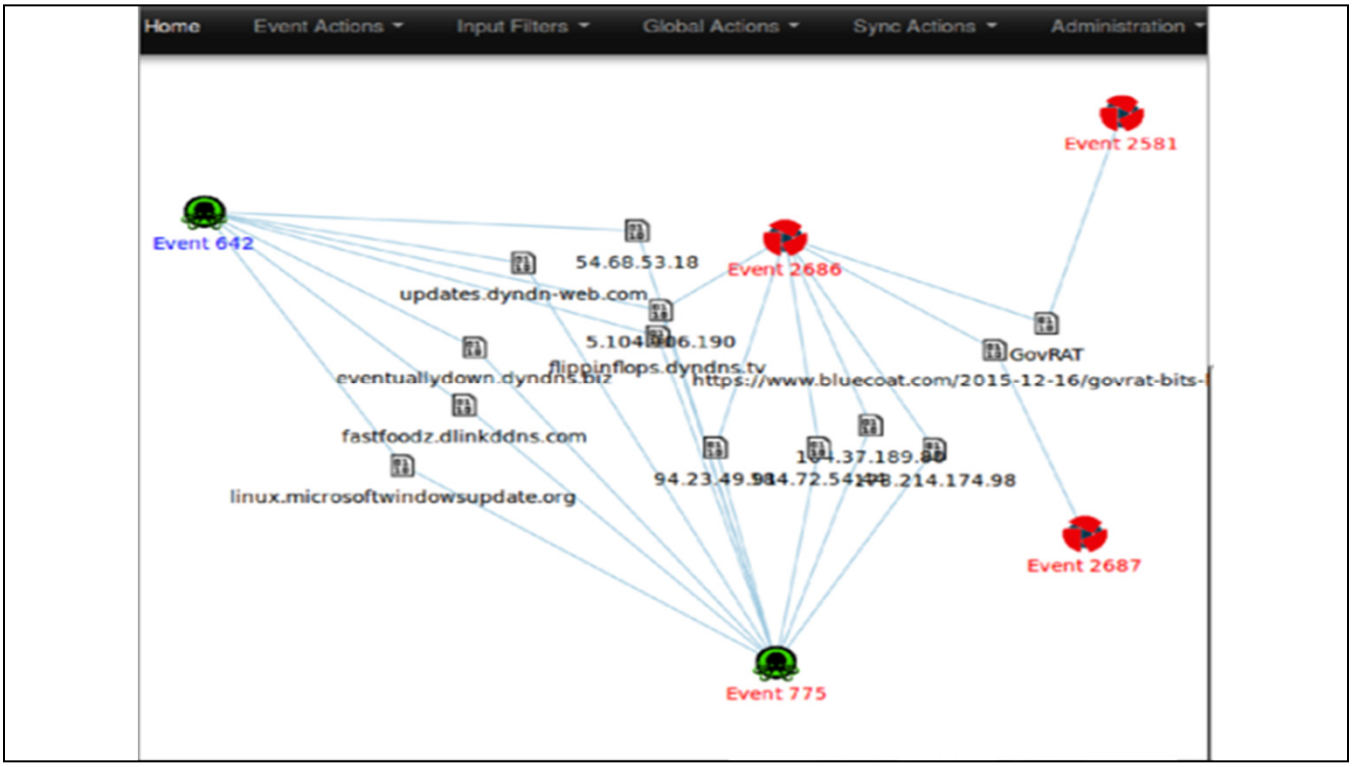

Visually Linking Indicators Between Events

In this sample graphic from the MISP team, it is easy to see the demonstration of visually identifying links between indicators. Visual correlation is particularly helpful in threat intelligence, as analysts often focus on the packets and granular information; being able to abstract the information and view it visually helps to identify patterns that would otherwise be missed.

In the previous slide, it was noted that indicators about the adversary could be entered with each event. They may not be used in the IOC itself, but it can be useful for the reason shown here: Linking information between multiple events. Adversaries will often use infrastructure not related to them: Russian adversaries using Chinese emails and IP addresses, US actors using Brazilian information, Chinese actors using Korean infrastructure and names, etc. (all hypothetical examples). But those choices are still chosen by humans. They can link events. They are the human fingerprint; analysts can and should be aware of past purely technical events.

Reference:

- https://raw.githubusercontent.com/MISP/MISP/2.4/INSTALL/screenshots/misp-panorama.png

Methods of Storing: Best Practices

There is a lot of valuable data in threat intelligence reports and in your information databases on threats and network characterizations. These are all valuable targets to adversaries and insider threats who might want to profit from the information. Be sure to set up your database in a secure manner that has unique authentication processes so that you can track who has access to the data and when. When the National Security Agency lost its data to Edward Snowden, some internal tracking processes helped it to identify *some* of the information he stole. Even the government-sized Intelligence Community is vulnerable to wishing for better practices after the fact—learn from their mistakes.

The method you employ for storing reports and data should also use a common format where your analysts submit reports following common formats and naming conventions. In addition, make sure the database and its resources are scalable for when your organization or team grows. Lastly, make sure to make the data shareable. You may not want to share it externally now but have a process you can implement when it is time for you to share with other partners—there will eventually come a time when you want to.

- Make a common format:
  - Ensure that your personnel store reports in a common format so that information can be quickly obtained easily even after a person leaves the organization.
  - Common formats should be for the reports themselves and the naming convention.
- Make the method scalable:
  - Internal servers such as SharePoint and SQL databases can be great tools for storing reports, but ensure that you can expand storage as needed.
- Make the method secure:
  - Threat intelligence reports are valuable resources of data that unauthorized users will want access to (inside and outside of your network).
- Make the method sharable:
  - Easy access to your users and network defenders is essential, as well as authorized third-party users; consider API access to deliverables such as the stored IOCs.

Lead-in to Exercise 4.1

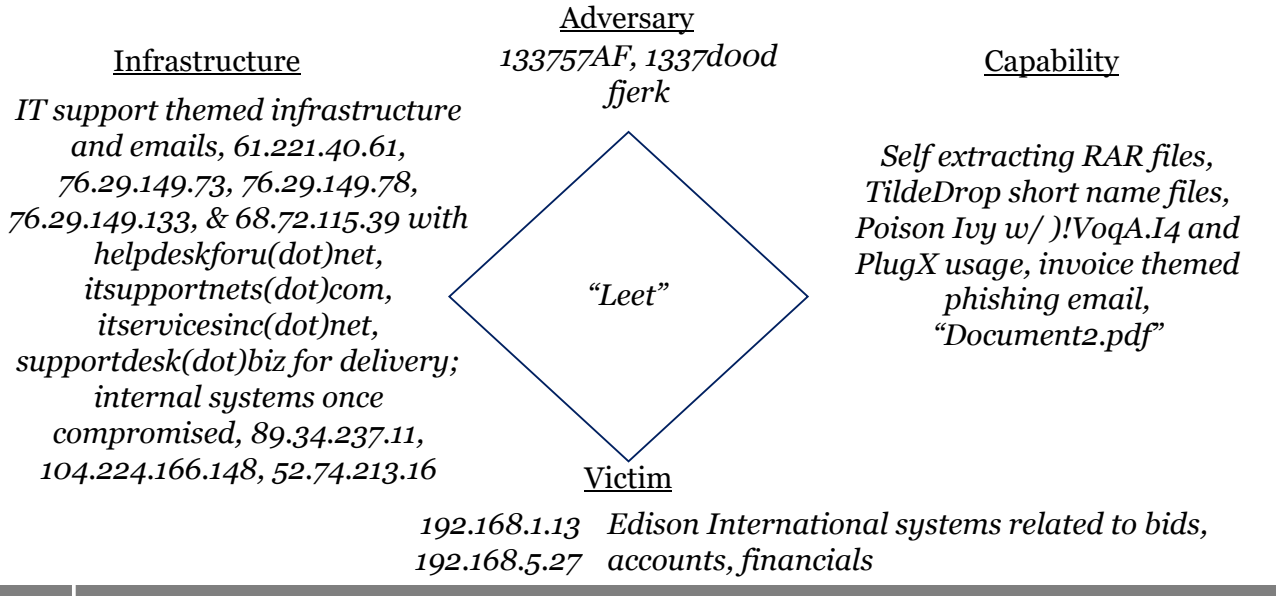

Lead-in to Exercise 4.1

We will take a subset of the "Leet" intrusion set and leverage these in Exercise 4.1. Our intent is to start storing the information for long-term usage, especially as this set of intrusions seems to be persistent against our organizations.

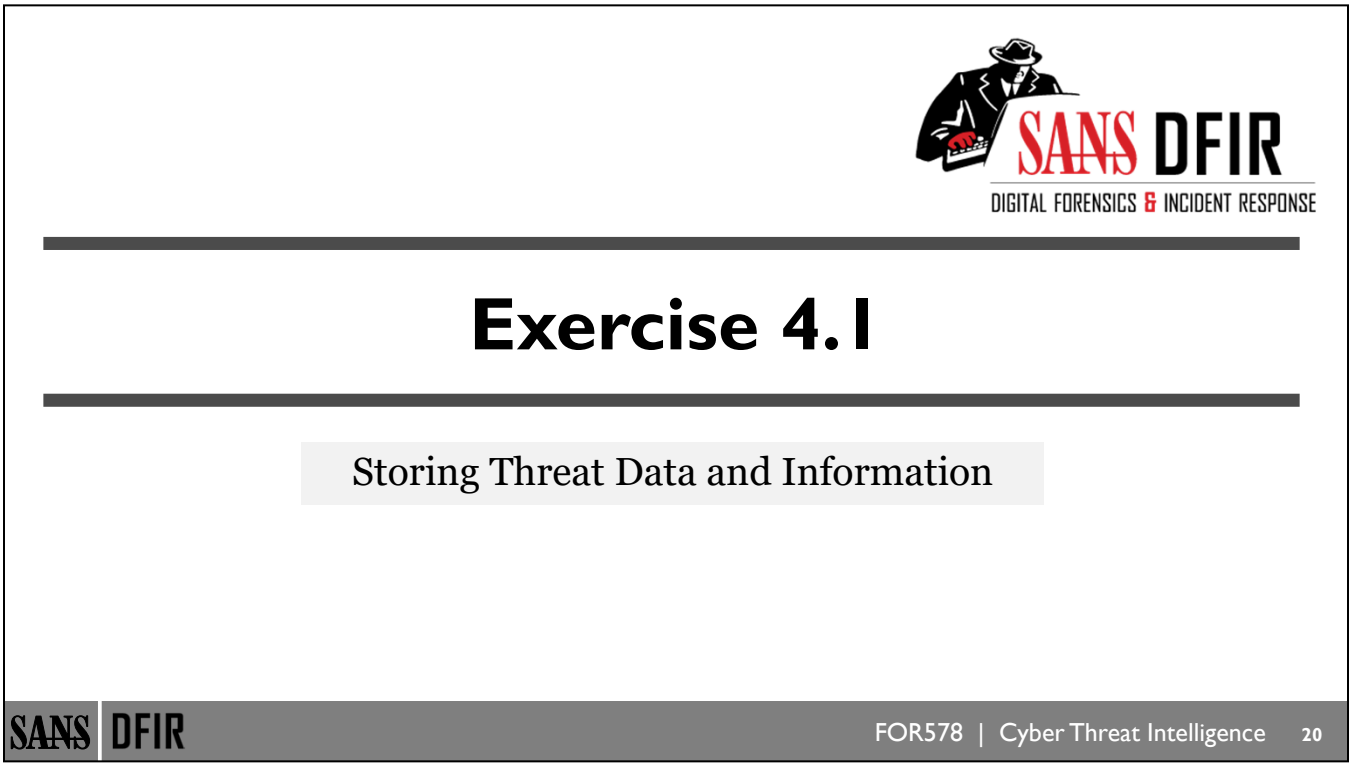

I tried to make this lab on my own and it really wasn't great; Katie, thank you for your contributions in making this lab a useful one for the FOR578 crew.

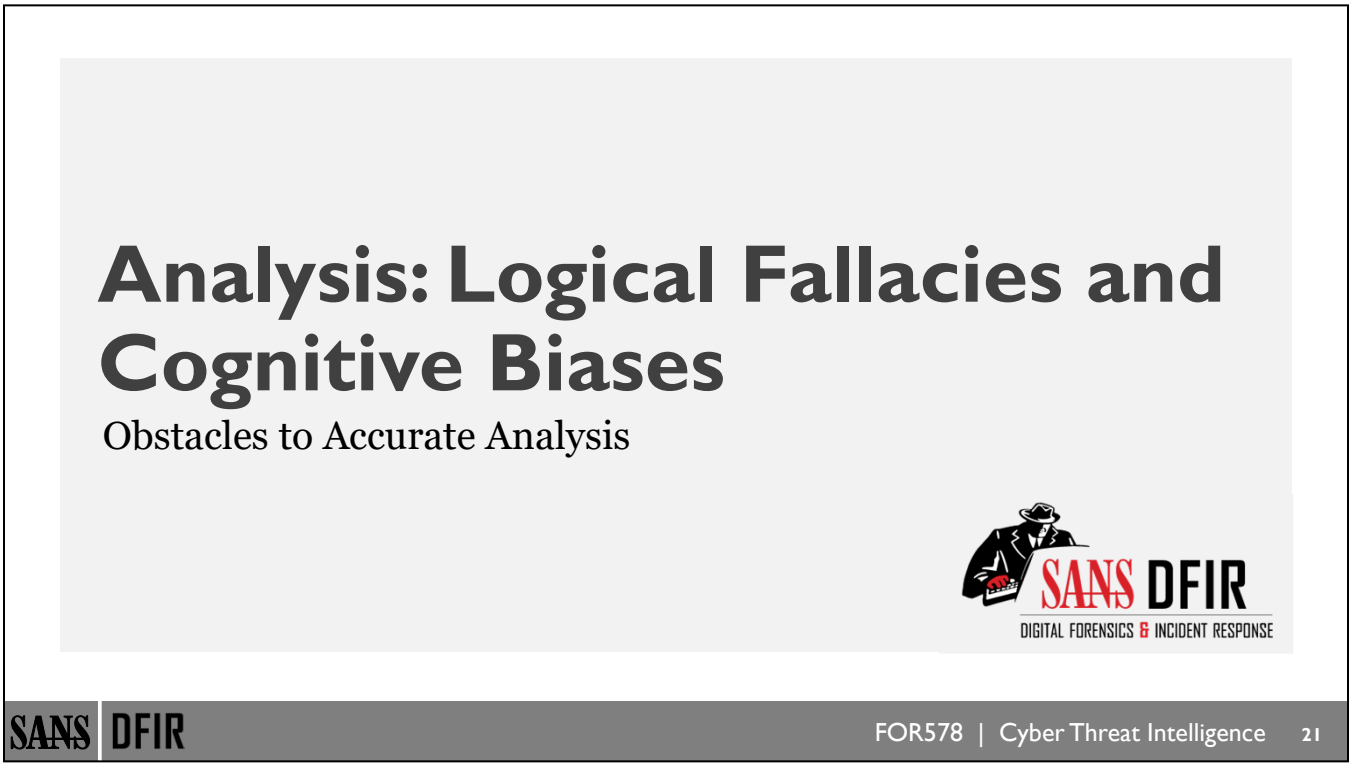

This page intentionally left blank.

Identifying and Defeating Bias

- All analysts have bias
- Analysis relies so heavily on the human mind and analyst understanding that bias poisons good analysis, especially at the strategic level and regarding attribution

Identifying and Defeating Bias

All analysts have bias. Where you are born, your political views, your religious views, your salary, your geopolitical affiliations, nationality, etc., all shape how you view the world. Understanding logical fallacies and cognitive biases help you avoid these issues.

Logical Fallacies

- Simply put, logical fallacies are flaws in reason
- Logical fallacies often (unfortunately) appear in cyber threat intelligence assessments

Logical Fallacies

Logical fallacies occur when arguments do not logically make sense. For example, good logic would state that if there are three pieces of evidence indicating an actor is U.S.-based but the analyst collected one piece of evidence stating that the actor was China-based, that the three pieces of evidence count for more (as long as they are of equal weight). But an anecdotal fallacy looks to use isolated cases or analysts' personal experience to encourage them to make a choice.

Likewise, just because something is likely the case, such as Russia being responsible for an intrusion, does not make it a logical choice. Logic would dictate that we need to fully analyze available information and make an assessment.

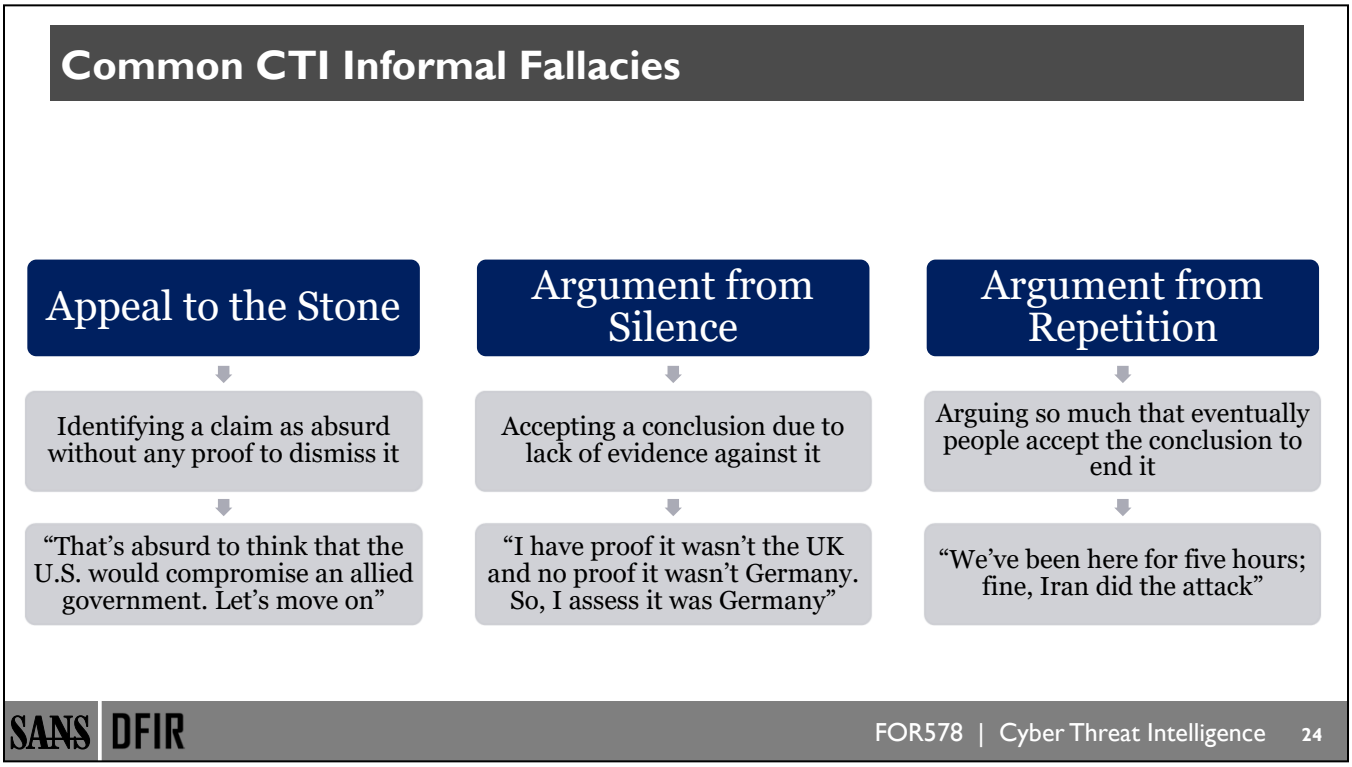

Common CTI Informal Fallacies

Informal fallacies occur when an argument does not support the conclusion. These are extremely common in CTI assessments. Appealing to the Stone as an example is a common informal fallacy where someone makes a claim and dismisses it as absurd without providing any proof. This is common among many analysts both intentionally, trying to move past an idea that they cannot prove or disprove, as well as unintentionally, not thinking outside the box enough.

Also, be careful of Argument from Silence or its near opposite, Argument from Repetition. Proof should never rely upon who provides the least proof or most proof; it should always require quality of evidence and analysis.

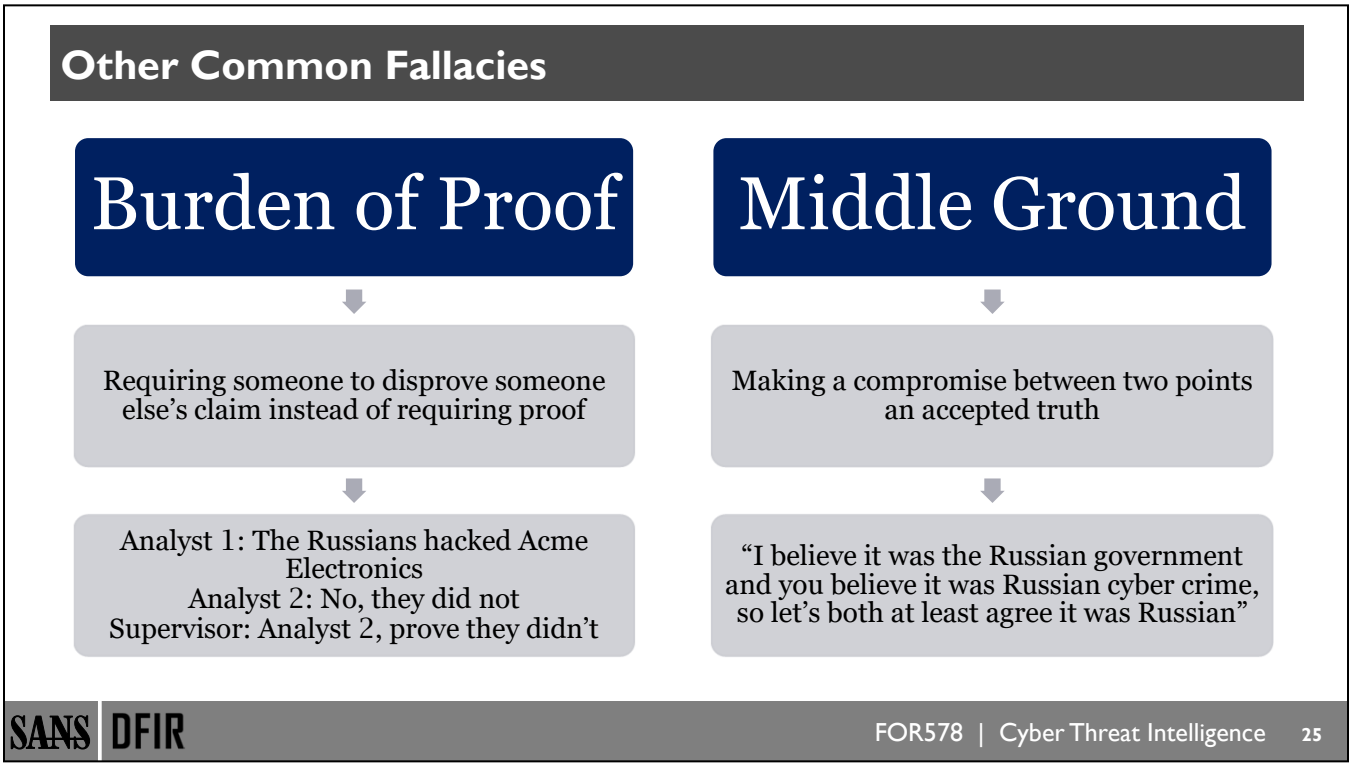

Other Common Fallacies

Other common fallacies to form, including the burden of proof fallacy and middle ground fallacy, commonly come up. Most of us have been in the situation where a supervisor, other analyst, or friend hears an argument from one person, or you relay it, and instead of asking for proof that the analysis is correct, they request you prove it did not. As an example, in our scenario, we have the Poison Ivy and PlugX variant that has targeted Acme Electronics. One analyst may assess that China is responsible. You may request evidence to support that conclusion and instead you are prompted to provide evidence that it's not China. "Why wouldn't it be China? Where's your evidence to support doubting his conclusion?"

Cognitive Biases

- Cognitive biases are constraints on how we as analysts think that influence incorrect decisions and assessments
- They allow analysts to create their own version of reality where inaccurate judgments and illogical interpretations occur

SANS DFIR FOR578 | Cyber Threat Intelligence 26

Cognitive Heuristics Introduce Bias

The human brain is a marvelous thing in its ability to adapt and lend greater understanding to our environment. However, the evolutionary leaps our species has undertaken over a million or so years that have enabled our remarkable progress also introduced biases in our daily lives that can adversely impact our goal of sound analytical reasoning. In addition:

"…the circumstances under which intelligence analysis is conducted are precisely the circumstances in which accurate perception tends to be most difficult" ["Cognitive Factors in Deception and Counterdeception." Donald C. Daniel and Katherine L. Herbig, eds., *Strategic Military Deception*, Pergamon Press, 1982]

Heuristics that we have developed to evolve to our environment are the basis for these biases ["Making Hard Decisions: An Introduction to Decision Analysis" Robert T. Clemen. Duxbury Press, Pacific Grove, CA. 1996, pp. 281–285]. There are many different cognitive biases—too many to cover in this course. Some of the most common biases we have seen in others and experienced ourselves are:

- Anchoring
- Confirmation bias
- Congruence bias
- Hindsight bias
- Illusory correlation
- *Cum hoc ergo propter hoc:* Correlative and Causal Confusion

We'll go in depth on a few of these because understanding them and recognizing them in your daily work is absolutely essential to your success as an analyst.

Understand that our biases are unavoidable. It is a consequence of our biology and evolution. As such, it is important to accept these errors will occur, and although they cannot be avoided, we can develop our own personal, individual approaches to managing these biases. To do so, we must be willing to see fault in our own reasoning. Any skilled analyst will tell you that hubris is antithetical to skilled analysis. Be honest with yourself, always treat analysis—even your own—with a skeptical eye—and be eager to find fault in your own reasoning so that it may be corrected. This isn't easy but will make you a talented analyst over time.

Mirror Image

- Common bias that forms when you analyze a situation, person, or entity with your context, background, experiences, etc., instead of theirs

tracted by the authorities had gained access to my Gmail account through an unsophisticated phishing scam targeting me and several people I knew.

THROUGH ALL OF MY ATTEMPTS TO EXPLAIN THAT THE IRANIAN AVOCADO QUEST WAS *NOT* SOME shadowy CIA mission dubbed "Project Avocado," one of the big sticking points for Kazem was the very concept of Kickstarter. He simply refused to acknowledge its existence. "Why would people give other people they don't know money to try to accomplish meaningless things?" he kept asking in the accusatory and rhetorical tone particular to those underinformed and self-righteous ideologues that Iran continues to produce year after year.

"Alan Eyre sent you to Iran, Jason," Kazem proclaimed. "The Great Judge can prove that."

"Please tell the Great Judge to prove it." I responded, but was

SANS DFIR FOR578 | Cyber Threat Intelligence 27

Mirror Image

One extremely common cognitive bias that is particularly hard to break is mirror imaging. It makes sense if you think about it. You are reviewing some situation or person or entity that you don't know well, and so your mind starts grabbing information and relatable findings that you do understand. I don't know how an Iranian operator would act, but I know how I would act in this situation, surely that's helpful right? Wrong.

Mirror images are when we fool ourselves into believing that the entity/person/etc. that we are analyzing would behave in any way similar to what we would given our experiences, biases, and life's context. We must analyze situations, people, and entities in their context not our own.

"If I were a Russian intelligence officer…"

Anytime you find yourself thinking or saying something similar, realize you're falling prey to the very common mirror imaging bias.

The snippet on the slide is from "Prisoner" which was written by Jason Rezaian who was held for 544 days in Evin Prison. His Iranian captors could not understand the concept of Kickstarter and that Jason's attempt to travel, which was funded by others, to get Iranian avocados—MUST be a CIA operation. Of course, it wasn't a CIA operation, but Kazem could not break from his world context to see Jason's. More importantly, this is a good example of Kazem's world context being so different than ours so when we analyze Iranian operations, as an example, we really have very little in common with them in context and world perspective to drive analysis thus forcing a deep need to understand them, not ourselves.

References:

- https://www.cia.gov/static/9a5f1162fd0932c29bfed1c030edf4ae/Pyschology-of-Intelligence-Analysis.pdf

Anchoring/Focusing

- Overvaluing one piece of information
  - "Anchored" on it
- Forces an analyst to be unable to seek new information or analyze competing information
- Anchored information is often the first information acquired

Anchoring/Focusing
Anchoring refers to beginning with an assumption or assessment and then adjusting one's assessment as new information becomes available, rather than taking the information as a whole for an assessment. The natural tendency with this bias is under-adjustment, which in turn fails to properly account for the possibility of what are perceived as extremely unlikely scenarios ["Making Hard Decisions: An Introduction to Decision Analysis." Robert T. Clemen. Duxbury Press, Pacific Grove, CA. 1996, pp. 281–285].

To put this in the context of competing hypotheses, you should not continue to operate with a set of (or single) hypothesis as new information becomes available. This results in "fitting," or tweaking, of the hypothesis that can result in errors in conclusion that might be more properly accounted for by starting over with a rethinking of all competing hypotheses in the context of new information.

In our world of CTI, new information rapidly emerges as analysis is conducted, which makes anchoring a bias that is difficult to avoid.

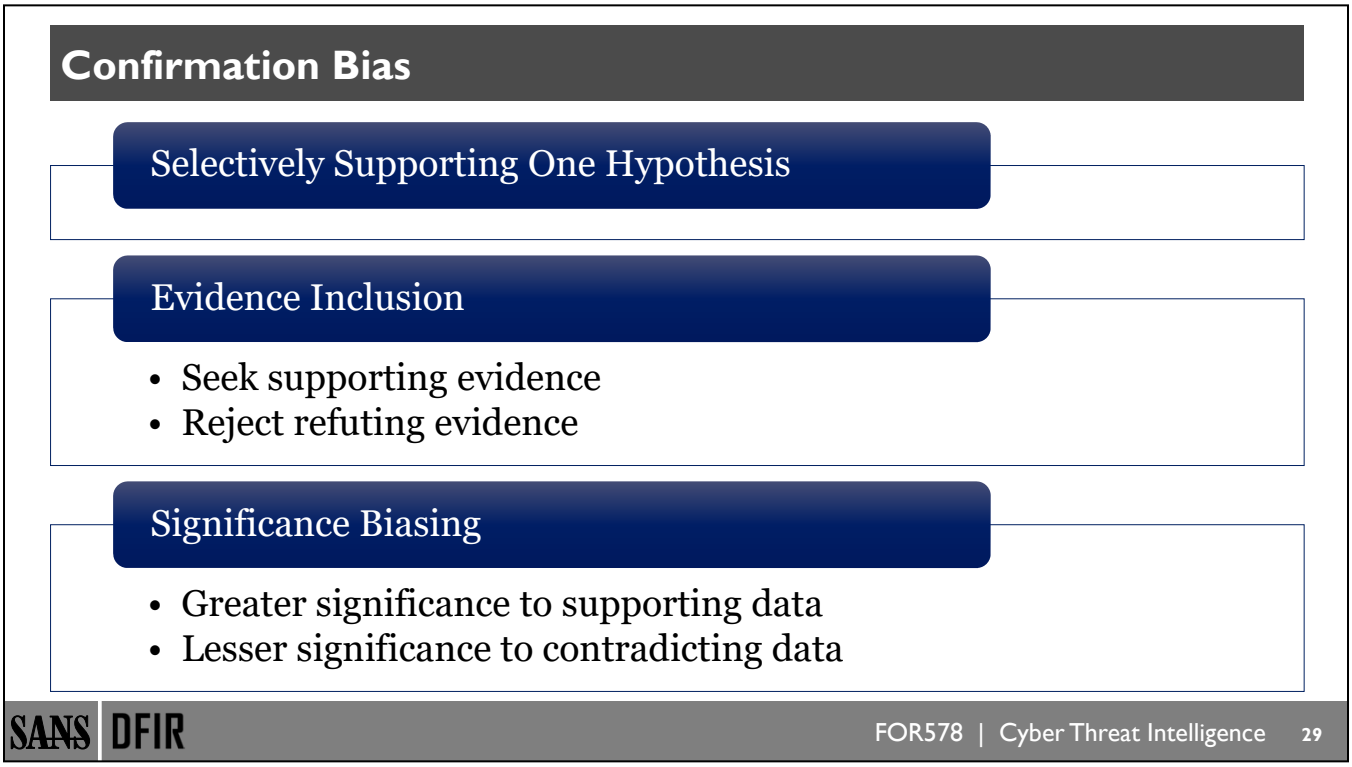

Confirmation Bias

Confirmation bias is fairly straightforward: This is the tendency to include or reject evidence based on its alignment to a preferred hypothesis. More subtly, confirmation bias also includes the tendency to ascribe more or less significance to evidence based on its support of a preferred hypothesis or outcome. This latter condition is common in our field and difficult to identify. The best way to avoid this bias is to fairly consider all hypotheses, regardless of their implications or perceived likelihood.

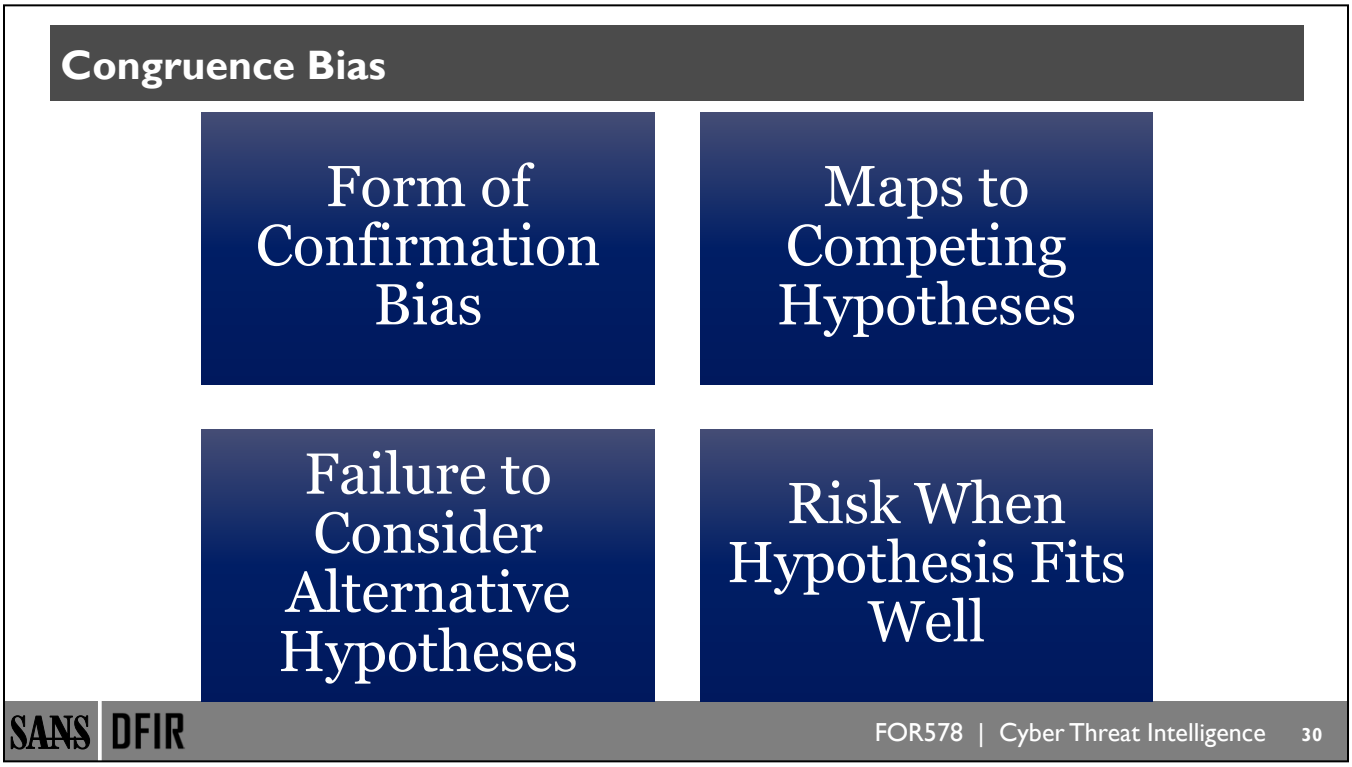

Congruence Bias

Congruence bias is related to confirmation bias but in a more abstract way. To put this into our intelligence lexicon, congruence bias is the failure to adequately present and test alternative competing hypotheses, and instead find different ways to present data that tests an existing hypothesis.

Now that you know about competing hypotheses, the congruence bias should be straightforward to identify and mitigate. That said, it is easy to fall into the trap of focusing on a single hypothesis, particularly when that hypothesis neatly fits or explains the intelligence you have on hand.

Reference:

- ["Heuristics and Biases in Diagnostic Reasoning II: Congruence, Information, and Certainty." Jonathan Baron, Jane Beattie, and John C. Hershey. Organizational Behavior and Human Decision Processes 43, 88−110. Academic Press Inc. 1988.]

Hindsight Bias

- "I knew it all along"
- Unlikely outcome seen as obvious
- Results in victim blaming
- Common in network intrusions
- Nation-state activity and APT intrusions are often difficult to predict and hindsight bias inappropriately simplifies that problem

Hindsight Bias

Hindsight bias is another self-explanatory bias. Hindsight bias is the tendency to see an unpredictable event as an obvious result of a set of conditions or parameters. It is important to bring up because it is *exceptionally* common for analysts looking at intrusions in a forensic capacity to disregard the overwhelming complexity of human behavior on computers and wonder how "anyone could have been so dumb as to let this happen." A significant effect of hindsight bias is victim blaming, which is currently endemic to nearly all network intrusions. The difficulty in managing and defending networks is extreme, and most often the reality of APT intrusions is that typically, it isn't reasonable to have expected the outcome that occurred given the information available to analysts and network managers prior to an intrusion.

Illusory Correlation

- Observe correlation when none exists
- Common when associating an unusual or extreme experience with all future experiences
- Stereotypes are the most common type of illusory correlation

Illusory Correlation

Illusory correlation refers to the tendency to observe a correlation between two observations when no such correlation exists, particularly when those observations are each relatively unusual. It is most commonly described as the basis for social stereotypes, but illusory correlation exists in intelligence analysis as well.

Illusory correlation can often be tied directly back to memory. Studies have shown that those with a higher load on their working memory experience more instances of illusory correlation. Psychology studies have shown that we tend to overestimate the importance of events we can easily recall and underestimate the importance of events we have trouble recalling. The easier it is to remember, the more likely we are to create a strong relationship between two things that are weakly related or not related at all.

Case Study: New York Stock Exchange (NYSE) Computer Glitch

- On July 8, 2015, the NYSE had to halt
  - A computer glitch was blamed publicly
- All United Airlines flights were grounded
  - An unknown computer problem was blamed
- The *Wall Street Journal*'s website then went down
- Illusory correlation:
  - These unusual events were seen as correlated under the concept of a cyber attack, which led to widespread concern
- Hindsight bias:
  - Individuals claimed they knew all along that it had been the WAN optimizer, and not a cyber attack. While the cyber attack was unlikely, it was also unlikely that a shared vendor between NYSE and UA caused issues
  - WSJ was DOS'ed by those rushing to it to see about NYSE and UA

Case Study: New York Stock Exchange (NYSE) Computer Glitch

July 8, 2015, the NYSE had to halt trading for a computer glitch, all United Airlines flights were grounded, and the *Wall Street Journal*'s website went down. Illusory correlation of these unusual observations led to concern of a widespread cyber attack, which was later disproven. United's outage was caused by a bad routing configuration that was propagated across all the company's networks; the NYSE trading system crashed when a recent database upgrade was placed under stress during trading hours; and the *Wall Street Journal*'s website subsequently became unavailable due to the overwhelming load placed on it by individuals trying to learn what was happening with the NYSE.

Cum hoc ergo propter hoc

- Confusion of correlation and causation is common
- Causation
  - **Is** correlation
- Correlation
  - **Often isn't** causation

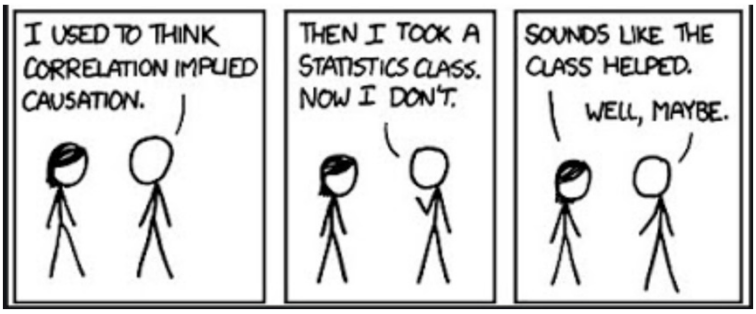

Cum hoc ergo propter hoc

Cum hoc ergo propter hoc ("With this, therefore because of this.")

One of my favorite phrases is "correlation is not causation." Inexperienced analysts often confuse the correlation of two events (these two things are related or can be mathematically correlated) with a causal relationship. It's true that two events that have a causal relationship will be correlated, but remember that it is much more difficult to establish causation than correlation. I think the most concise quote available on causal and correlative confusion is this:

Two events occurring in close proximity does not imply that one caused the other, even if it seems to make perfect sense. ["Causation vs Correlation" https://senseaboutscienceusa.org/stats/ George Mason University Statistical Assessment Service. Retrieved 4 October 2014].

In short, causation implies correlation, but correlation does *not* imply causation.

Image:

xkcd

Case Study: Turkey Pipeline Explosion

- A 2008 explosion occurred at the Baku-Tbilisi-Ceyhan (BTC) pipeline
  - Turkey blamed extremists and the extremists took credit
- In 2014, the news organization Bloomberg claimed that Turkey and the extremists were wrong and that the event was caused by a Russian cyber attack
- Incident responders at the time had found malware in the control center
- Attribution to Russia was based on Russian IP command and control
- Cum hoc ergo propter hoc:
  - The IR team claimed that because the explosion occurred, and because the malware was present, that therefore the malware caused the explosion

Case Study: Turkey Pipeline Explosion

In 2014, Bloomberg reported on the Baku-Tbilsi-Ceyhan (BTC) explosion of 2008. At the time, the Turkish government claimed the attack was a physical terrorist attack by extremists. The extremists came forward and claimed attribution of the attack. However, Bloomberg reported that the attack was actually due to a cyber attack by Russia. The story of the Russian cyber attack has been debunked (see the references below if you are interested); however, one thing observed by the incident responders was useful to the discussion of cum hoc ergo propter hoc.

According to personal interviews, the incident responders had discovered malware in the control center of the gas pipeline. The control center is a central point of communications for these types of operations and relies upon traditional Windows-based systems. The incident responders identified the malware as being of Russian origin (it is not known if the Russian origin was meant to be nation-state, cyber crime, or just Russian malware). The incident responders fell prey to the bias of believing correlation meant causation; they were called in after the BTC pipeline's explosion, and once they found malware on the network, they correlated the two events. No analysis was ever published, nor were samples presented to indicate that the malware had the capabilities required to cause a physical explosion at the pipeline, a rare feat and the subject of immensely targeted code. However, we now know that the attack had nothing to do with any "cyber" event. The incident responders believing that because there was an explosion, and because there was malware, that the malware must have been involved with the explosion was classic correlation does not equal causation bias.

References:

- https://www.bloomberg.com/news/articles/2014-12-10/mysterious-08-turkey-pipeline-blast-opened-new-cyberwar
- https://ics.sans.org/media/Media-report-of-the-BTC-pipeline-Cyber-Attack.pdf
- https://www.sans.org/blog/closing-the-case-on-the-reported-2008-russian-cyber-attack-on-the-btc-pipeline/

Bias and Experience

- Security personnel often use their experience as useful bias to quickly come to conclusions
- Security personnel are often paid for their well-formed biases
- Cyber threat intelligence analysts are often paid to defeat their biases

Bias and Experience

It is important for cyber threat intelligence analysts to be able to identify the difference between bias and experience. Many analysts and security personnel are paid for the experience they bring, often in a way that could be perceived as bias. For example, an incident responder that has seen dozens of cases where an adversary has compromised the VPN and then moved to get administrator credentials on the domain controller is going to be more adept in a case due to that experience. He or she will have a certain set bias to moving to the domain controller quickly after identifying an adversary that has compromised a VPN into the environment, if that was what the responder has observed over and over again.

Cyber threat analysts, though, get paid to defeat their bias. They may have seen a piece of malware being used in campaign after campaign associated with Chinese-based actors. The next time they see that malware used, they cannot jump to the conclusion that the actors behind the campaign are Chinese based. The analyst should recognize this bias and apply proper practices such as Analysis of Competing Hypotheses, which will be explored in the next section.

Image:

Deadpool; all rights owned by Marvel

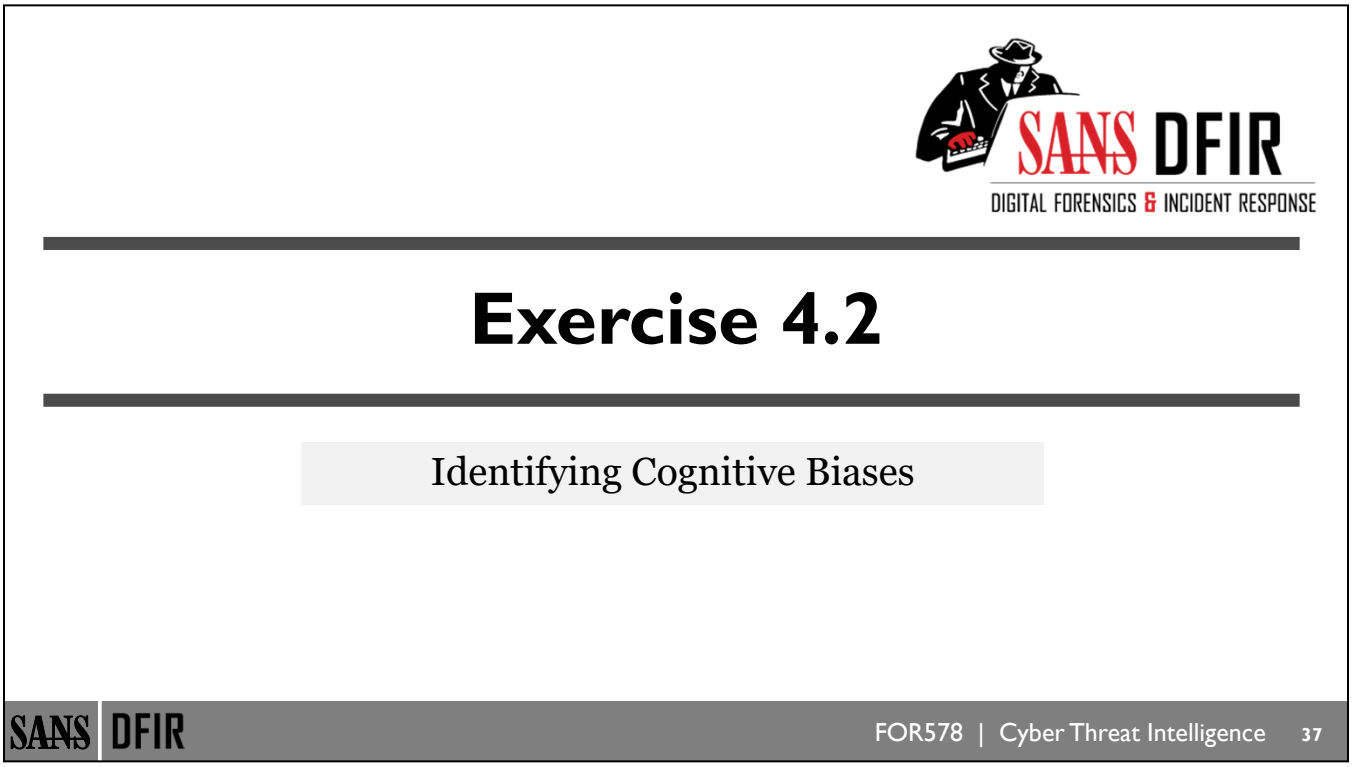

This page intentionally left blank.

Analysis of Competing Hypotheses

An Analytical Process by Former CIA analyst Richards J. Heuer, Jr.

This page intentionally left blank.

Analysis of Competing Hypotheses

- Developed by Richards Heuer, Jr., a 45-year veteran of the CIA
- Method for evaluating hypotheses and choosing the best one based on observed data while mitigating biases
- Seven basic but time-consuming steps:
  1. Hypothesis
  2. Evidence
  3. Diagnostics
  4. Refinement
  5. Prioritization
  6. Sensitivity
  7. Conclusion and Evaluation

Analysis of Competing Hypotheses

The basic premise of the Analysis of Competing Hypotheses is a structured method to identify all potential hypotheses, collect all the evidence, compare the evidence with the hypotheses, and then rank hypotheses and identify a potential best choice. It's also important to identify hypotheses (once already into the process) that do not make sense and to identify any potential pitfalls in analysis and evidence that exist. The end product should result in the best choice and a solid evaluation driven by evidence and facts instead of biases.

ACH Process Steps

A fantastic document that I recommend everyone in the class read is *Psychology of Intelligence Analysis* by Richards J. Heuer, Jr. The book is an easy read and discusses in detail a number of aspects of intelligence analysis from the perspective of a former CIA analyst. The book was published by the CIA's Center for the Study of Intelligence in 1999 [Heuer, Richards J. *Psychology of Intelligence Analysis*. CIA Center for the Study of Intelligence. https://www.cia.gov/static/9a5f1162fd0932c29bfed1c030edf4ae/Pyschology-of-Intelligence-Analysis.pdf Retrieved June 18, 2015].

In this book, Heuer describes one aspect of analysis as *Analysis of Competing Hypotheses*. As analysts, when we are given intelligence and asked to make an assessment based on it, we essentially create a number of different hypotheses and then choose from among them that which we believe the most valid. It is important that we are cognizant of this process, and exhaustively and fairly assess all hypotheses so as to reduce the amount of bias in the conclusion we draw. Heuer outlines a seven-step process for analysis of competing hypotheses:

1. Enumerate all of the possible hypotheses.
2. Support: Seek supporting and refuting evidence for each hypothesis.
3. Compare the evidence for and against each hypothesis as more or less helpful in determining the most valid hypothesis. Build this as a matrix of hypotheses and evidence.
4. Refine the matrix by removing evidence that has little value in determining the most valid hypothesis.

5. Prioritize the hypotheses by their relative likelihood; build this list by seeking additional evidence refuting them.
6. Dependence: Determine the degree to which your conclusion relies on a small amount of evidence, and consequences of that evidence being invalid, misinterpreted, or misleading.
7. Report your conclusions, including all competing hypotheses and their comparison.

(Additional) Identify future circumstances under which the conclusion reached might change; if assumptions are proven incorrect, factual data ends up being temporally bound, etc.

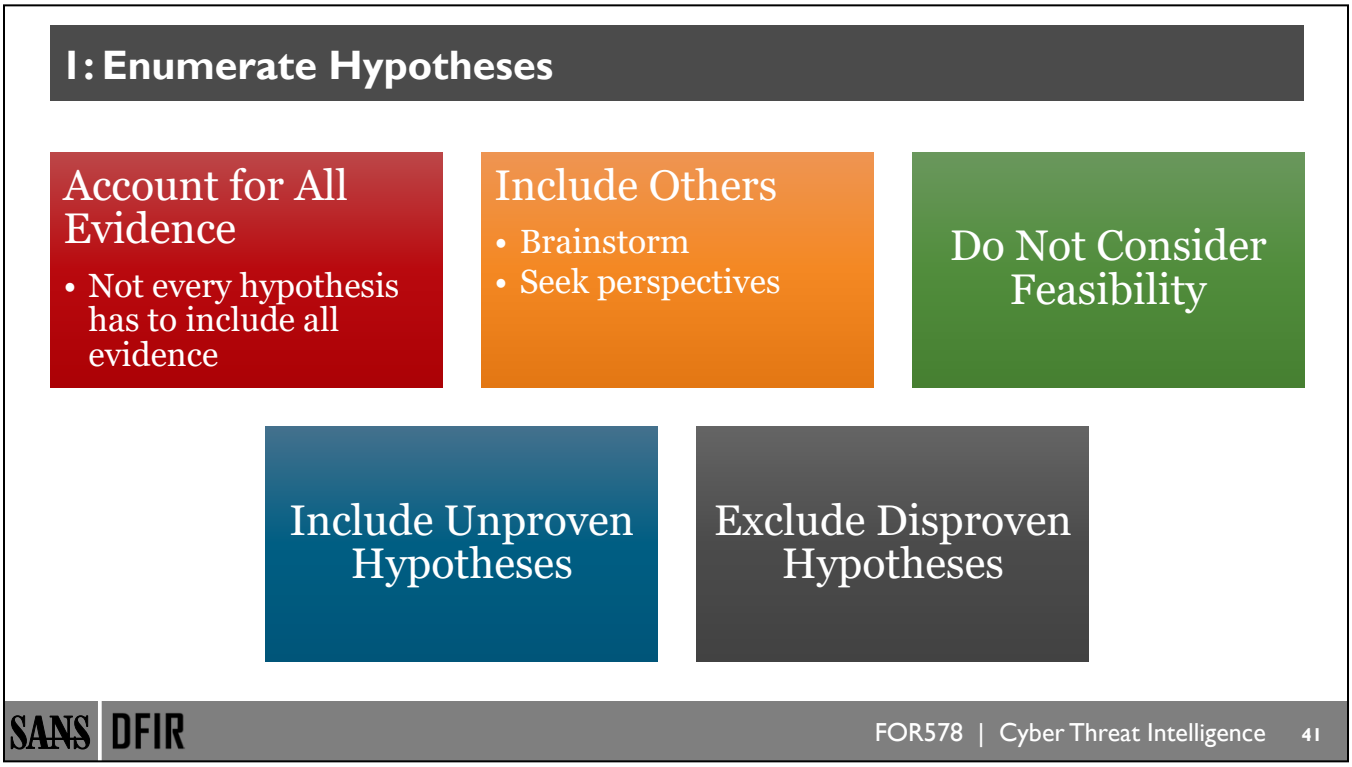

Enumerate Hypotheses

The first step in analyzing competing hypotheses is to develop the hypotheses themselves based on the available intelligence. As new intelligence becomes available, *all* hypotheses should be re-evaluated. (We discuss why in our discussion on cognitive biases). Create as many hypotheses as necessary to ensure inclusion of all the available evidence, even if you cannot fit all the evidence into a single hypothesis. Include others in the development of your hypotheses; take particular care to brainstorm with those who can bring a variety of perspectives. Do not yet consider feasibility in the formulation of your hypotheses. Exclude only hypotheses for which evidence exists that preclude the possibility of them being valid. Heuer helpfully distinguishes between *disproven* and *unproven* hypotheses thusly:

For an unproven hypothesis, there is no evidence that it is correct. For a disproven hypothesis, there is positive evidence that it is wrong. (p. 98)

Remember that as much as scientists love to distinguish between science and art, there is an art to scientific evaluation. Part of that art is in the formulation of hypotheses. Heuer offers this advice in determining how many hypotheses are appropriate:

The greater your level of uncertainty, or the greater the […] impact of your conclusion, the more alternatives you may wish to consider. (p. 98)

2: Support the Hypotheses

- Seek additional evidence:
  - Supporting
  - Refuting
- Include as evidence:
  - Deductions
  - Assumptions
- Discuss missing evidence

SANS DFIR FOR578 | Cyber Threat Intelligence 42

Support the Hypotheses
Although this section is concisely titled "Support the Hypotheses," this includes seeking evidence and making arguments that both support *and* refute the hypotheses developed in step 1. Although evidence is powerful in this step, one should not be limited to evidence alone. In the context of this activity, assumptions and deductions serve as "evidence," as they dictate the outcome of the process. Assume each hypothesis is true, noting which evidence supports it and which pieces of evidence are expected but missing. Discuss why expected evidence is missing, and note it as such. Do not overly focus on the presence of evidence and sacrifice consideration of its absence.

3: Diagnostics

Do FOR578 Students Pay Attention?

1. Students pay attention
2. Students do not pay attention
3. Students are not even present
4. There are no students; the cake is a lie

Analysis

	H1	H2	H3	H4
E1. Students on Facebook	-	+	-	-
E2. 80% of FOR578 students pass the certification	+	-	-	-
E3. There are some empty chairs	-	+	+	+
E4. Students are asking questions	+	-	-	-
E5. The hotel serves us snacks	0	0	0	-
E6. Students are maintaining eye contact	+	-	-	-

SANS DFIR FOR578 | Cyber Threat Intelligence 43

3: Diagnostics

To compare the available evidence gathered from the last step, Heuer recommends building a matrix of hypotheses (across the horizontal) and evidence (down the vertical) collected thus far. Use this matrix to determine which data points are the most helpful in assessing the likelihood of the presented hypotheses. To do this, consider each piece of evidence at a time, and assess the degree to which it supports or is consistent with each individual hypothesis.

The point of this step is to assess the degree to which each piece of evidence is diagnostic in determining the relative likelihood of the hypotheses. Evidence that supports, or does not support, all the developed hypotheses to the same degree is not helpful in a diagnostic sense for determining which is the most likely, no matter how interesting the evidence may be. We say in this case that **analysis proceeds horizontally, across the hypotheses, for each piece of evidence individually.**

In the matrix shown here, "+" and "-" indicate supporting and not supporting, whereas you could use "++" and "--" to indicate strongly supporting and strongly not supporting, respectively.

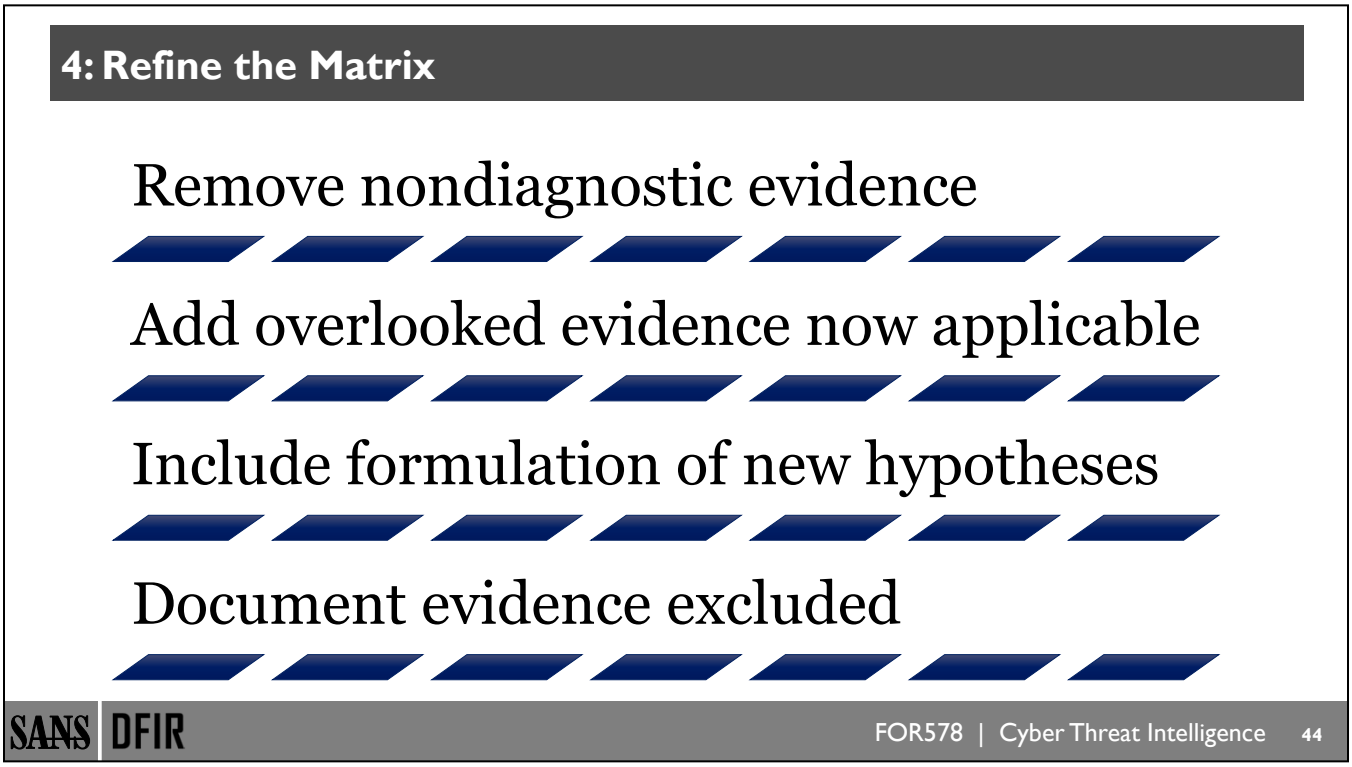

Refine the Matrix

At this point, it should be clear which evidence is not helpful in determining the relative likelihood of the hypotheses. Remove this evidence from the matrix, and then review the matrix. Sometimes, this process results in the identification of new pieces of evidence mistakenly excluded from the process earlier. Add this evidence in. The removal of evidence with no diagnostic value may also result in the formulation of new hypotheses. Add these hypotheses in as well. Be sure you document the evidence removed from the matrix so that your assessment can be reproduced should it be later questioned or found to be invalid.

5: Prioritize the Hypotheses

Do FOR578 Students Pay Attention?

1. Students pay attention
2. Students do not pay attention
3. Students are not even present
4. There are no students; the cake is a lie

Analysis

	H4	H3	H2	H1
E1. Students on Facebook	-	-	+	-
E2. 80% of FOR578 students pass the certification	-	-	-	+
E3. There are some empty chairs	+	+	+	+
E4. Students are asking questions	-	-	-	+
E5. The hotel serves us snacks	-	0	0	0
E6. Students are maintaining eye contact	-	-	-	+

SANS DFIR FOR578 | Cyber Threat Intelligence 45

Prioritize the Hypotheses

In this step, the hypotheses in the matrix are evaluated vertically, considering each hypothesis (rather than each piece of evidence) for its relative likelihood based on the evidence presented. Think of this as the hypotheses competing for your preference, selling themselves based on the evidence in the matrix when compared with one another.

Always start by looking for pieces of evidence that reduce the likelihood of certain hypotheses, evidence that seems to preclude (but not necessarily outright reject) hypotheses. These will be at the bottom of your priority list. By taking this reject-first approach, you can manage confirmation bias. Consider supporting only evidence after all disproving evidence has been used to prioritize the list, with hypotheses mapping to the most disproving evidence obviously prioritized the lowest.

In this step, we say that analysis proceeds vertically, looking at the total evidentiary support for each hypothesis. The hypothesis with the most contradicting evidence, H4, ends up listed first, followed by H3 and H2. H1 now becomes our top hypothesis: That the students are paying attention.

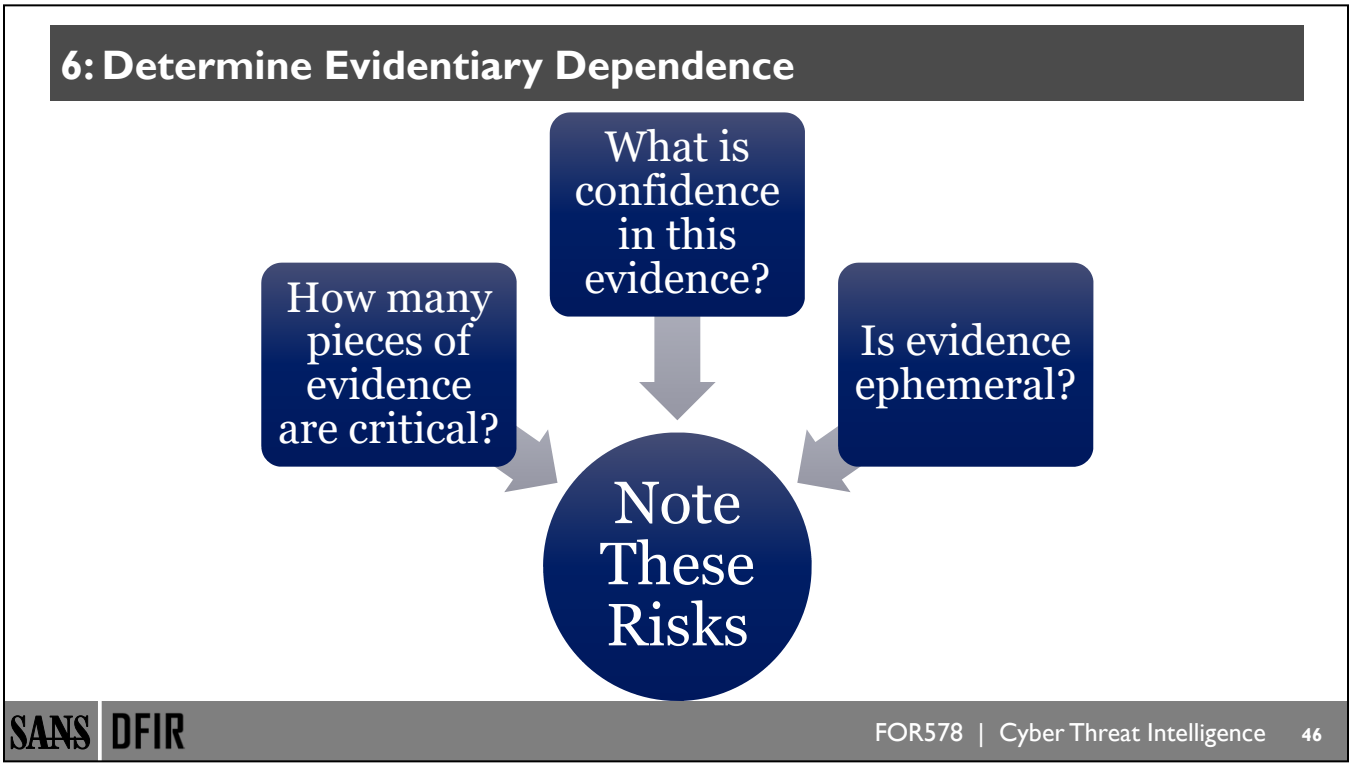

Determine Evidentiary Dependence

Now that your hypotheses are prioritized, look at the evidence most significant in the prioritization: Are one or a small number of pieces of evidence critical in the prioritization? If so, what is the level of confidence that this evidence is accurate? Are there assumptions underlying the evidence that need to be reconsidered? Might the evidence change in time? Note these assumptions and evidence as significant for inclusion in your final assessment.

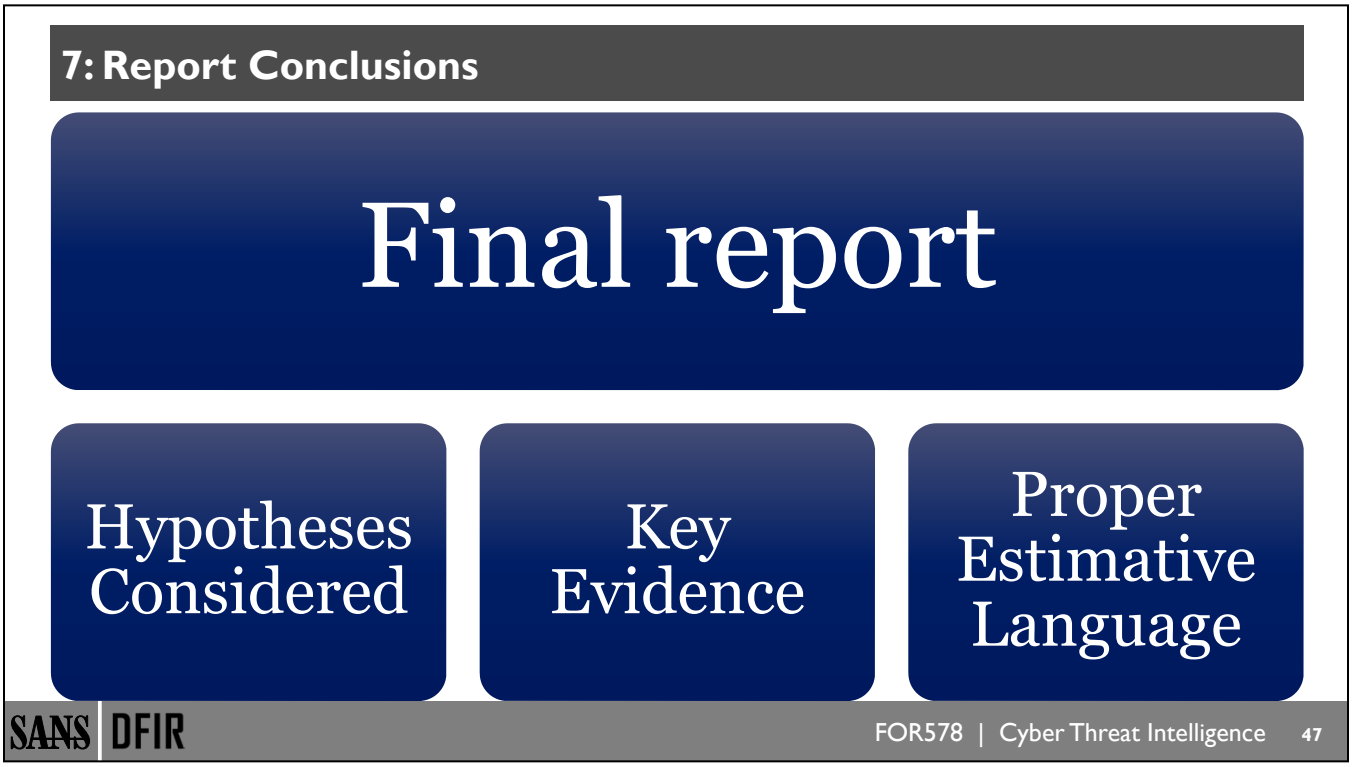

Report Conclusions

When reporting your conclusions, be sure to include the hypotheses considered and the most important pieces of evidence in your conclusion. Be sure to include a discussion about key evidence if only a few pieces of evidence were highly instrumental in your decision. Properly qualify your assessment using clear language, but do not attempt to enumerate probabilities if they cannot be calculated. ("I'm 90% confident this is right" is misleading because it suggests precision when there is none if that number was not calculated but guesstimated.) Properly use estimative language.

Identify Milestones

Analytical conclusions should always be regarded as tentative
—Heuer, p. 107

- Evidence may change in time
- Changes may affect outcome
- Note circumstances under which evidence may change
- Note how changes would affect conclusions

Identify Milestones

Analytical conclusions should always be regarded as tentative (Heuer, p. 107).

As more evidence becomes available, the facts of the existing evidence change, other circumstances pass, or the conclusions you draw may change or become invalidated. These must be identified and called out so that not only your conclusion is properly qualified, but it is also clear to your audience the circumstances in which the assessment would change.

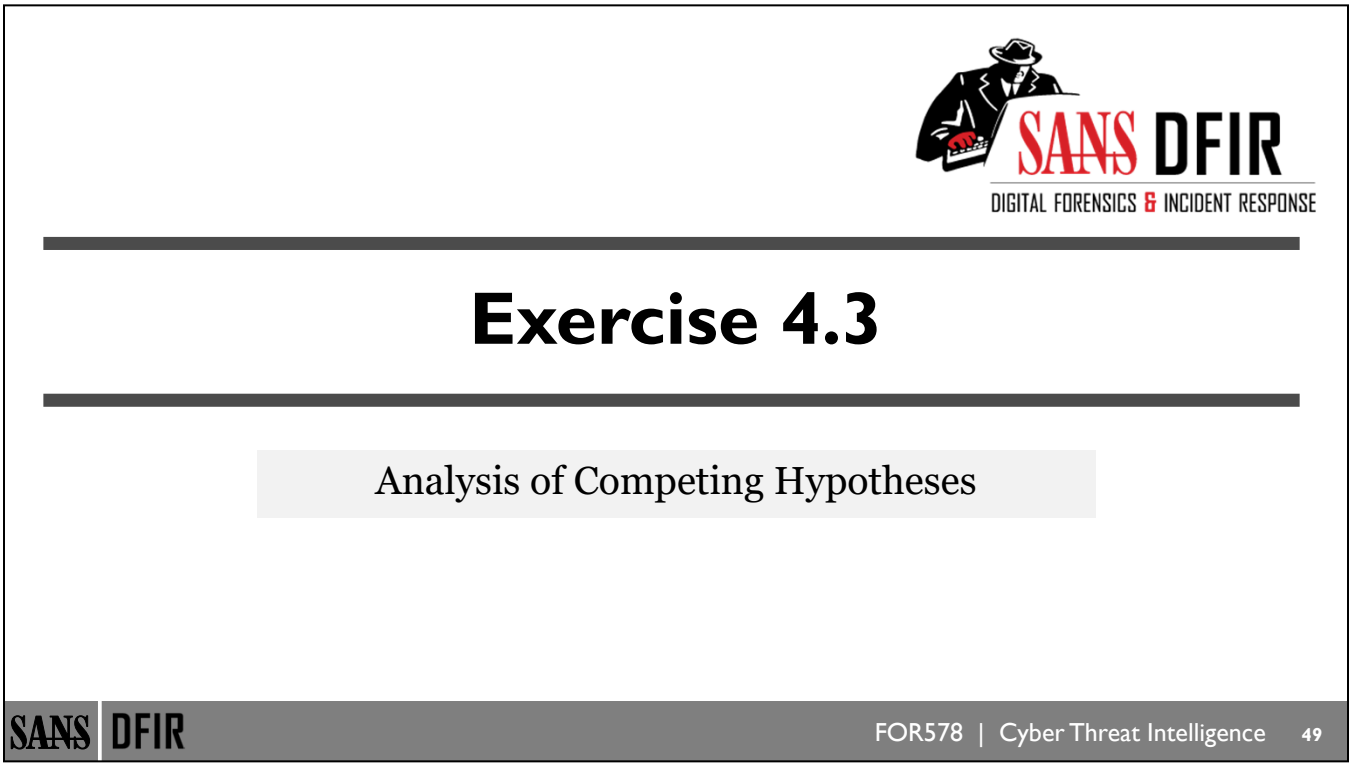

This page intentionally left blank.

Analysis:
Different Types of Analysis

This page intentionally left blank.

Leveraging Different Types of Analysis

Know Thyself

- Everyone has a favorite type of analysis for given situations
- Learn what analysis types facilitate your process

Know the Team

- Learn your team members' analysis types
- Ensure your tools and approaches play to everyone's strengths

Inject New Approaches

- Try new types of analysis, especially on critical cases
- Ensure you do not only leverage one type of analysis

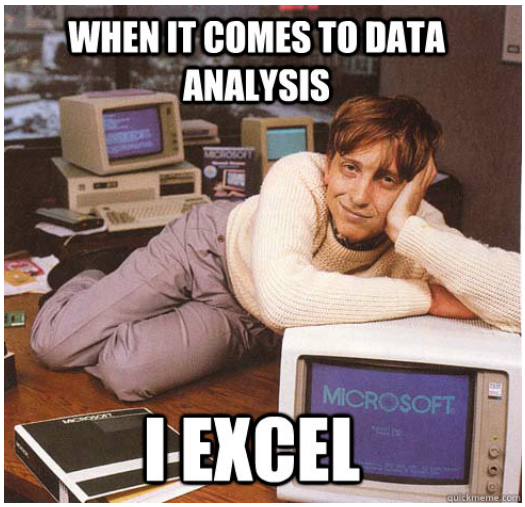

Leveraging Different Types of Analysis

Analysts all respond differently to different types of analysis and inputs. You should make a conscious effort to learn what you respond well to and what you do not respond well to; as an example, do you really find visual analysis useful? Your tools and efforts should likely help you complement that. Every now and then, you might try injecting different types of analysis, like to make sure you do not overly rely on one type and hinder your approach. Additionally, you should learn the types of analysis on your team and make sure that you are all able to work effectively together by focusing on each other's strengths.

Link Analysis

- Analysis of relationships between data points
- Visualization tools support analysis of large datasets:
  - Can also be used to represent smaller relationships
- For maximum effectiveness:
  - Entities need sources
  - Links need context (that is, domain resolved to IP at a certain date/time)

Link Analysis

The previous two charts demonstrate a fundamental analytical method referred to as link analysis. Specifically, we demonstrated pivot link analysis, where a pivot is performed around each entity within the graph to gain additional associations for each. Link analysis essentially refers to two entities that are related to one another by some data point. In the previous chart, you can notice that the second-level registered domains are all related to the email account cpyy.chen@gmail.com. The links supporting this association indicate that cpyy.chen@gmail is the registrant of each of the linked domains. This graph was assembled manually to serve as a visual aid to express the data in a format that is more digestible for most humans. However, link and data visualization tools can scale up to display large datasets that the tool can access.

Whether displaying relatively small amounts of data or large amounts of data, it is critical to verify that entities have sourcing information ("Where did this come from?") and that each link should provide context to the relationship ("How are the two entities related?"). In some cases, arrows can help to convey the directionality of the relationship between two entities or perhaps which entity is subordinate to the other. This might seem like a trivial concept, but this small nuance can affect how different display configurations, such as the hierarchical view, realign your graph.

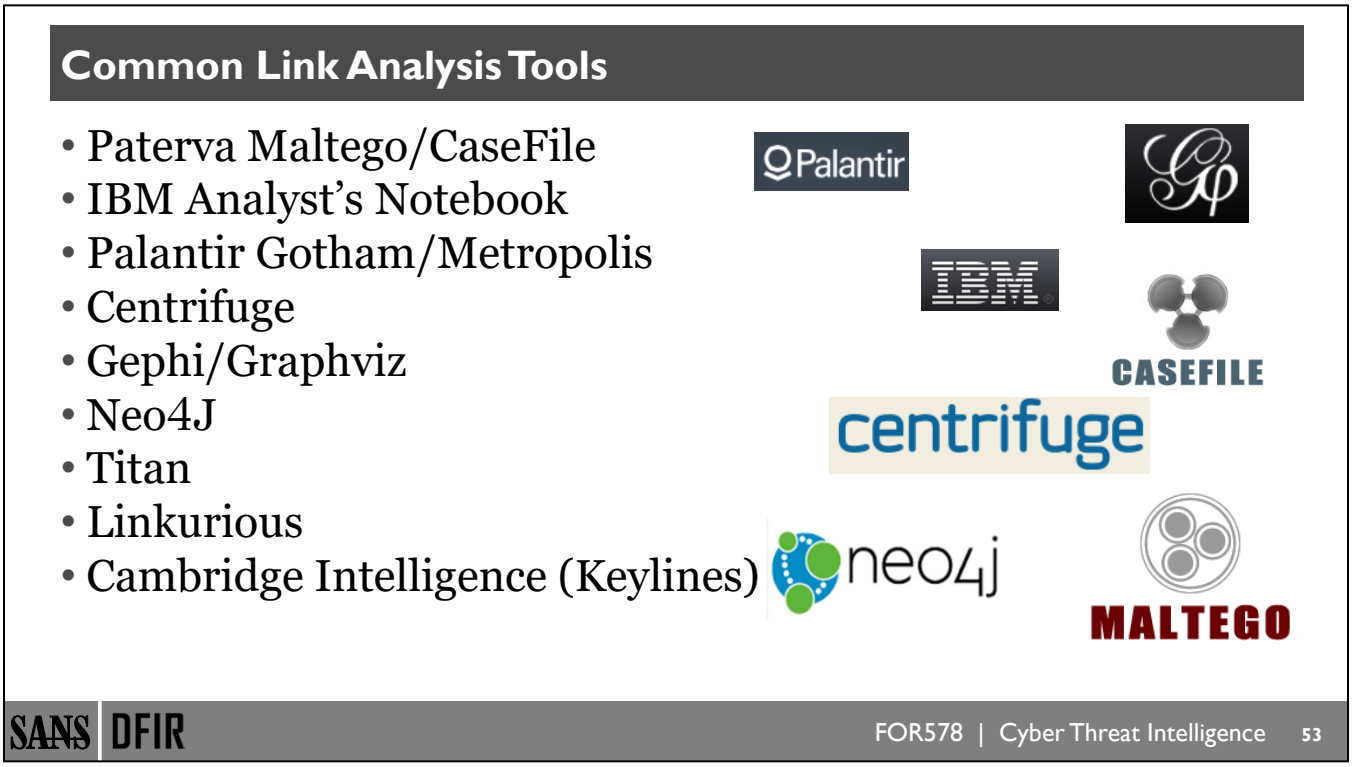

Common Link Analysis Tools

These are some of the most common visualization and analysis tools used today. Although specific capabilities vary between software vendors, some commonalities include the capability to graphically display large datasets with links in various formats, such as hierarchical, spherical, and so on. Some, such as Maltego, include a bubble chart view, as shown in the next slide. It is important to understand that the majority of link analysis tools are only truly effective with normalized, structured data. They make great options for asking questions of this type of large structured data such as netflow, telephone records, or other transactional or relationship types of data. They don't handle unstructured data well, although some tools do provide proprietary backends and tools that can be used to parse unstructured data, such as IP addresses out of a narrative-style report. An important takeaway is that analysts must be confident in their data sources and the accuracy/consistency of that data to successfully use these types of tools.

References:

- Paterva: https://www.maltego.com/?utm_source=paterva.com&utm_medium=referral&utm_campaign=301IBM:
- https://www.ibm.com/products/i2-analysts-notebook
- Palantir: https://www.palantir.com/products/
- Centrifuge: https://www.centrifuge.tech/
- Gephi/Graphviz: https://gephi.org/
- Neo4J: https://neo4j.com/
- Titan: http://titan.thinkaurelius.com/

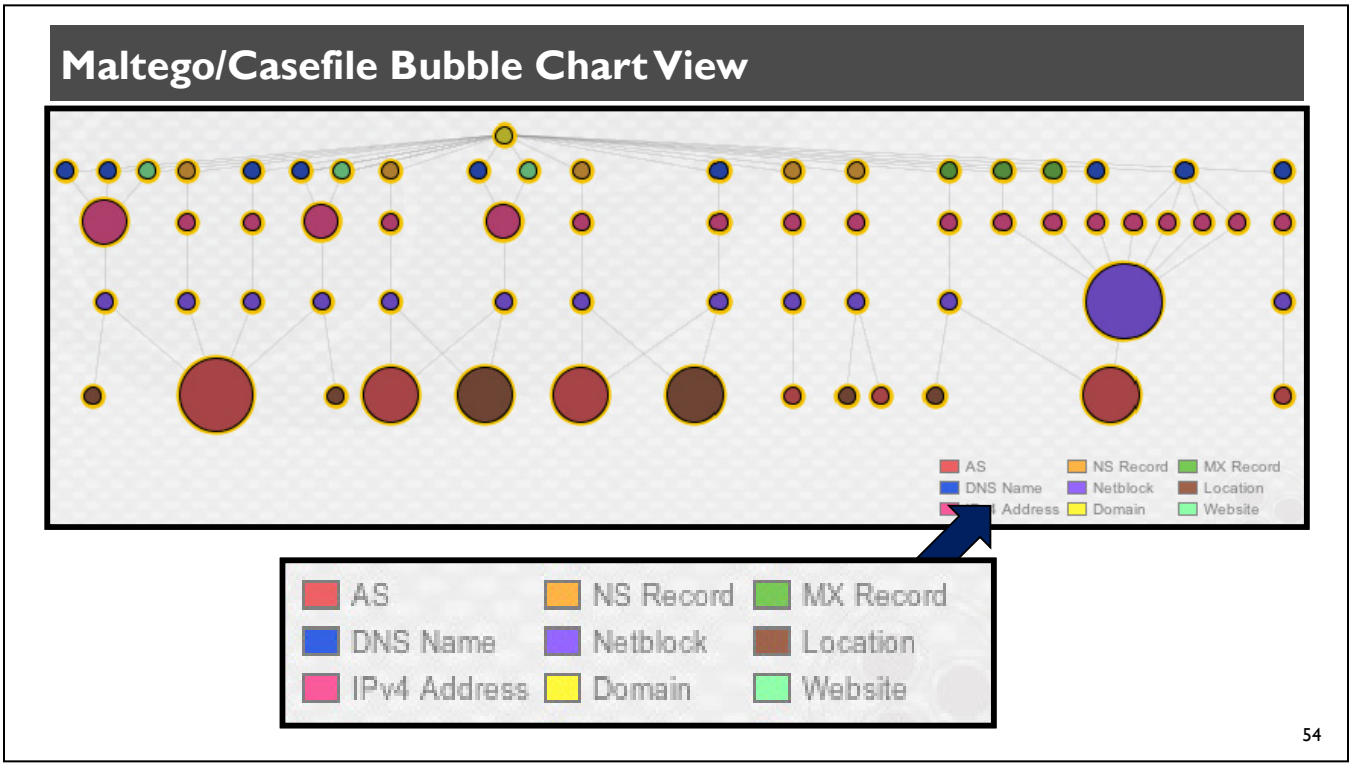

Maltego/Casefile Bubble Chart View

Some tools provide different visualization schemes and layouts. Maltego and Casefile, for instance, provide a bubble chart view that creates different size bubbles, depending on the weight or prevalence of a particular entity in the dataset. As with all visualization tools, the point is not to simply "make a picture," but also to use the visual representation to learn about a large set of data that you can't process by viewing it in a written format or to convey to your audience in a more concise manner.

So, what can we determine from the chart on the slide without even fully understanding the specific entity values within the chart? First, we know the domain has multiple subdomains with corresponding IP addresses that are in different locations. One block of the resolution IP addresses falls within the same netblock as depicted by the large purple circle on the right side of the chart.

Let's shift away from link analysis and take a look at how analysts can use temporal analysis to gain insight into datasets when time becomes an independent variable.

Data Analysis

- The cleaning, transforming, and modeling of data
- Insights revealed through new techniques, models, and correlations between datasets
- Numerous ways to do data analysis, many of which tend to be heavily complemented by structured and unstructured models and machine learning
- Data science is a growing field and often complements threat intelligence very well

Data Analysis

Data analysis includes the cleaning, transforming, and modeling of data, especially for the purposes of revealing patterns and new insights into the data itself. The field of data science largely utilizes data science and modeling through various methods, including machine learning to drive new value out of sometimes disparate datasets.

Data analysis, especially on intrusion trends and data, can be incredibly powerful. Often, data scientists complement threat intelligence teams very well.

Temporal Data Analysis (1)

- Analysis of data over time
- Reveals patterns of activity that reoccur
- Useful for trending adversarial activity
- Instrumental in proactive CTI analysis
- Requires data elements to include a date/time

Temporal Data Analysis (1)

Temporal simply refers to time, and in this case, it is an independent variable in our dataset. Viewing datasets along a timeline can reveal patterns within the data that might not be readily apparent through visual link analysis. Some tools contain built-in functionality to reorganize entity-link depicted data into a timeline, as long as either the entities or links contain a date/time element.

Temporal Data Analysis (2)

Figure 2: APT10 domain registration times in UTC+8

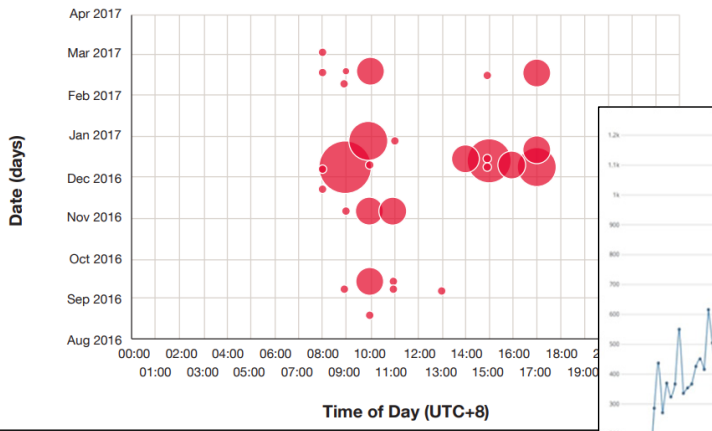

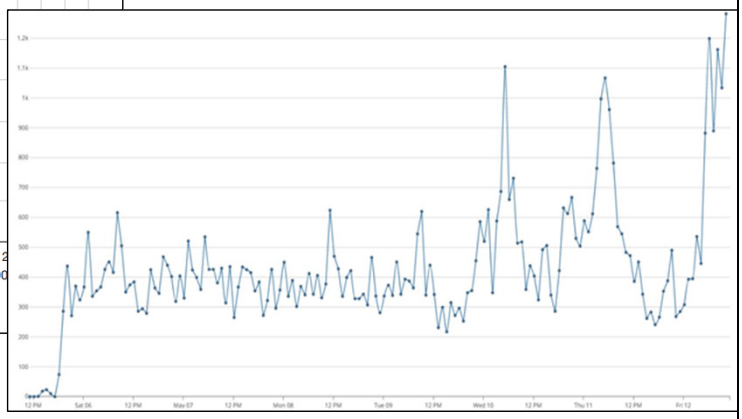

Trend Analysis

- Kill Chain or Diamond Model completion yields intelligence
- Intelligence over time reveals patterns between intrusions:
  - Range from general to specific, ephemeral to immutable
- Assimilation of external intelligence clarifies patterns
- Leveraging different types of analysis (Visual, Link, Temporal) can assist in identifying patterns

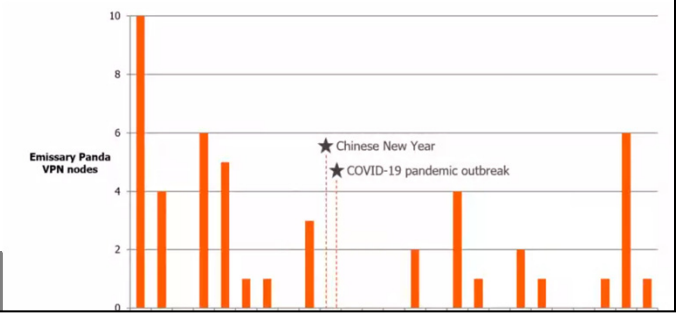

SANS DFIR

Trend Analysis

As you detect and respond to intrusions over time, complete Kill Chains or Diamond Models, execute the indicator life cycle, digest external intelligence, and build your intelligence corpus, you'll begin to notice trends in intrusions. Some of these trends are general, some are quite specific. Some are ephemeral and passing, but some will be lasting. The inclusion of external intelligence makes these patterns more robust and clearer. It is these trends that you observe that are the basis for the definition of clusters of intrusions.

Reference:

- https://www.sans.org/webcasts/star-webcast-constant-change-tracking-adversary-trends-115525

Case Study: Panama Papers

This page intentionally left blank.

John Doe

On May 6, "John Doe" leaked the internal files from over 214,488 offshore entities to German journalist Bastian Obermayer to fight against "economic injustice"

The files represented the largest information leak of history of organizations avoiding tax and sanctions, laundering money, and making deals for international favors

The International Consortium of Investigative Journalists (ICIJ) validated the leaks and set up a team for a year-long effort to analyze the files

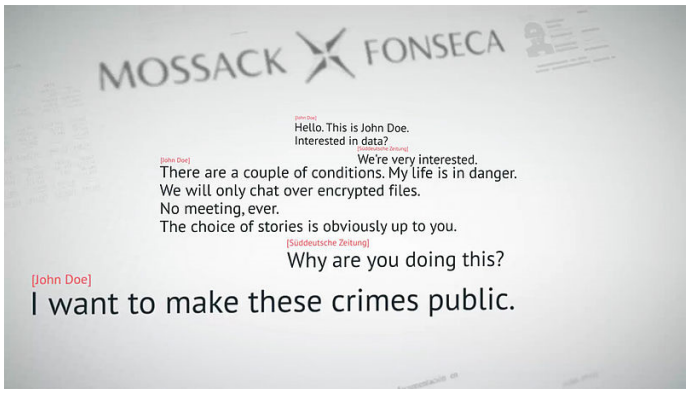

107 media organizations from 80 countries across 25 languages organized to analyze the leaks under the ICIJ, and it was released April 3, 2016

John Doe

A whistleblower that identified himself as "John Doe" came forward on May 6, 2015, to leak files from the corporate service provider and law firm, Mossack Fonseca, a Panama-based company used by wealthy individuals around the world. John Doe noted that he was coming forward because of economic injustice; this was during the height of the 1% movement that was taking place around the world as a protest against the "upper 1%." German journalist Bastian Obermayer of *Süddeutsche Zeitung* received the documents and turned them over to the International Consortium of Investigative Journalists (ICIJ). After confirming their legitimacy, a special team of over 100 media organizations across 80 countries and 25 languages was formed to analyze the data. However, that would, of course, be a complex challenge that required data analysis skills identical to those we are talking about today.

References for Case Study:

- https://www.bbc.com/news/world-35954224
- https://www.icij.org/investigations/panama-papers/data-tech-team-icij/
- https://linkurio.us/blog/panama-papers-how-linkurious-enables-icij-to-investigate-the-massive-mossack-fonseca-leaks/
- https://www.firstpost.com/world/panama-papers-verdict-pakistani-anti-corruption-body-seeks-bank-details-of-nawaz-sharif-and-his-family-members-3946803.html
- https://www.icij.org/investigations/panama-papers/

The Challenge of Data

- 2.6TB of leaked data, reporters from 100+ media outlets, 25 languages
- 11.5M files total including 4.8M emails, 2.2M PDFs, and 1.2M images
  - Journalists indexed the documents in Apache Solr and Apache Tika
  - Nuix for optical character recognition to make them machine-readable
  - Coupled with Neo4J and Linkurious to do link analysis

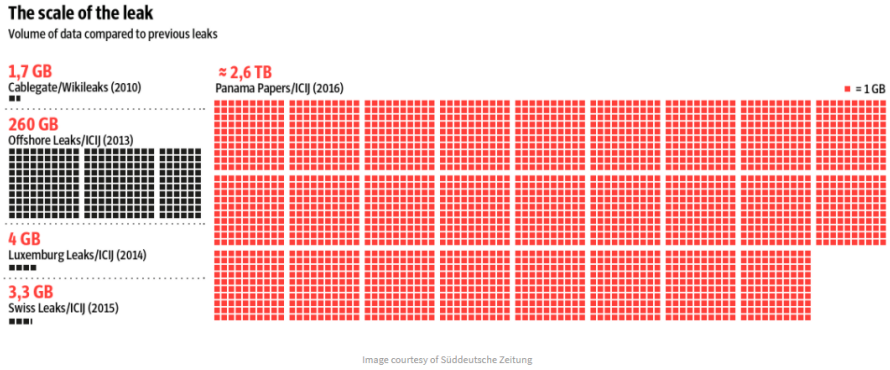

The Challenge of Data

The journalists were faced with an initial challenge that all of us who are threat analysts know too well: Too much data is overwhelming. It's not hard to find things of importance in a few intrusions. To look across a lake of data is an issue, though, especially if it wasn't put into a structured schema to begin with. The journalists had 2.6TB of leaked data with over 11.5 million files to go through. Before even beginning analysis, they had to choose a technology stack. They decided to use open-source tools like Apache Solr and Apache Tika to index all the data. Nuix donated its technology for optical character recognition (taking a scanned document and making it machine-readable/searchable/indexable), and coupled all that with Neo4J and Linkurious to do the bulk of the analysis, which for them would be link analysis.

Before they started with the technical stack, they would have had to make requirements: How will we use this data? Do not let the technology stack get ahead of the requirements.

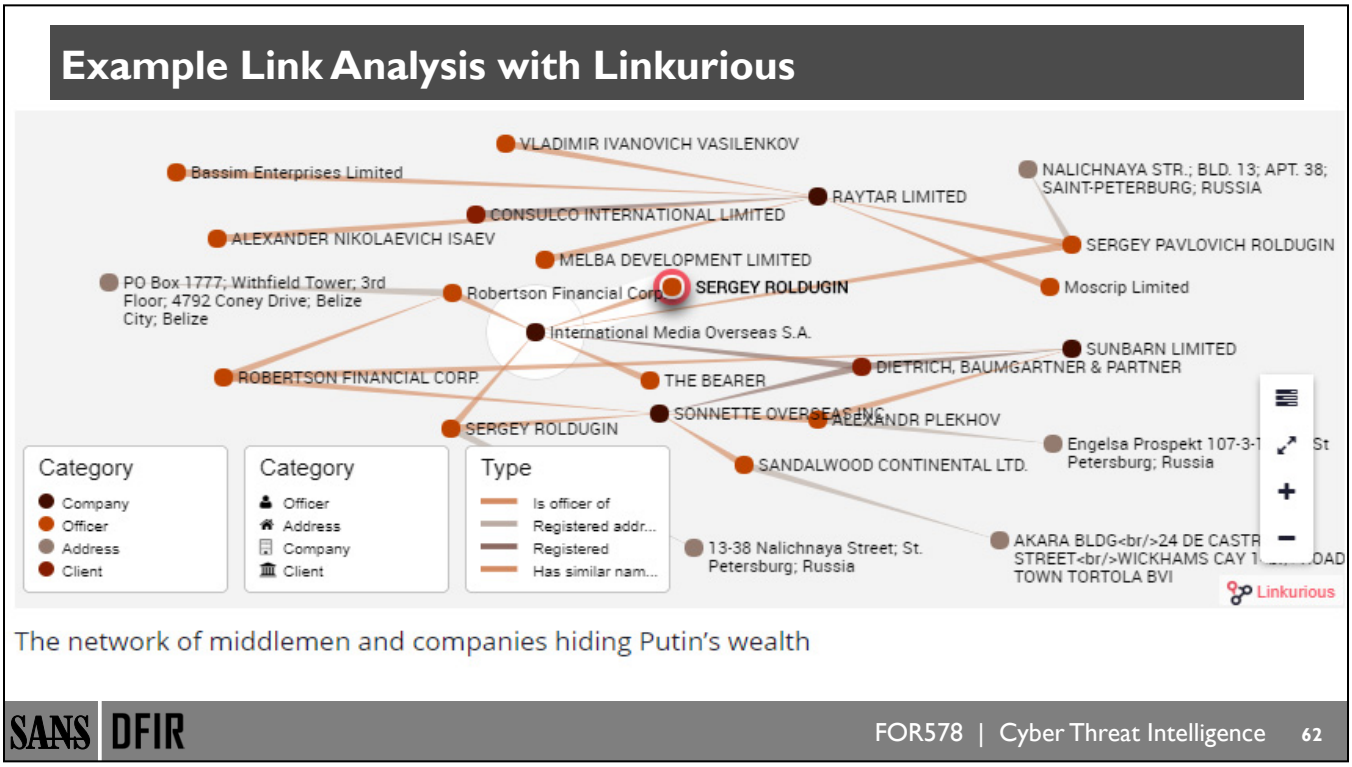

Example Link Analysis with Linkurious

Here we can see one of the sample files and graphs from the data as compiled by the ICIJ. The focus of this graph was Sergey Roldugin, one of Putin's closest friends. The journalists used the documents to make a link of all the data associated with him, which helped reveal shell companies, where money has flowed (to and from), and what countries, companies, and people benefited from the laundering of the money. This is a good example of how link analysis can work in general, even though there's no immediate obvious value for CTI in this dataset. However, as we will see toward the end of this case study, there are actually some good crossovers with CTI other than the core analysis skills, which are themselves identical.

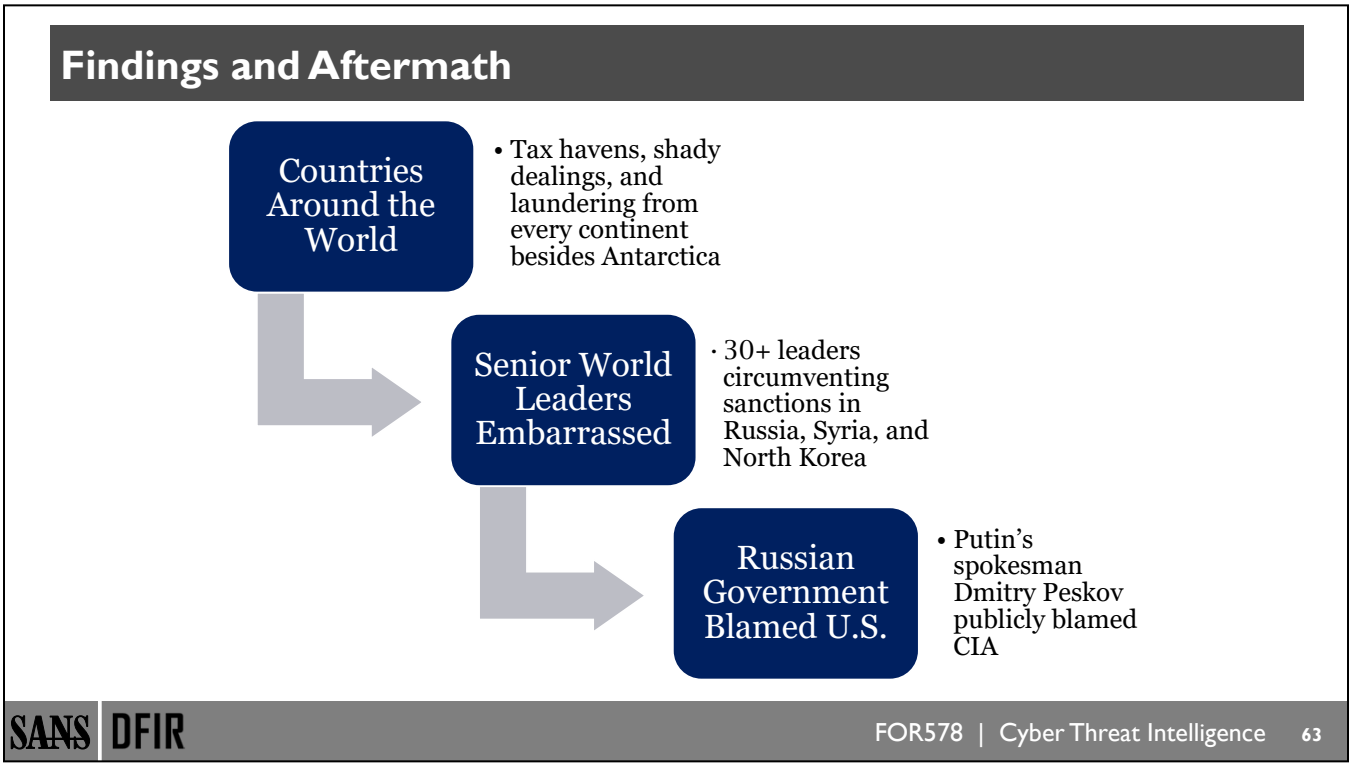

Findings and Aftermath

While some individuals were tried and convicted for illegal dealings, many of the "offenses" were technically legal. The leaks showed nepotism and shady deals between countries and individuals of power in those countries, tax avoidance, and profiteering, but most was not technically illegal because of the way international laws work with respect to host nations' laws, such as those of Panama. The International Monetary Fund (IMF) estimated the dealings cost developing countries $213 billion a year. The shell companies being stood up largely impacted the poorest countries in the world and helped dictators and their close allies avoid sanctions.

The U.S. Treasury Department used the data to identify more than 30 leaders in Russia, Syria, and North Korea circumventing sanctions and highlighted previously undisclosed dealings in Russia, including hundreds of millions being given to Putin's allies for favors and positions; an example was a revelation that $230 million in tax funds by Moscow tax inspectors was stolen by individuals with ties to the Russian government. The official spokesman for Putin came out (after at first denying and avoiding the leaks were accurate) and claimed the leaks were an operation by the State Department and the Central Intelligence Agency (CIA). Scholars have noted this would become the prime motivation behind another case study we will explore later: The Shadow Brokers.

Reference:

- https://en.wikipedia.org/wiki/Panama_Papers#cite_note-ABC_explainer-20

CTI Angle: Intelligence-Driven Hypothesis Generation

- Case study is an interesting one of data analysis and link analysis
  - However, there are also tangible tie-ins to CTI
- By using requirements to drive their analysis, journalists around the world were able to analyze a large dataset sigh major findings
- Leverage major events that might change the targeting patterns of campaigns and threat groups of interest to you
  - Nation-state-backed teams from around the world had the motivation to target ICIJ
- Sample Intelligence-Driven Hypotheses:
  - Journalists covering the Panama Papers will be targeted by threat groups
  - Panama Paper-themed phishing emails will be used by opportunistic threats
  - Financial support and sanction evading will reveal trust relationships
- Leverage findings to help satisfy intelligence requirements for your own intrusion and campaign analysis

Intelligence-Driven Hypothesis Generation

Intelligence analysts need to leverage major events and items of interest to nation-states to identify threat groups that have operated with the motivation of nation-states before. As an example, if you were tracking specific campaigns that have previously operated in the perceived interests of the Russian state, looking for their TTPs and key indicators in datasets related to other areas of interest to them could reveal new patterns. In major crisis-like events, adversaries do not possibly have the time to fully prepare and plan operations, leading to plenty of potential for OpSec-like issues and reuse of tradecraft and capabilities. Identifying this can help satisfy knowledge gaps in your own intrusion and campaign analysis.

Exercise 4.4: Visualizing Large Datasets

- In this lab, you will use Maltego for visualization
  - Visual representation expedites pattern identification
  - Different visualizations reveal more subtle patterns
  - Pivoting more intuitive through UI
  - Visual representation of data to management, other analysts
- Analytical setup more involved
  - Excel/bivariate analysis as "first cut" on data
  - *Maltego*/graph-based tool to round out analysis, documentation

Exercise 4.4: Visualizing Large Datasets

After the initial analysis performed by your team in Excel, you are able to secure funds to purchase licenses for the data graph creation tool *Maltego.* This tool will allow you to use more sophisticated analytics to identify patterns that remain somewhat opaque to simple bivariate analysis (such as what you did with a Pivot Table in Excel). The visual representation of the data will allow you to pivot and dig into details in a different and more intuitive manner and represent the findings visually to a broader audience (such as leadership or other analysts unfamiliar with the dataset).

The complexity of the tool, however, means the analyst incurs a higher up-front cost in terms of effort and time to properly manicure the data for consumption by the tool. For this reason, sometimes a simple bivariate analysis will occur as an initial view on the data to expedite courses of action selection for the "easy wins," followed later by the use of a graph-based tool like *Maltego*.

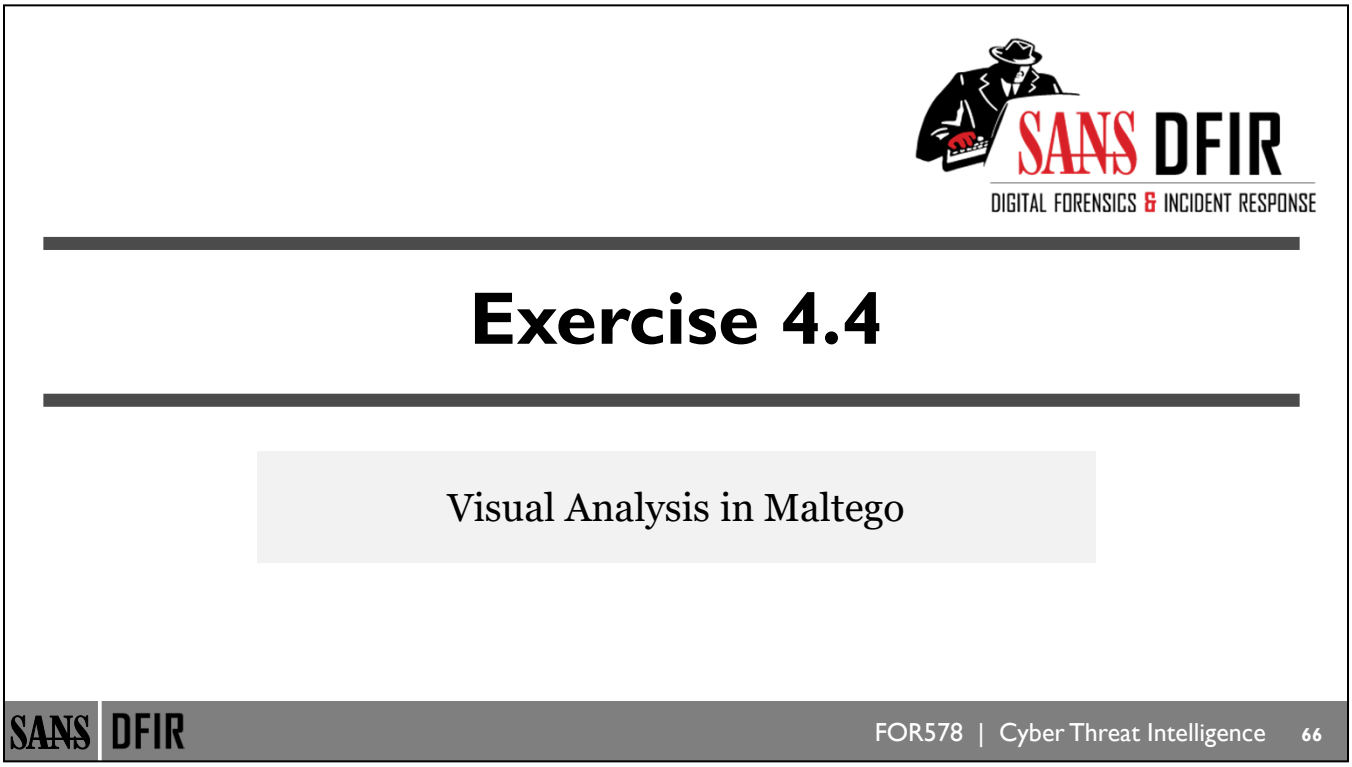

This page intentionally left blank.

Analysis: Clustering Intrusions

This page intentionally left blank.

Style Guide

Throughout the class, we've been progressing from a single intrusion without a style guide, without intelligence requirements past Day 2, and yet have been linking intrusions into sets. This was the natural progression for many CTI teams. If you have felt frustrated and felt like you didn't know exactly what should or shouldn't go into a set of intrusions, you are not alone and that is absolutely part of the point. This was how many analysts felt before the field became more of a profession and started formalizing some methods and structured analytic techniques like the Diamond Model. Schools of thought formed in different companies doing this type of work; many fantastic firms like Mandiant, Palo Alto's Unit 42, Symantec, Microsoft, Kaspersky, and more have formed their own unique schools of thought with community sharing taking place to trade ideas and examples. That's one of the key roles of FOR578—a school of thought that's transparent in its objectives and open to the community to give one approach to CTI.

The most important thing for you to do, though, in starting your CTI team or formalizing an existing one is to establish a Style Guide. A Style Guide can contain anything that's useful to you that's going to be mid-to-long-term value to your team but, at a minimum, it must include:

- Team structure
- Accepted lexicon
- Words, phrases, and actions to not do
- Sample structure analytical techniques
- Sample intelligence requirements with example outputs
- Guidance to analysts on how to create clusters and finalize intelligence products
- Key processes to follow

It's incredibly important to ever being able to cluster and track intrusions over the long term to have a style guide that helps guide your analysts. Something as simple as setting formal definitions for groups will keep folks focused and able to communicate among themselves.

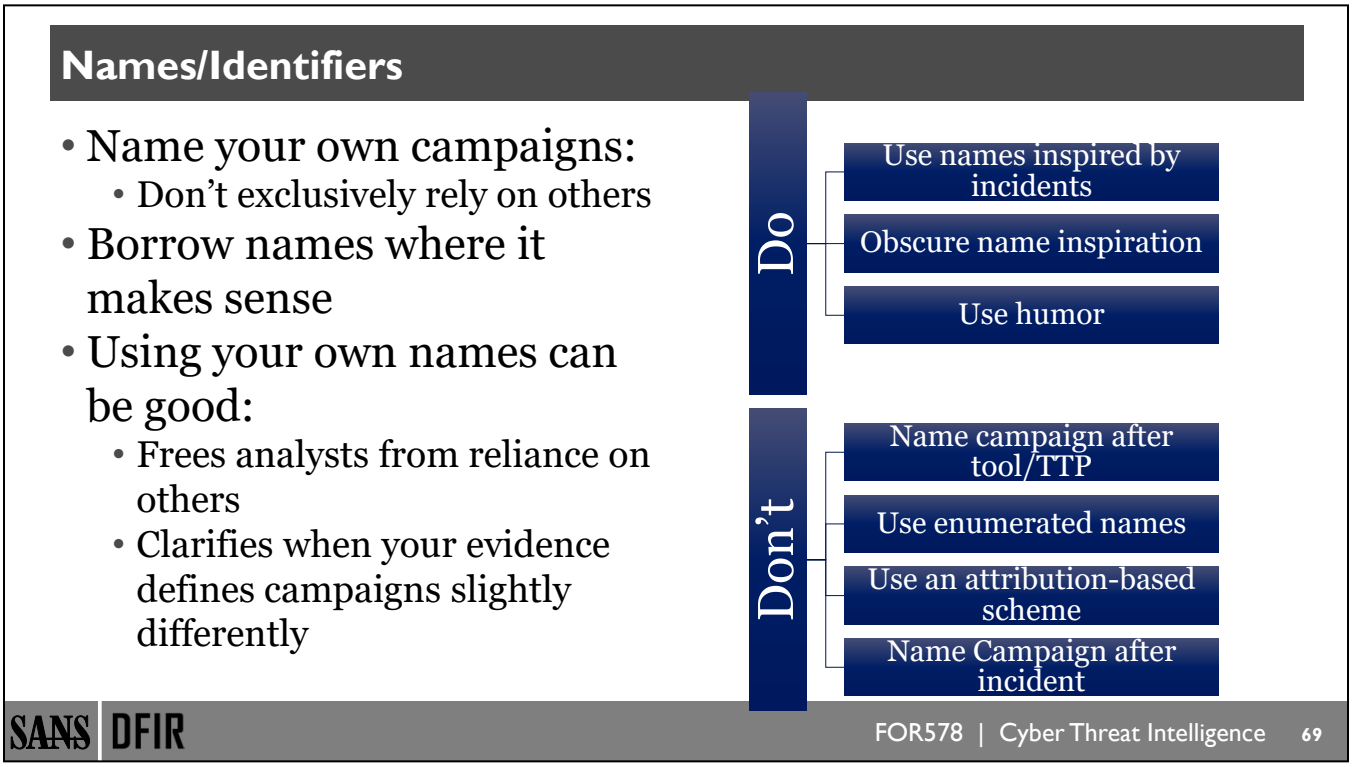

Names/Identifiers

At first, it might seem silly to have a slide dedicated to naming campaigns, but the experience of many has been that this can actually be problematic. This is not so different from the nomenclature problem that has plagued the antivirus industry for decades: What do you call it, what do other people call it, and how do you translate between these names?

Years of experience and many failures later, possibly the best advice is to accept that these translation issues will occur. It's advisable that you use names you come up with—or that trusted peers have come up with—because your campaign definitions (in terms of key indicators and behavioral TTPs) will inevitably differ from those reported by vendors. This is simply because you will have more direct access to the raw data and intelligence to formulate them, as well as TTPs for a single actor, which can differ between targeted industries.

So, in formulating a campaign name, there are some "DOs" and "DON'Ts" that can help avoid problems in the future. (Helpful tip: Renaming a campaign after you realize it's been poorly named is much more difficult to do than you'd think.)

DON'Ts:

- **Name a campaign after a tool or technique**: Tools are shared, and naming a campaign after a tool that becomes shared creates confusion among analysts. There is a story about a team of CTI analysts who, in the early days of the Poison Ivy backdoor, named a campaign Poison Ivy. A year later, there was the Poison Ivy campaign and Poison Ivy 2. Then came Poison Ivy 3. At that point, the team learned a lesson and renamed all three, with many headaches resulting.
- **Use enumerated names**: Be creative. Generic names such as "APT-1" (no offense to Mandiant) make it difficult for analysts to keep them all distinct when you get up to 20 or 30 campaigns.

Risks of Clever Naming Conventions

- CrowdStrike employs a clever naming convention for campaigns
  - Countries receive animals that are easily remembered
  - China has Pandas, Russia has Bears, etc.
  - Allows customers to quickly know "Sparkling Bear" is a Russia-based group, whereas "Feisty Panda" is Chinese
- This is smart business, but what's the CTI issue?
  - If you are ever wrong about attribution, your campaign is now stuck to "bear" or "panda" or "kitten," and changing it can be difficult as well as embarrassing
- Takeaway: Allow flexibility with your naming convention

Risks of Clever Naming Conventions

The security company CrowdStrike has a lot of talented analysts and a great intelligence team—however, their naming convention has been controversial in the CTI community. They employ animal names for countries. So, when a customer hears anything "panda," they associate it with a Chinese-based group. Each country has different animals. The problem, though, is that if you are wrong on your attribution, you will find yourself in a difficult spot trying to change the naming convention, changing the group name, or just explaining the outliers. Each become confusing quickly. Additionally, what happens if your Jedi Panda group turns out to be the same as the Fluffy Kitten group you attributed to Iran? Do you just try to have the Panda and Kitten have a baby? That's not going to work out very well. Fluffy Jedi are no galactic heroes.

MITRE ATT&CK Groups Page

- Previous page version shown below used "Aliases"
  - Current page uses "Associated Groups"
  - MITRE tracks what names teams create but does not pass judgment

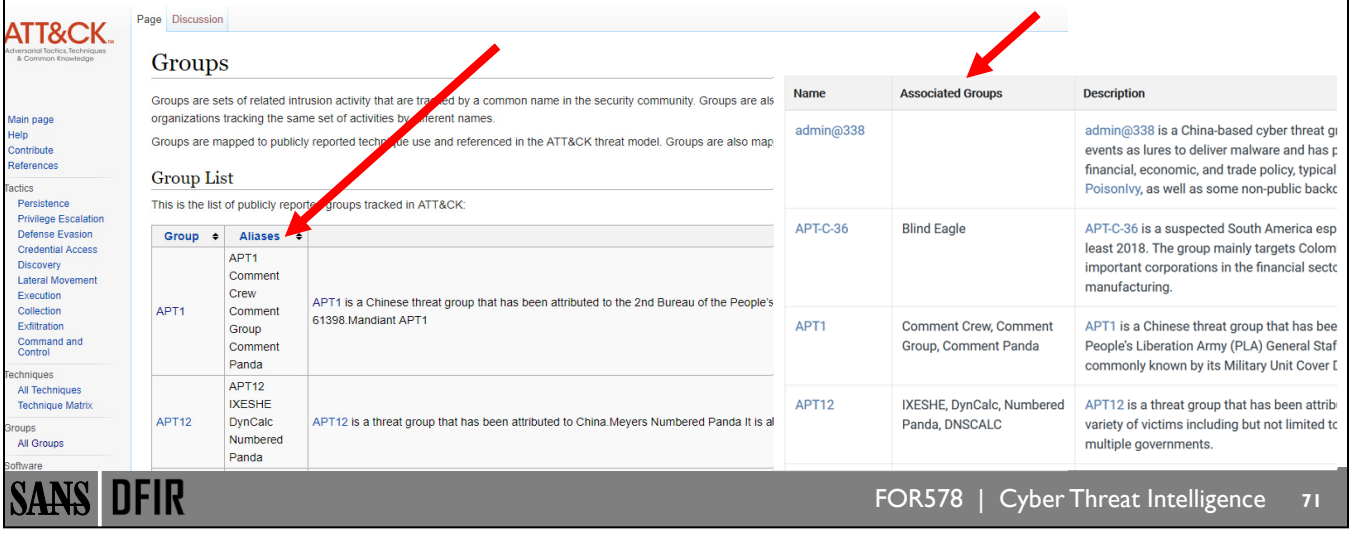

MITRE ATT&CK Groups Page

The MITRE ATT&CK Groups page is a pretty useful resource covering some of the most well-known groups and their associated names. On the slide, you can observe the "old" ATT&CK Groups page, which used the term "Aliases." Next to it, you can observe that the "new" ATT&CK Groups page uses the term "Associated Groups". If you read the text at the top of the page, it might sound familiar—this is because the thinking behind it was inspired in part by this class! The current ATT&CK methodology of tracking the references for each name may provide a useful example of how you could approach tracking group names.

Reference:

- https://attack.mitre.org/groups/

Rosetta Stone: APT Groups and Operations Matrix

- Maps known attribution, campaign names, and malware used
- Loose correlations, but useful in communicating external to the CTI team, as defenders get hit with various names

China						
Common Name	CrowdStrike	IRL	Kaspersky	Dell Secure Wo	Mandiant	FireEye
Comment Crew	Comment Panda	PLA Unit 61398		TG-8223	APT 1	
	Putter Panda	PLA Unit 61486			APT 2	
UPS	Gothic Panda			TG-0110	APT 3	
IXESHE	Numbered Panda			TG-2754 (tentat	APT 12	BeeBus
					APT 16	
Hidden Lynx	Aurora Panda				APT 17	Deputy Dog
Wekby	Dynamite Panda	PLA Navy		TG-0416	APT 18	
Axiom			Winnti Group			
Shell Crew	Deep Panda		WebMasters		APT 19	KungFu Kittens
Naikon		PLA Unit 78020	Naikon		APT 30	
Lotus Blossom						Spring Dragon
	Hurricane Panda					
	Emissary Panda			TG-3390	APT 27	
	Stone Panda					
	Nightshade Panda				APT 9	
Hellsing	Goblin Panda		Hellsing			
	Night Dragon					
Mirage	Vixen Panda	Ke3Chang		GREF		Playful Dragon

APT Groups and Operations Matrix

A document located at the link below was created by Florian Roth and then contributed to by a number of members in the community (check the contributors tab on the XLS document) for the purpose of mapping together all the different campaign names. As companies call groups different things, it can lead to a confusing scenario where CTI analysts lose track of the different naming conventions.

It is advised that you keep such a document up to date on your team, as well as your internal naming conventions and where they overlap with existing identified campaigns. It is embarrassing to learn that two or three different threat actors you've heard about or are communicating about are actually the same threat but just named differently at different companies. This is an incredibly lazy way to do name tracking because the analytical model that Kaspersky uses to name its group is not exposed and is absolutely not the same model as used by Mandiant. Merging groups together has significant issues in the insights and recommendation that can come off of those groups. This way of tracking is much more about attribution than anything else. However, each of these companies absolutely have better and more focused ways of tracking threats internally. It is not like the public name is the only method and they are all tracking things incorrectly; it's that in how we communicate about them and the public names tends to be very poorly done, and as organizations try to consume data off of the insights from these macro public names, it can quickly misguide their defense efforts.

However, it is a useful tool because others will communicate and use these names interchangeably. Therefore, it is recommended to maintain a Rosetta Stone and have a column on the group you are tracking internally and understand the links they may have to other external groups and names. That way, when people in your organization try to communicate about threats, you can translate it and refocus them on what matters to your organization that you are tracking. From an intel perspective, this approach can misguide you, but as a tool to communicate to others, it can be very effective. Likewise, if your focus is simply on attribution, then this can be an effective tool, as whether the two groups are the same or not may not matter so long as the overarching attribution is.

Reference:

- https://docs.google.com/spreadsheets/d/1H9_xaxQHpWaa4O_Son4Gx0YOIzlcBWMsdvePFX68EKU/pubhtml#

There is No One-to-One Mapping

- Each organization's intelligence team *should* create an analytical model
  - From that model, they analyze intrusions, select key findings, and cluster them
- The output, a cluster/group, is unique to each organization and not reversible
  - The intrusions they see, the weights they apply, the findings, etc., are unique
  - Clusters cannot have a one-to-one relationship
  - Clusters can have links to other clusters
    - Links can be one or more of the Diamond Model vertices

Source	Title	Group
Crowdstrike	Bears in the Midst: Intrusion into the DNC	APT28/29
ESET	En Route with Sednit version 1.0	APT28
ESET	Visiting The Bear Den	APT28
FireEye	APT28: A Window Into Russia's Cyber Espionage Operations?	APT28
FireEye	HAMMERTOSS: Stealthy Tactics Define a Russian Cyber Threat Group	APT29
FireEye	APT28: At the Center of the Storm - Russia strategically evolves its cyber operations	APT28
F-Secure	BlackEnergy & Quedagh the convergence of crimeware and APT attacks, TLP: WHITE	APT28
F-Secure	The Dukes 7 years of Russian cyberespionage	APT29
F-Secure	COSMICDUKE: Cosmu with a twist of MiniDuke	APT29
F-Secure	OnionDuke: APT Attacks Via the Tor Network	APT29
F-Secure	COZYDUKE	APT29
Kaspersky	Sofacy APT hits high profile targets with updated toolset	APT28

Wrong

SANS DFIR FOR578 | Cyber Threat Intelligence 73

There is No One-to-One Mapping

It is very common for analysts to try to merge names. It seems so silly to have to track all these different made-up names. And at times, there are well-meaning analysts that argue, "why can't we all just pick a name and stick to it." That can work with malware most of the time, because the malware is a defined entity. It is a fact: "this is a sample of X malware." But clusters of intrusions under any name (Campaign, Activity Group, Threat Actor, Intrusion Set, etc.) are abstractions of data according to an analytical model, whether it's poorly or well defined by analysts. You are no longer comparing facts but instead are comparing assessments. And assessments on even the same data may reasonably be different, but assessments on different data such as different intrusions are not going to be the same. What an analyst abstracts away is a choice that dictates the final product and how the users can consume that intelligence. It may seem like semantics, but there are real impacts on what you value and what other organizations value based on the intelligence requirements levied. The entire intelligence process ends up being different based on what the requirement is, how it drives collection, and the impact of the analysis and production of that collection.

Image source is from the US Government's Enhanced Grizzly Steppe report on Russian activity. It is wrong because ESET, CrowdStrike, F-Secure, and Kaspersky do not track APT29 or APT28. That is a FireEye clustering. And attempting to note that APT28 and APT29 are one-for-one representations of what those firms track is inaccurate and simplistic. It still may serve the same function, though, of discussing attribution (i.e., it's wrong, but does it matter?).

A longer form of this discussion can be found here as a SANS webcast:

https://www.youtube.com/watch?v=3CUNlgQBwc4

One-to-One Mapping Issues (Example)

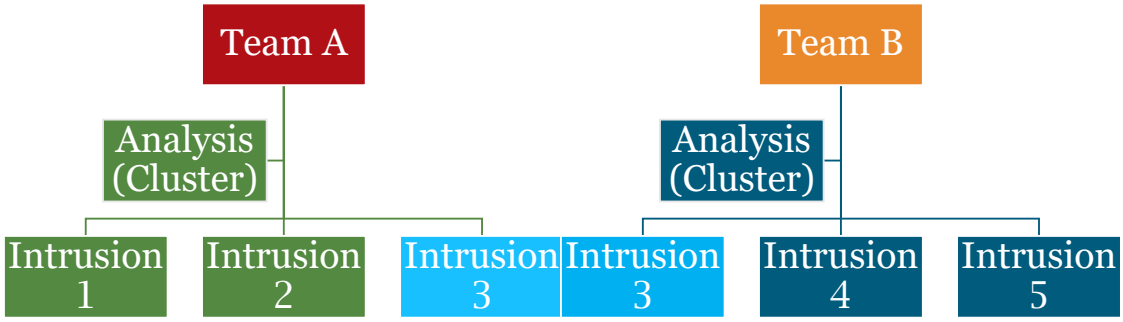

- Everyone has different collection even if sometimes it overlaps
  - Team A and Team B won't work all the same intrusions
  - Different intelligence requirements drive different analysis of that collection
- Team A's clustering of intrusions (Activity Group) cannot be the same as Team B's
- These differences have impact on how you utilize the final intelligence product

One-to-One Mapping Issues (Example)

Here is a more visual example of the previous concept if an organization or team comes up with an Activity Group or Threat Actor that is inherently a cluster of activity. What is captured for that cluster is not all the details of each intrusion but instead what the analyst values for their analytical model. The team or organization cannot possibly work every same intrusion as another team as well because of collection differences. Therefore, the clusters are based off of different data, even if there are links, and what the analysts find value in those datasets is different and possibly for different goals.

At the Strategic level, it may not matter. If both groups represent activity by "Russia" and you're trying to tell a story about Russian activity (i.e., attribution-based analysis), then it may be good enough. Operationally, it can pose an impact; if the tradecraft is similar or the same in the groups, then it may be useful to the defenders in that way, although the victimology may be different, as an example, and have different takeaways for different consumers. At the tactical level, the indicators and expected observables in the environment can be wildly different. It is okay to view two different groups and then abstract what's important to you for defense. It is not okay to merge the groups together to track future intrusions. You are now adding your own analytical model and weighting on top of someone else's, which will drive completely different takeaways along each diamond model vertex over time.

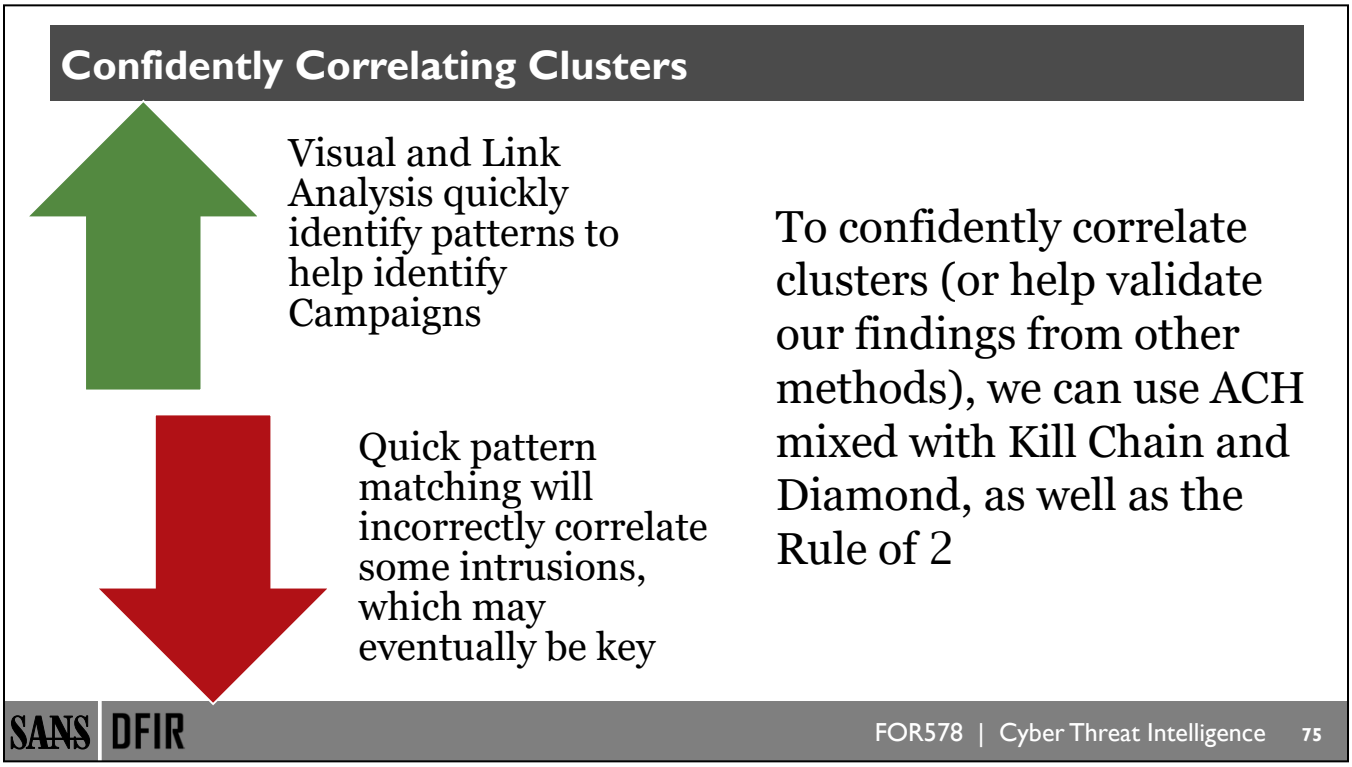

Confidently Correlating Clusters

Different types of analysis on top of intrusion data can absolutely help identify campaigns. Visual analysis, link analysis, clustering, and other forms of analysis on top of structured datasets can help reveal patterns. However, if we want to confidently tie intrusions (to intrusion sets, to campaigns, to groups), we need to use structured analytical techniques. There are many you could use, but we will leverage ACH and the Rule of 2 over this next section. In the lab, you will leverage the Rule of 2. We will do this using the Kill Chain and Diamond Model, but note that you can use the techniques with other structured schemas as well.

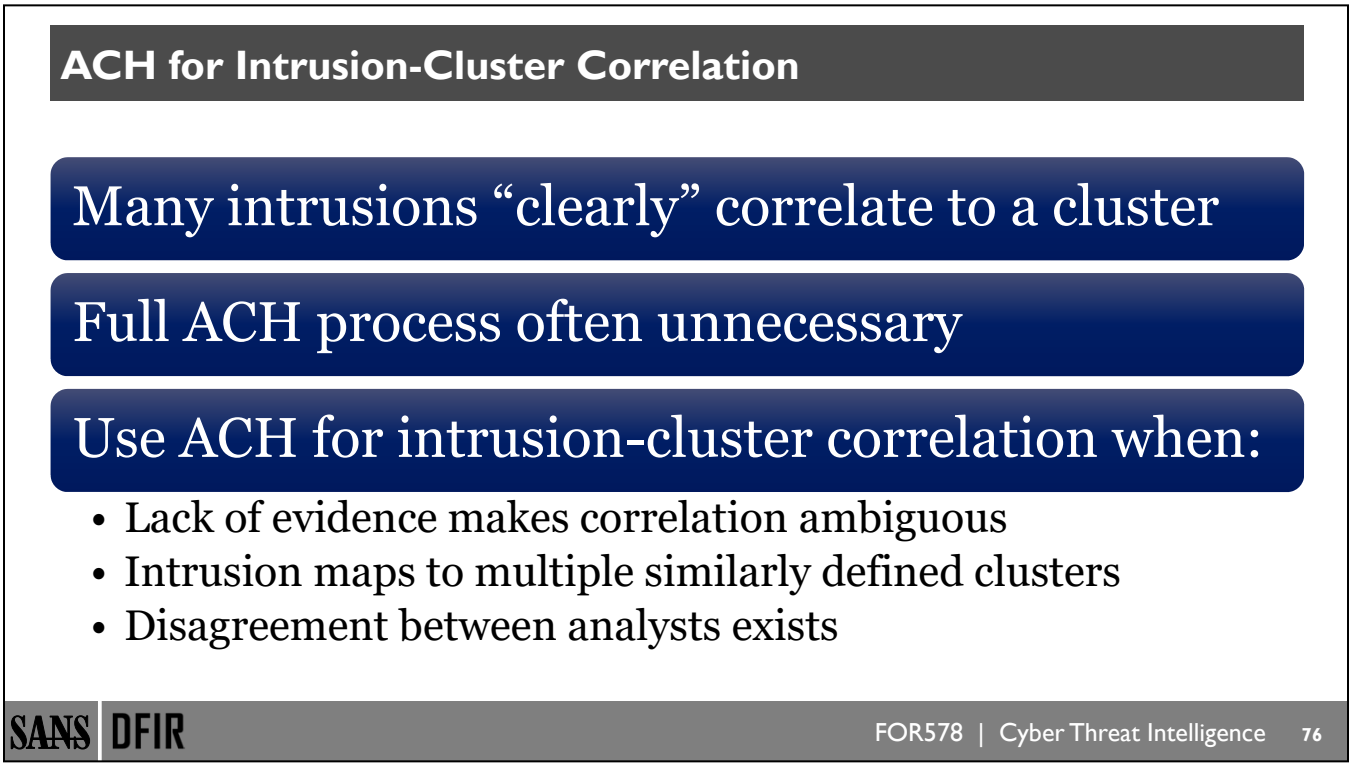

ACH for Intrusion-Cluster Correlation

The process of attributing a single intrusion to a cluster can sometimes be straightforward, but often it is not. This could be due to two clusters that operate similarly or because of a lack of specific indicators for a cluster. This ambiguity drives the need for a more formalized, rigorous exploration of attribution in which ACH makes sense. Keep in mind, ACH for cluster correlation probably doesn’t make sense for every intrusion. Often, the attribution is “obvious” based on the intuition of analysts. Where disagreements or ambiguity occur, this process will help resolve, or at least qualify, that uncertainty.

The Basics

- Follow ACH steps
- Classify evidence based on intrusion definition:
  - Kill Chain
  - Diamond
- Confidence in assessment informed by support in each clustering of evidence

The Basics

When determining what intrusion(s) contribute to a single campaign, it is important to follow the analysis of competing hypotheses process and classify the evidence from the intrusions according to the kill chain and the diamond model. From there, it is possible for you to make a confidence assessment on whether or not the intrusions are linked by a single adversary campaign.

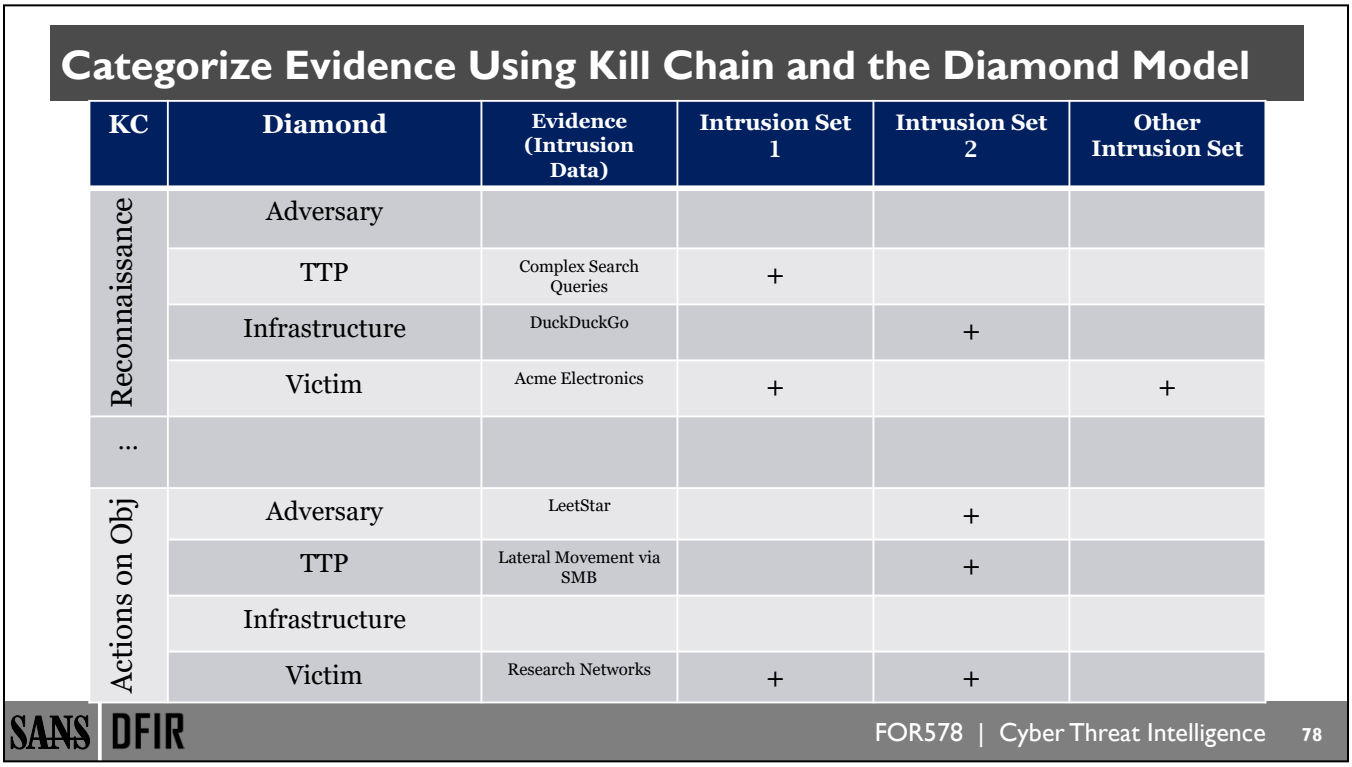

Categorize Evidence Using Kill Chain and the Diamond Model

One method to move from intrusions to intrusion sets, to campaigns, to groups exist in using the kill chain and diamond model phases in conjunction with the ACH process. Here, we structure categories for diamond and kill chain for our structured schema. Then we would include evidence across them for intrusions. We would not take every possible intrusion but the ones where we found some key indicators or behavioral TTPs. We would leverage this process to move intrusions into intrusion sets that we are already tracking. This is a complicated and time-consuming process but a way to be highly accurate. We will discuss shortcuts later on.

Always include "other intrusion set" and make sure that the information we are tracking is actually descriptive of the intrusion set we are tracking. As an example, the fact that someone targeted our organization Acme Electronics is not descriptive of specific intrusion sets, but all the ones we are interested in tracking internally. However, specific types of lateral movement using SMB commands or the specific targeting of Research Networks might be descriptive of one or multiple intrusion sets we are interested in tracking.

When something is not descriptive of anything, we will highlight it to see if it's useful later on but remove it from our process. We will only keep track of data that has a + in a category; not all of them or none of them.

Enumerating Intrusion-Campaign Hypotheses

- Take key indicators/TTPs/findings from intrusion sets as your evidence to position against campaigns as hypotheses
- Evidence may most strongly support correlation to one un-attributed intrusion
- Always include "other campaign"
  - Lots of information there may indicate the need for a new campaign

Evd	Campaign A	Other Campaign
E1		+
E2		+
E3		+

Evd	Campaign A	Campaign B	Other Campaign
E1		+	
E2		+	
E3			+

Enumerating Intrusion-Campaign Hypotheses

Enumerating campaign hypotheses is not as straightforward as it might seem. First, of course, identify candidate campaigns for which you have sufficient evidence to perform correlation: Those which you have key indicators and TTPs of. The more evidence, the more confident your assessment will be!

Often, evidence aligns with intrusions that you have observed, for which you do not have any campaign yet defined. This should be captured in your hypotheses! Initially, it makes sense to include "other unattributed intrusion" as a hypothesis. If you notice that evidence keeps supporting a single other, noncorrelated intrusion, it makes sense to add that intrusion as its own hypothesis, at which point you may have:

- Campaign A
- Campaign B
- Other Campaign

If, in the end, your evidence supports correlation to some other, unattributed intrusion more strongly than one of your hypothesized campaigns, you might have just discovered a wholly new campaign!

External Intrusion Reports

Complement Knowledge Gaps	Address methods and behaviors you did not previously know
	Think operationally (leverage Diamond Model)
	Inspire a threat-based hypothesis on how you hunt in your network
Do Not Merge With Your Data	Not all intelligence is created equal
	Marketing sometimes wins out
	Using the vendor info or name as your campaign name forces you to lose control of the narrative

External Intrusion Reports

External reports from trusted third parties, industry collaboration, or vendor reporting can amplify your understanding of a campaign in important ways. It carries the risk of misleading your analysis should the findings in those reports be problematic themselves. Value the observable data reported over analytical conclusions, and use this to amplify your own knowledge of a campaign.

This data should be documented along with your campaign so it's clear what components of the campaign are the result of direct versus indirect observation. Be mindful that others may attribute an intrusion to a campaign using different criteria than you, and what is reported as being related may not, in your estimation, be so.

Do not use the vendor's name as your name; it will cause you to lose control of the narrative as your security people or executives hear of high-profile cases that may not be important or relevant to the campaign you're tracking. That is, if they call it "APT28" publicly, come up with some name internally for it if you see a similar campaign you're tracking. If you aren't tracking it (it really is just the vendor's data), keep it the same name then. The way to do this effectively is having an Intel Rosetta Stone.

Diamond Model Deeper Dive: Meta-Features

- The Diamond Model's core features are:
  - Adversary, Capability (TTPs), Infrastructure, and Victim
- The Diamond Model's meta-features are:
  - Timestamp (start and end), phase, result, direction, methodology, and resources
- Each core feature and its meta-features should have a confidence value
  - Definition up to each user; utilize at minimum High, Moderate, and Low weightings

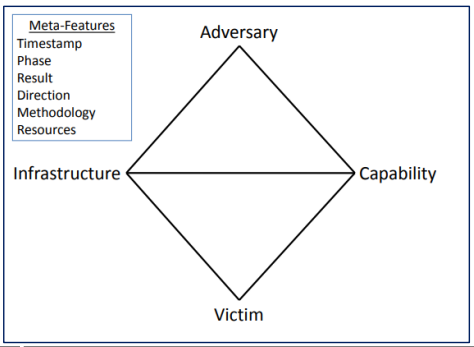

Adversary Event is formally defined as "E"

$$E = \langle \langle Adversary, Confidence_{adversary} \rangle, \langle Capability, Confidence_{capability} \rangle, \langle Infrastructure, Confidence_{infrastructure} \rangle, \langle Victim, Confidence_{victim} \rangle, \langle Timestamp_{start}, Confidence_{timestamp_{start}} \rangle, \langle Timestamp_{end}, Confidence_{timestamp_{end}} \rangle, \langle Phase, Confidence_{phase} \rangle, \langle Result, Confidence_{result} \rangle, \langle Direction, Confidence_{direction} \rangle, \langle Methodology, Confidence_{methodology} \rangle, \langle Resources, Confidence_{resources} \rangle \rangle$$

Diamond Model Deeper Dive: Meta-Features

The diamond model has four core features, which at various times in the class have been referred to as the vertices. These core features are Adversary, Infrastructure, Capability (broadly defined to include TTPs), and Victim. The meta-features are optional additions to the model that can help guide analysts in their use of it. As an example, recording the start and end of the observed activity can be done with the timestamp feature to help keep a temporal aspect to the various intrusions observed and order the activity (especially for TTPs that have a time component like one step being completed before the next). The phase can map to the kill chain but does not have to; it simply means defining a phase of activity, but the kill chain construct is a useful tie-in. The result is the impact or action that the adversary was able to achieve or what occurred because of it. The direction dictates the flow, such as a victim calling out to infrastructure or infrastructure being leveraged to target the victim (infrastructure -> victim; victim -> infrastructure). The methodology can be thought of as the tactic; it is the general class of activity, such as phishing. The resources help identify what were needed to be successful, such as funding, knowledge, software, access, etc.

Each of these features and meta-features can include a confidence score. The score and weighting of it, as well as its use on just features or features and meta-features, is up to the analyst to help them achieve the goal they set out. The only requirement is to be consistent in its use for each problem you are trying to solve.

Reference:

- http://www.activeresponse.org/wp-content/uploads/2013/07/diamond.pdf

Creating an Activity Group

- Six distinct steps to creating an Activity Group
  - Step 1: Analytical Problem (define what you want to solve, your intelligence requirement)
  - Step 2: Feature Selection (event features and weighting of what's important to you)
  - Step 3: Creation (analyze events/intrusions and compare against the model to cluster)
  - Step 4: Growth (compare new events and classify them into the Activity Group)
  - Step 5: Analysis (analyze the Activity Group itself to address the Analytical Problem)
  - Step 6: Redefinition (redefine the model as your needs change and more to new cluster)

Problem: Identify unique adversary tradecraft to create scalable detections

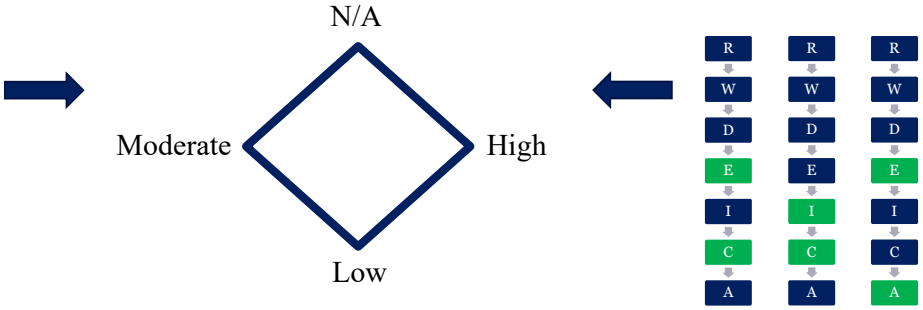

Creating an Activity Group

There are six distinct steps in creating an activity group. Essentially, you are defining an intelligence requirement, determining your weighting of the features you care about to solve that problem, and then analyzing intrusions against that model to abstract out the adversary events that matter to you. In the image, the adversary feature is weighted as not applicable to the requirement, infrastructure is weighted moderately, victims are low, and Capability/TTPs are high. Next to it we see a mock high-level abstraction of intrusions analyzed along the kill chain construct. There may be elements we highlight in the kill chain that are common across the intrusions (i.e., clustering the intrusions based on diamond model features that we weighted). This would help us abstract the intrusions away to a few components each that matter to us and their meta-features to define an Activity Group against the analytical model we've defined.

As intrusions are observed, so long as they match the model, we can add them to that Activity Group. If the victimology among the intrusions is different, we may not change to another Activity Group. If the infrastructure is similar but not identical or has similar aspects, we may keep it in the same Activity Group. But if the Capability/TTP changes significantly, we'd likely note that there is a second Activity Group because our goal here is to break them out into answering the problem of having scalable detections. So we'd want the clustering of activity to each drive different types of threat behavior detections. However, if the victimology and infrastructure change, we might define that the low and moderate weighting together justifies a new Activity Group just like a significant change in Capability would.

One of the unintended but useful benefits of this approach is that false flag attacks (one group attempting to convince the victim that another group attacked them) do not matter unless we weight and value the adversary feature. If the problem being solved isn't related to attribution, then the "who" of the attack doesn't matter.

Different Examples of Diamond Models for Different Req's

- The Diamond Model is a structured analytical technique, there is no 1 way to use it
- The Analytical Requirement may be an intelligence requirement, it may be your own requirement to help further your ability to satisfy other intel requirements

Consider the "K. Lazutin" set of intrusions

- Clustering as a set on overlaps is not really satisfying any specific requirement
- Edison International, MoneyBin, and SpaderTech have all been targeted
- Our Analytical Requirement may simply be developing a defensive strategy against all human operated ransomware groups focused on Capability and Victim
- Requirement: Identify groups that leverage human operated ransomware against Edison International companies and peer organizations

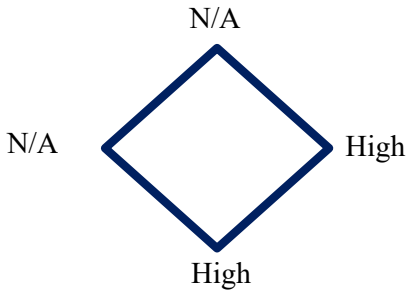

Different Examples of Diamond Models for Different Req's

The Diamond Model can be leveraged to satisfy any range of analytical requirements that you have. It is incredibly useful, though, but not always required, to formally define your activity groups either to a precise level (using all the features and meta-features in the Diamond Model paper) or at a high level (such as the high-level feature selection shown in the slide).

If you decide to formalize your Activity Group by choosing an analytical requirement and selecting the features you care about, realize that you can still record observables in the features you don't select (in the slide, it shows Adversary and Infrastructure are N/A, but we can still record the observables there that we can later pivot on of the intrusions that match to a high degree the Capability and Victim features). Or said simply, the Capability and the Victim in the example shown are the filter which you compare intrusions. If you have matches on both of those features (in this example common Capability and common Victims such as Edison International, Spader Tech, and Money Bin or peer companies) then you would add the intrusion's details to the Diamond Model cluster.

We should be looking for intrusions that meet both of the following requirements. For this feature selection of the group, just meeting one is not enough:

- Consistent TTPs/Capabilities
- Consistent Victims (we could make it specific victims here as we're looking for specific victims and their peer organizations)

So as an example, if we saw someone using WannaCry ransomware against a bank, and another intrusion with Ryuk against MoneyBin, it'd form a group.

If we saw someone using Wannacry ransomware against MoneyBin, and another intrusion using Ryuk against an oil and gas company, it would not meet the requirements for a group.

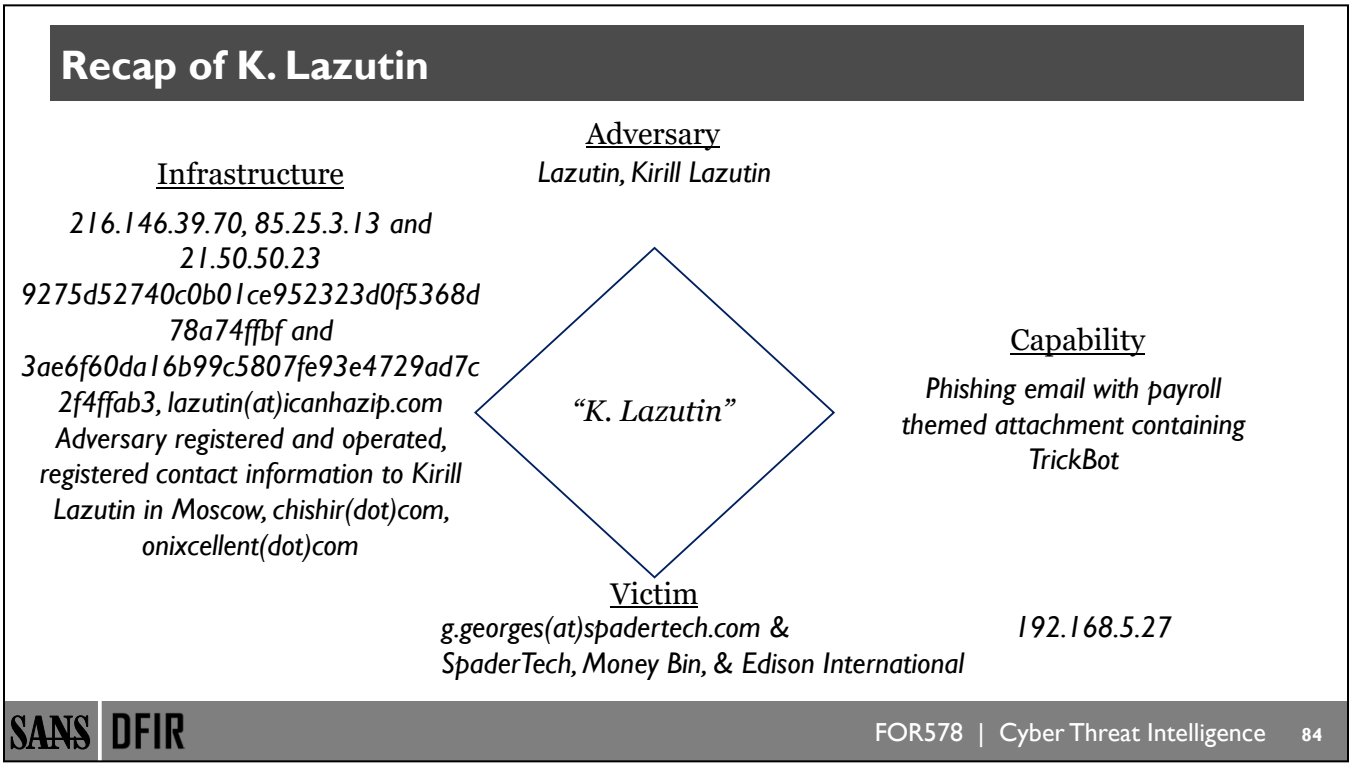

Recap of K. Lazutin

Recap of what we know about K. Lazutin cluster of intrusions. To analyze this against our analytical requirement for the activity group, we really only care to see commonality among the Victim (to make sure it's relevant to our organizations) and Capability. All aspects of "Lazutin" really don't matter anymore. It may be an interesting observable to pivot off of, but we wouldn't add intrusions that have Lazutin-themed emails/domains and personas to the cluster, as it doesn't meet our actual objective. Or said differently, document the components of the Diamond Model that you can see, but only of the intrusions that get added to the group because they meet the filter set by the feature selection.

New Intrusion – Does it Fit?

- Notice the following intrusion: Does it meet our Feature Selection of the DM?
- The theming (human choices) aren't related to Lazutin at all, but that's not a Feature we chose
- The Capability of Phishing email overlaps, the theme (COVID vs. payroll) is different but doesn't change the TTP leveraged
- The victim is a neighboring bank, which makes it a peer to Money Bin
- The Dridex malware is not something we've seen but as a tactic it fits perfectly

	Discover	Detect
Recon		
Weaponization		Malicious PDF "COVID19Results.pdf" w/ Dridex
Delivery		Jan@healthhellper(dot)com
Exploit	Social Engineering	
Install	Dridex	
C2	34.32.22.11	
Actions	CobalStrike and then Ryuk ransomware	

New Intrusion – Does it Fit?

If we look at the kill chain details of the information shared from a peer bank to Money Bin, we see that the malware choice, theming of the phishing email, and human fingerprints are all different than what we've seen in K. Lazutin. This is very likely a completely different team than the one behind K. Lazutin. But realize, we aren't tracking the team behind K. Lazutin. In fact, we don't care. We explicitly prioritized the Capability and Victim feature. And while the malware is different and the theming is different, the actual tactics being employed highly overlap with the intrusion set we had, including the deployment of ransomware at the end; the ransomware, Ryuk, is also exactly what we'd expect to see in TrickBot infections. So, we don't get focused on the technical indicators of the malware choices, but that they are similar tradecraft and capability usage for what the adversary is achieving. The overall tradecraft over the adversary and the overlaps in the Victim feature are what matter.

This means we can confidently add this intrusion to our Activity Group with the other data optionally collected in case it helps provide a pivot point in the future.

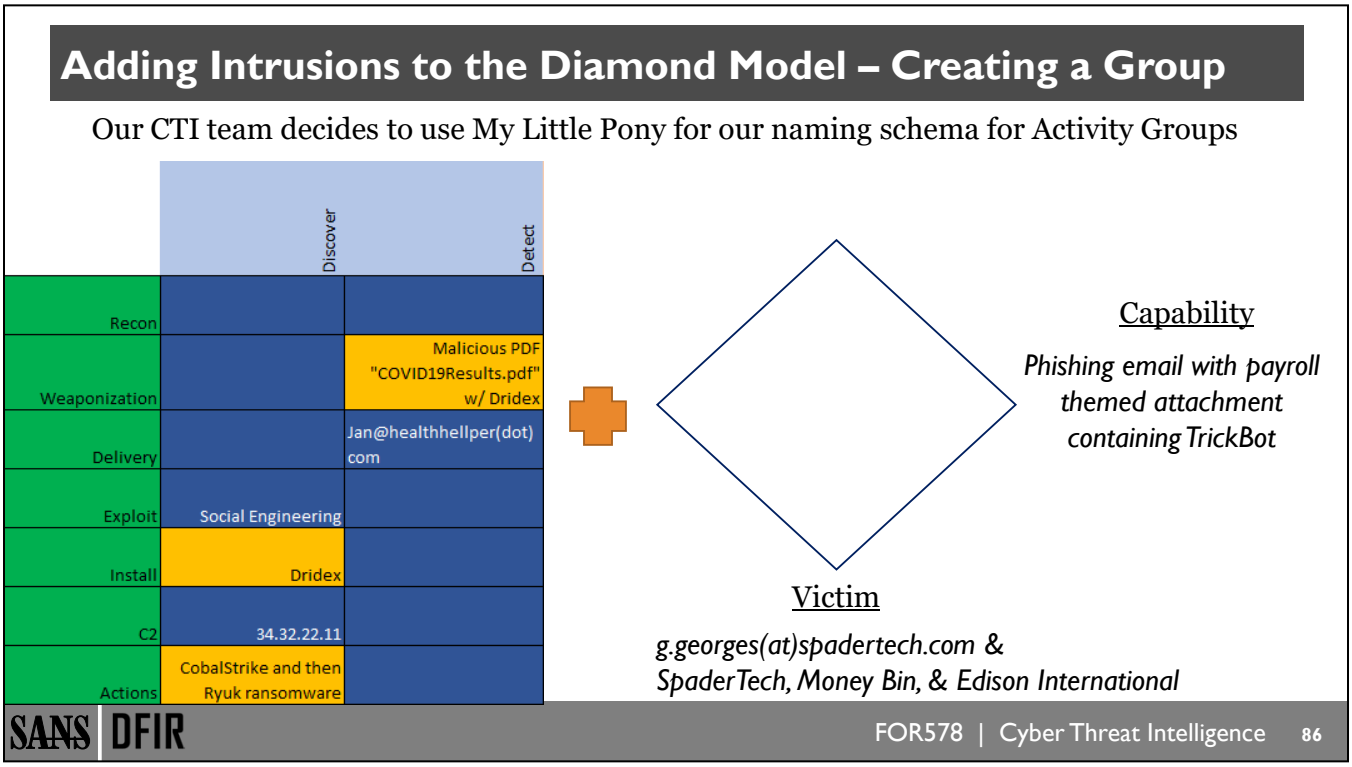

Adding Intrusions to the Diamond Model – Creating a Group

Therefore, we'd take the components of the intrusion that matter and add them to the components of the intrusion sets we were tracking before that matter to the Feature selection of the Diamond Model. Together they'd form an Activity Group. Our team decides to use My Little Pony for our naming schema.

If we saw just Dridex and didn't see any ransomware usage, it wouldn't have met our requirement to track groups that perform human operated ransomware operations. Additionally, it's our call as an analyst on how specific we want to get. As an example, we might note that different tradecraft results in different groups. Where Dridex deploying CobaltStrike and Ryuk would be different than TrickBot directly deploying Ryuk. This is where the analyst really needs to understand not only the requirement, but how the consumer wants to consume the intel; what's the dissemination look like and does that meet their goal? In this case, as long as it's human operated ransomware, it's ok for the capabilities and tradecraft to be different. That's the requirement here.

While we would add an intrusion that just had Ryuk and targeted a bank, the strongest correlations are when we are not just adding any 1 indicator to the clustering. For us to confidently add an intrusion, we really want to see multiple indicator overlaps. Ideally, there are 2 steps or more in any kill chain that overlap to ensure it's an intrusion we care about. As an example, simply using Ryuk is a bit weak (but works), but the combination of Dridex and Cobalt Strike and Ryuk makes this a strong correlation.

Introducing: PINKIEPIE

Adversary

Optionally Collected Indicators for Pivoting
(Irrelevant to the Defenders)

Infrastructure

Optionally Collected Indicators for Pivoting
(Irrelevant to the Defenders)

Capability

Phishing emails with malicious documents that contain first stage malware. Observed payroll and COVID19 theme. Observed TrickBot and Dridex.

Victim

Technology, Financial, and R&D organizations. Observed g.georges(at)spadertech.com & SpaderTech, Money Bin, & Edison International

Introducing: PINKIEPIE

Our team chooses to use My Little Pony as the naming convention for our Activity Groups. This helps us avoid biasing ourselves with the names but having an easily remembered naming convention. We clustered (defined the features) on Capability and Victim.

We could continue to collect and document information and indicators in the Adversary and Infrastructure features, though we are clustering (using them as a filter) off of Capability and Victim. Here, for simplicity and to avoid confusion, nothing is shown in the Adversary and Infrastructure features.

This analytical requirement and selection of features would allow us to cluster intrusions that we could then use to inform a set of tailored defensive playbooks. These playbooks could serve as an opportunity for prevention, detection, and response use-cases across our consumers.

What bothers most people the first time looking at clustering this way is that there could be numerous different teams, even entirely different state actors, criminals, etc., all fitting into the same Activity Group. But that's actually fantastic. If we can create an Activity Group that encompasses numerous different adversary teams where the defensive playbook and recommendations are the same, we are then saving tons of time and resources for our defenders. If that is the requirement.

Remember, the point isn't how much we can cluster and create cool groups that we track and get intimate with revealing state adversaries. The point is to satisfy the intelligence requirement. The intelligence requirement this time around has nothing to do with any element of state level attribution and is instead focused on giving intelligence-driven defense recommendations to the defenders so we can get ahead of awful ransomware operations.

Analysts first think about analyzing and clustering intrusions to track adversaries. You can do that, but that is not the mindset to take. The mindset to take is that your job is to satisfy an intelligence requirement, and the analyzing and clustering of intrusions is a means to do that. Sometimes that will relate to tracking

a specific team, sometimes it'll relate to tracking specific TTPs even if used by multiple teams, sometimes it'll focus on unique styles and approaches to gathering infrastructure, etc. Your feature selection in the Diamond Model is a component of meeting your requirement, not just an attempt to track teams.

This concludes the PINKIEPIE group formation that started all the way back with a single indicator in Exercise 2.1.

Shortcut: The Rule of 2

- One shortcut to clustering is simply applying the Diamond Model
  - Look for overlaps between two vertices in intrusions or campaigns
- The goal is to identify unique characteristics (key indicators of behavioral TTPs)
- Map the unique characteristics to the Diamond Model

Shortcut: The Rule of 2

Another way of confidently creating campaigns is to apply the Rule of 2. The Rule of 2 is simply looking for consistency in intrusions in some key way (key indicators or behavioral TTPs as an example) to create an activity group. If the victims are the same or similar, you might have also identified a specific campaign.

This is not an exact science, but, ultimately, you're choosing to hunt for adversary activity in your dataset from any one of the four vertices of the Diamond Model and then combining that with combinations of other vertex-focused hunts. Let's look over the next two slides to get a better understanding of this concept.

In this way, you are not clustering on a formalized definition, and the components of the two features can change intrusion to intrusion if you would like (they do not have to). As an example, you could identify two intrusions that overlap on the Infrastructure and Capability features and then identify another intrusion that overlaps on the Adversary and Infrastructure components of the previous two intrusions you clustered. This all could be one Activity Group. In this way, though, the Analytical Requirement is simply sorting through the intrusions to identify clusters. You are not likely satisfying an intelligence requirement of a user in doing this but instead clustering intrusions that are highly correlated together to then determine if there are observations you can further produce to help defenders.

This is a much looser approach to clustering but can be quick and still have value, especially to continue the analytical process for the analyst. Just do not get trapped in loving to find clusters more than satisfying requirements.

If you do not want the features to rotate intrusion to intrusion, that is OK as well; the rule of 2 still works where you are quickly separating your intrusions into specific clusters of selected features but without a formal definition yet. This can be used to quickly sift through lots of intrusions (e.g., 100's or 1000's) to then be able to step back and look for patterns and commonalities that pique your interest or overlap with standing intelligence requirements.

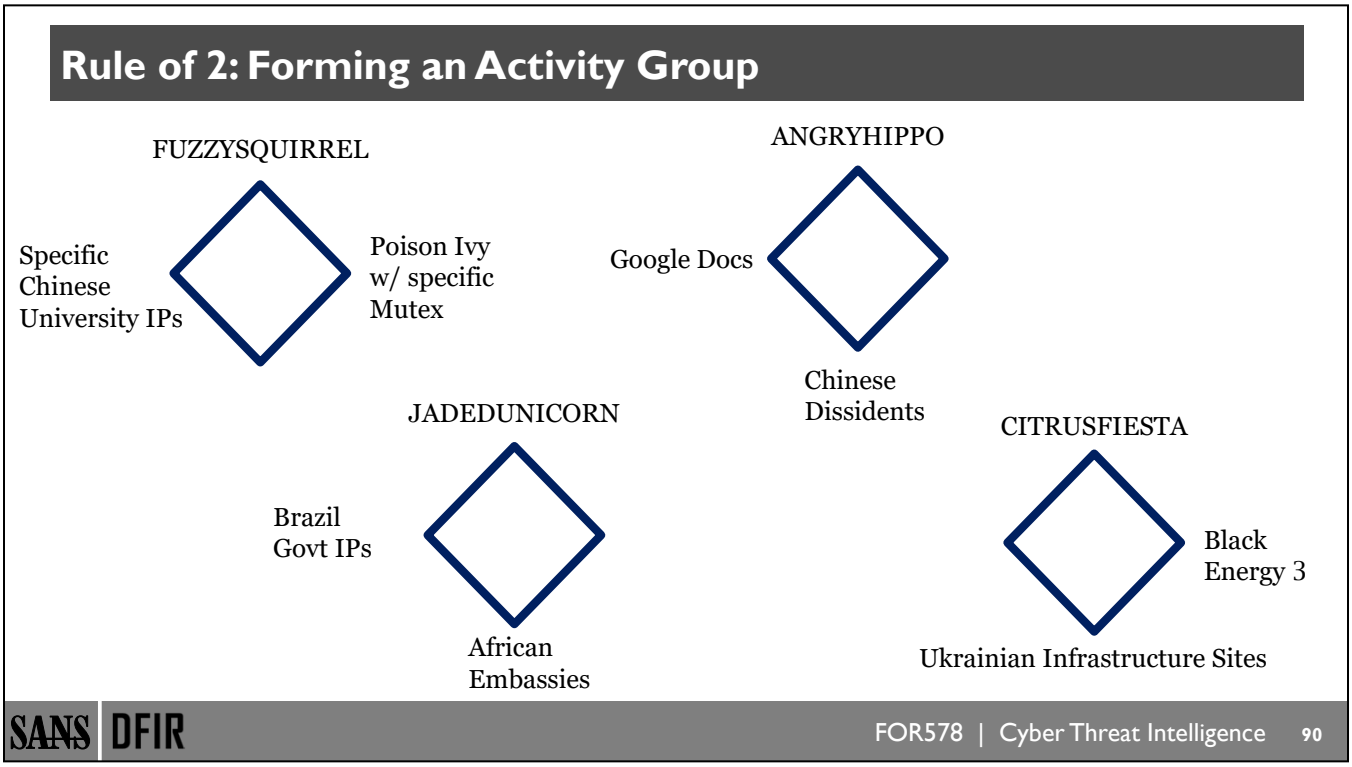

Rule of 2: Forming an Activity Group

What we are ultimately looking for is things that uniquely describe an intrusion or intrusion set and placing those key values on the Diamond Model. We are, in essence, abstracting out all the Kill Chains and Diamond Model overlaps into a single Diamond Model of unique attributes. Where we identify combinations that are consistent across multiple intrusions and have at least two shared vertices, we can cluster it into a Diamond Model.

As an example, if we look across 100 intrusions and find that 5 are using specific Chinese university IP addresses, and those 5 also all use Poison Ivy with a specific mutex, we could cluster those together and make up an activity group name for that group, such as FUZZYSQUIRREL.

We could do this across all of our dataset looking for patterns and trying to identify activity groups to track. If we can add enough data over time and also understand the adversary campaigns, we can possibly correlate different activity groups together as one larger group. In this way, it almost seems backward—we are going from groups to campaigns instead of campaigns to groups. This is okay because what we are identifying here is unique groupings that describe the "group" behind the activity, not the mission they have. Understanding the missions they conduct, though, will be vital to confidently distinguishing between groups or combining them together.

When to Retire Clusters?

- Campaign states:
  - Active
  - Inactive
  - Dormant
- Keep all information pertaining to clusters indefinitely
- Future intrusions can illuminate past:
  - Redefining clusters may fit old intrusions together more logically
- Future clusters may correlate to past clusters

When to Retire Clusters?

Another question analysts are inevitably faced with is when to retire, or obsolete, clusters. Remember, these are projections of people. The people at the other end of the wire probably haven't gone away. They've probably just reorganized and retooled to the point in which they now appear as new clusters. Sometimes, there are long gaps in activity, as long as a year or more, in which clusters have simply refocused their operations to meet objectives against a different industry or even a different part of the world.

For these reasons, it's advisable to classify clusters like volcanoes: active, inactive, and dormant. A dormant volcano leaves a huge scar on the landscape, looming over the ground below, silent. And it may remain silent forever. Or there's a chance it might again become active, at which point everyone unprepared is totally up a creek… of lava.

Keep all the information pertaining to a campaign indefinitely, including documentation on its constituent intrusions and, if possible, the corresponding raw data collected. There is at least one example of a "new" cluster correlating to a "dormant" campaign 4 or more years back that, upon closer inspection of the incidents and data, was refined to be a single cluster over 6 plus years.

A good rule of thumb:

If you feel you have good collection and actually can tell if the adversary is active from your collection:

Active = any linked intrusion within 6 months

Inactive = haven't seen linked intrusion for more than 6 months

Dormant = haven't seen linked intrusion for more than a year

If you don't feel you have great collection, then double the timelines (12 months, 12-24 months, 24+ months).

Case Study: APT10 and APT31

Group Names Aren't Names, They Are Definitions

This page intentionally left blank.

Recorded Future and Rapid7 Attributed Breaches to APT10

- Recorded Future and Rapid7 identified a breach into a Norwegian managed service provider and a U.S. Law firm on the heels of a US indictment of APT10 actors
- APT10 and the Cloud Hopper report discussed targeting of managed service providers, the use of a variety of C2 servers, and the use of Trochilus malware
- Recorded Future and Rapid7 found these overlaps and attributed it to APT10 and identified new variants, TTPs, and a Chinese shell company

Recorded Future and Rapid7 Attributed Breaches to APT10

This case study is not meant to embarrass Recorded Future or Rapid7; actually, where they made a mistake in my opinion isn't entirely their fault. It's a consistent mistake in the community when teams decide to use other CTI teams' nomenclature and names of threats not realizing they are not names, they are definitions.

Recorded Future and Rapid7 published a report stating that APT10 targeted a Norwegian managed service provider (MSP) and a number of US companies including a law firm. There was strong technical overlap and victim overlap with indicators and findings in PWC's Cloud Hopper and FireEye's APT10 reports. Even the Trochilus malware was used.

References:

- https://www.documentcloud.org/documents/5638889-APT10.html
- https://www.cyberscoop.com/apt10-apt31-recorded-future-rapid7-china/
- https://www.recordedfuture.com/apt10-cyberespionage-campaign/
- https://www.pwc.co.uk/cyber-security/pdf/cloud-hopper-report-final-v4.pdf

Group Names are Definitions not Often Publicly Known

- While there were obvious overlaps, the Cloud Hopper and APT10 "definitions" are not public, only elements of what is used
- PWC (Cloud Hopper) and Microsoft analysts called out Recorded Future and Rapid7 noting that while there were some overlaps, the methods of the adversaries, the style of registered domains, and more were actually APT31

"None of the stuff that we were tracking as APT10 overlaps with what Recorded Future and Rapid7 have reported," McConkey said. His company published a detailed account of APT10's compromise of remote IT service providers in 2017.

SANS DFIR

Group Names are Definitions not Often Publicly Known

However, Microsoft analyst Ben Koehl and PWC analyst Kris McConkey (who's an awesome person and regular at the CTI Summits, by the way) both came out noting that what Rapid7 and Recorded Future found was not APT10 but in fact APT31. There were definitely a lot of overlaps, but they were definitively different groups, even if likely both Chinese-state teams.

Notice the language of Ben: "made subsequent changes in specific ways that we attribute (with other information)." It wasn't just the technical overlap but something about *how* the adversaries maintain their infrastructure and non-public data that led to the attribution to APT31. This wouldn't have possibly been known by Recorded Future.

The Problem Isn't just a Recorded Future / Rapid7 Problem

- It would be easy to scoff at Recorded Future and Rapid7, but that's not the problem
- The victim overlaps, technical overlaps, malware overlaps, etc., were all there
- The problem is that to PWC and Microsoft who followed these intrusions closely and had their own definitions for them, the differences in one group (APT10) and another (APT31) were entirely obvious (and not much was public on APT31)
  - They were not obvious to those not intimate with those cases and definitions

> McConkey said the command-and-control infrastructure listed in the Recorded Future-Rapid7 report is that of APT31, not APT10. His team, he added, has not seen APT10 deploy Trochilus in the manner described in the report (Recorded Future and Rapid7 described it as a "new variant").

The Problem Isn't just a Recorded Future / Rapid7 Problem

The reality is this happens a lot. Every team that creates a cluster of intrusions, whether they use the Diamond Model and call them Activity Groups or have their own schools of thought and use Threat Actors and custom definitions, ultimately create unique definitions known to themselves and not often publicly stated.

Even with exact malware and victim overlaps, it simply wasn't enough in the broader dataset that PWC and Microsoft were working to attribute these intrusions to APT10 in the way that Recorded Future and Rapid 7 did.

Notice Kris McConkey's language choices: "has not seen APT10 deploy Trochilus in the manner described." Similar to the Microsoft analyst, Kris's team isn't just associating the malware as the item to cluster intrusions but something more specific about the malware, because in their dataset, the malware isn't unique enough for the definitions they are using. For example, they might see behavioral indicators and choices, human fingerprints, or similar that need to be seen before the malware choice is enough for them to use.

Everyone's a Little Wrong

- Not to insult anyone, but actually, everyone here was a little wrong
- APT10 is not the same thing as Cloud Hopper, though obviously PWC and FireEye have worked closely enough together that they feel comfortable with these links being considered a 1:1 translation
- APT31 is not actually the same as ZIRCONIUM, though obviously Microsoft and FireEye worked closely enough together that they feel comfortable to say that
- The reality is the definition of a threat is difficult enough to maintain in one team, let alone spread across multiple teams; it is even harder when it is not transparently defined
  - Everyone is taking their own view of related intrusions using different definitions for different intel requirements. Links exist. 1:1 overlaps do not.

Everyone's a Little Wrong

The reality is this happens way more than people realize. But often, when we see firms publishing, they're doing it in collaboration with the original author who checks their work. For instance, it's a safe assumption that before PWC published on Cloud Hopper, they coordinated with Microsoft or FireEye, which was able to check some of their work for them and confirm what they were tracking was APT10. However, it is unlikely that every intrusion that went into Cloud Hopper actually met the definition of APT10. The reality is there's a lot of analytical baggage that abstracts these things away.

But the problem gets worse inside everyday companies. When your CTI team is tracking intrusions and tries to attribute them to known groups of other teams and vendors, you're likely making the very same mistakes that Recorded Future and Rapid 7 made, because you don't have intimacy on the group definition the other teams are using. But because you're not publishing your cases publicly, you likely don't get called out for it.

It's always a best practice to use your own group names if you have your own unique collection and intrusions (don't just rename groups for the purpose of renaming them). It's perfectly fine to say there are "links" or "overlaps" with other known groups. You should call that out, but what you are tracking, against what collection, against what requirements, with what analytical tools, with your biases, etc., is never a 1:1 overlap with someone else's definition.

This page intentionally left blank.

"Top Energy" Intrusion

- Previously identified "Intrusion 2" in Day 2 compromising SpaderTech
- The SOC has detected two new intrusions that may be linked to this intrusion
- Analytical Requirement: Define a cluster focused on unique victimology type and tradecraft to inform defense recommendations
- Feature Selection:

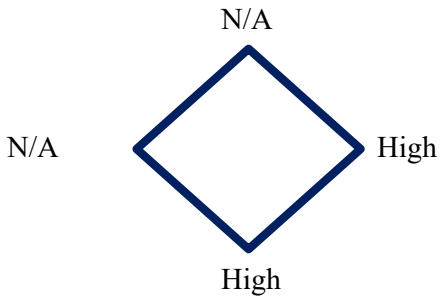

"Top Energy" Intrusion

For the Top Energy intrusion we worked, which was the Intrusion 2 that was in the Exercises on Day 2 and compromised SpaderTech, we want to figure out if this intrusion is linked to other intrusions.

However, if the requirement is: Is this linked? Then any pivot, any connection, any overlap, would link the intrusions. That doesn't actually solve any problem other than creating an interesting looking dataset. Pivoting will find pivot points but may not satisfy any requirement.

What is challenging, what is necessary, is ensuring that you set a requirement first and then only focus on the data points that meet the requirement. Staying focused helps you satisfy the requirement.

Here, our requirement is to identify clusters of intrusions specific to the victimology that we care about (SpaderTech and companies like it focused on R&D, engineering projects, and energy related companies) and the unique tradecraft used by the adversary (tools, tactics, and techniques).

So, if we analyze other intrusions and they have overlaps on Infrastructure or Adversary features, then it wouldn't be enough to link the intrusions together. On the intrusions that do have overlaps on Capability and Victimology though (we'd need both not just one) then it would be completely appropriate to record Infrastructure and Adversary feature data as well. We are not limiting our collection. We are applying a filter on the intrusions we care about to stay focused on the intelligence requirement instead of the natural tendency of analysts to analyze.

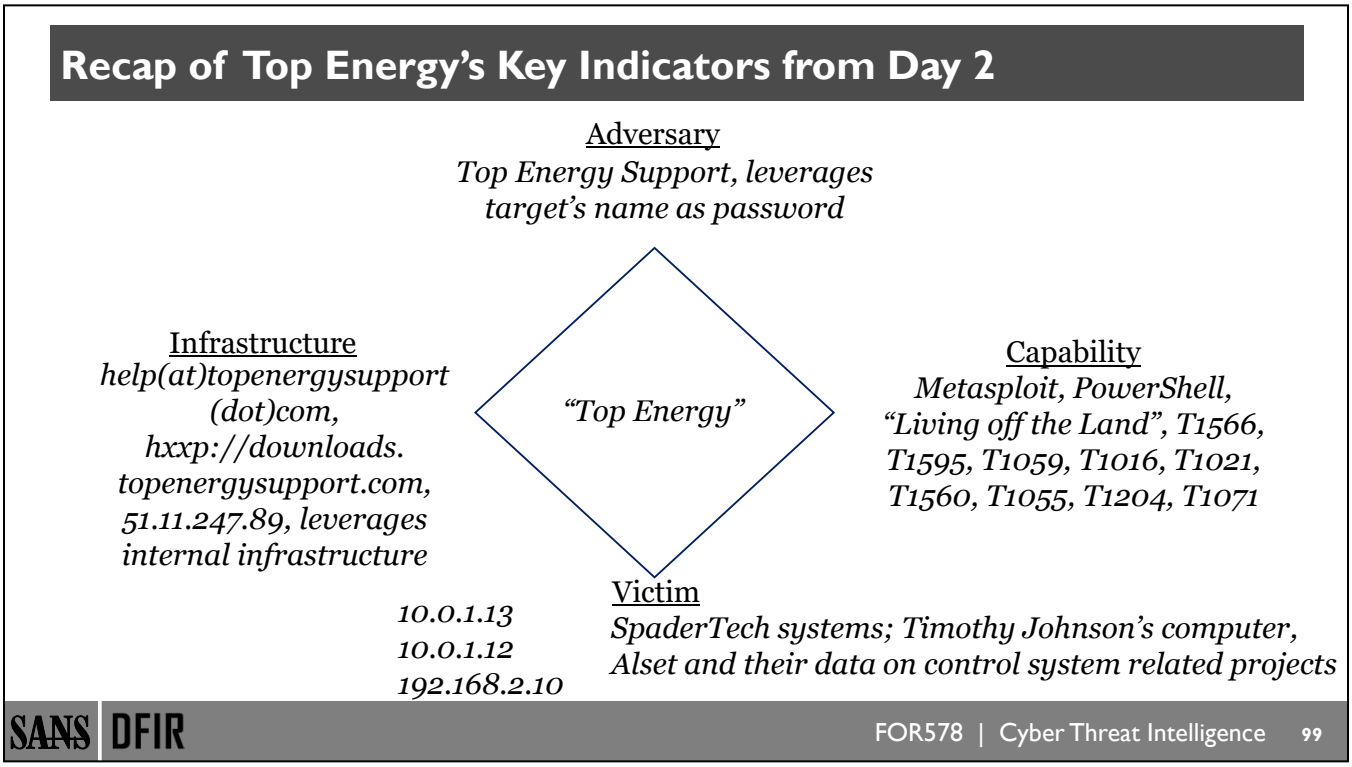

Recap of Top Energy's Key Indicators from Day 2

Here's a singular view on the key indicators and diamond model characteristics from "Top Energy" as a recap from what was showed previously. What we care about is determining what in the Capability and Victim components overlap with the new intrusions that we're introducing.

New Intrusion #1 Key Indicators

- Edison International Security Operations Center
- Target: Acme Power
- Delivery: Phishing email from "help@cleanenergysuport[.]com" with malicious URL "downloads[.]cleaneenergysuport[.]com"
- Installation: Encoded PowerShell
- C2: 38.72.114.16
- Actions on Objective:
  - T1021 – Remote Services
  - T1560 – Archive Collected Data

Analyst Notes:
The adversary targeted Acme Power's internal systems that contained engineering drawings. They pre-positioned on an internal system for exfil and the password on the archive was AcmePower. Leveraged PowerShell heavily.

New Intrusion #1 Key Indicators

This is a new intrusion our SOC has informed us about. They captured the details of the Target, the Kill Chain steps they observed, and any analyst notes.

The target is Acme Power, which does overlap with the Victim feature we care about. The C2 isn't something we're going to focus on as infrastructure, though, if this intrusion does fit then we would record it as well. The PowerShell usage does overlap with Top Energy as well. Further, the misspelling of "support" in the email address and the choice of "help@" for an energy themed domain is very interesting. The infrastructure itself, though, again doesn't matter. The adversary choice doesn't matter. It's very helpful to see these connections, but they alone wouldn't be enough to link the intrusions. But the Adversary's choice indicating something about the Victim (energy focused, as is proven by Acme Power) and the Capability overlapping so well (including the TTPs observed in MITRE) makes a strong case. The analyst notes that the internal systems contained engineering drawings and the files were pre-positioned on an internal system for exfil really helps us fit this intrusion.

The password being the same as the target name is also identical to what happened with "Alset"; however, that's an adversary choice. What this helps is make the case that the adversaries may be the same—but that's not the point nor the intelligence requirement here. So, it's an interesting observable but not what we're trying to solve. Based on our requirement, the attribution of "who" did this isn't important. So even if there were 15 different teams targeting our organization, it doesn't matter to us; what matters is if we can create a reliable defense strategy against a very clear cluster that we're seeing and choosing to prioritize. And if focusing on this cluster allows us to protect ourselves from numerous teams operating the same way, that's a huge win.

New Intrusion #2 Key Indicators

- Edison International Security Operations Center
- Target: Money Bin
- Delivery: Phishing email from "Todd(at)wholesalesuplies[dot]com" with malicious URL "downloads[.]wholesalesuplies[dot]com"
- Installation: EvilGrab malware as ~.exe
- C2: 185.117.88.80
- Actions on Objective:
  - Utilized shared folders and native functionality to steal credit card numbers

Analyst Notes:
The adversary targeted Money Bin and stole credit cards as well as account numbers and information on high worth individuals banking at Money Bin

New Intrusion #2 Key Indicators

In Intrusion 2, we see that the target is our financial services company, Money Bin. This doesn't meet our victimology, so we're pretty confident we're not going to use this intrusion in the cluster. However, that doesn't mean we can't see if indicators are useful to pivot around and find other intrusions we care about.

As an example, the misspelled "supplies" in the email address is interesting and the TildeDrop executable overlays nicely with what we saw in the slide scenario on Day 2. But even if so, this would belong to a different cluster at that point. The Capability and Victimology are just too different. And if we pivot around on the indicators and take all of our insights from Day 3 and Day 4, we can realistically come to the conclusion that this group is linked to Cloud Hopper and APT10. We could also link this with the other intrusions we're tracking, but again that would satisfy an intelligence requirement that's completely different than the one we're trying to solve, and the clustering of all those intrusions would not have tailored any actionable recommendations for the defenders against their requirements.

Which Intrusion Overlaps?

- Both intrusions overlap with both the "Leet" intrusion from the slides as well as the "Top Energy" intrusion from the slides and exercises in Day 2
  - You ***could*** make the case that these are all related to CloudHopper
  - However, you ***would not*** satisfy the analytic requirement laid out and instead would be just clustering on all possible pivots and opportunities thus not satisfying your intelligence requirement

- Intrusion 1 and Intrusion 2 both contain misspellings in domains as we saw in the Top Energy intrusion
  - Only Intrusion 1 was targeting engineering like projects
  - Only Intrusion 1 was leveraging PowerShell and similar TTPs
  - Our analytical requirement didn't prioritize Adversary or Infrastructure overlaps so any such overlaps in Intrusion 1 and Intrusion 2 are irrelevant

Which Intrusion Overlaps?

A very difficult concept for many is that just because you have two clusters related to one adversary doesn't mean you need to group them together. Further, even if it appears to be "one" adversary, you don't know behind the intrusions what's really happening. You could easily have two different teams (or more) sharing a development team, an infrastructure team, or even multiple adversaries working together on operations. Just like our PROMETHIUM case in the beginning of the class, there are also mercenaries and hackers for hire that get involved in operations, as there are allies and others. Clustering intrusions for the purpose of clustering doesn't help anyone. If you do have the need to track a specific adversary (while achieving attribution on who they are or not), that is perfectly fine to pay attention to such links, but you'd be prioritizing based on the Adversary feature not just pivoting off of Capability and Victim overlaps.

Here, we can assess with confidence that Intrusion 1 overlaps with the Top Energy intrusion we previously tracked. Now that we have unique observables and two intrusions, it's completely appropriate to form our Activity Group.

Introducing: RAINBOWDASH Activity Group

Adversary

Energy themed naming conventions, leverages target's name as password, link to Cloud Hopper

Infrastructure

Emails themed as help desks at energy themed domains with malicious URLs from the downloads[dot]domain.com, leverages internal infrastructure to stage files for exfiltration

Capability

Metasploit, PowerShell, "Living off the Land", T1566, T1595, T1059, T1016, T1021, T1560, T1055, T1204, T1071,

Victim

Energy companies and those companies working on their projects, engineering drawings and project details

SANS DFIR FOR578 | Cyber Threat Intelligence 103

Introducing: RAINBOWDASH Activity Group

Our team chooses to use My Little Pony as the naming convention for our Activity Groups. This helps us avoid biasing ourselves with the names but having an easily remembered naming convention. We clustered (defined the features) on Capability and Victim. We will continue to track Infrastructure and Adversary features as well of the intrusions that meet the Victim and Capability requirements. This will allow us to continue to build out and flesh out RAINBOWDASH so that it forms a lens to view future intrusions and as a mechanism to satisfy our analytical requirement and develop defensive strategies.

Because the focus is Victim and Capability, we would be mistaken to think that RAINBOWDASH is a 1:1 relationship with any adversary. It is a cluster that is essentially a defensive playbook. Multiple adversaries could fit this playbook. Which is hugely valuable to defenders. It's also important to note that RAINBOWDASH has a link to CloudHopper. That is not saying this group is the same as CloudHopper; it's just being analytically honest as well as helpful to analysts with other intel requirements that this group does share overlaps with that other group.

If we really wanted to track the actor behind this group, we would define that as an analytical requirement, select features such as Adversary and Capability, or Adversary and Infrastructure, or Adversary and Victim (you can choose more than two features as well but, in this example, Adversary would be involved). But you wouldn't just start tracking RAINBOWDASH in that way. You'd create an entirely new Activity Group and leverage what you knew of these intrusions to fit that filter. Or said differently, two intrusions could create two different Activity Groups depending on your intelligence requirements and what you're tracking. Just make sure it's a requirement that you need; don't create Activity Groups just for the purpose of doing it, and retire those groups when you're done with the analytic requirement or you stop seeing intrusions for a prolonged period of time that do not meet the group's definition.

Remember, group names are not names in cyber threat intelligence. They are definitions.

Exercise 4.5 Lead-in

- In Exercise 4.5, you will be given information about intrusions
- You will look to cluster those intrusions together into Activity Groups
- You will define the Activity Groups based on the requirements given
- This lab will draw on the Leet intrusions from the Slides in Day 2 as well as key indicators we observed across Day 3 and Day 4

- You will use both a formal definition as well as the Rule of 2

This page intentionally left blank.

Recap of Leet Intrusion Set

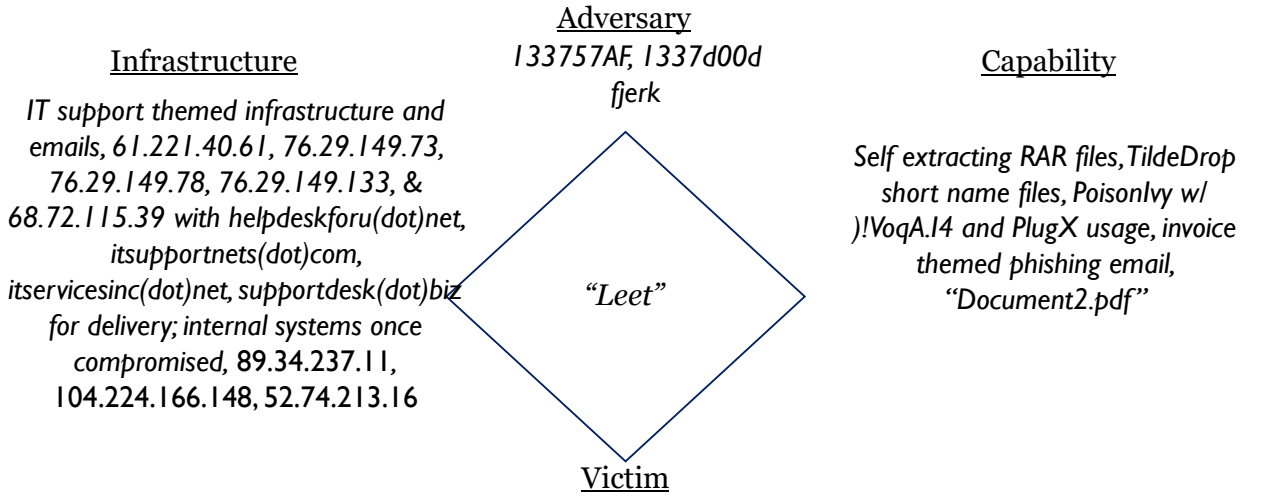

Recap of Leet Intrusion Set

We will take a subset of the "Leet" intrusion set and leverage these in Exercise 4.5. Our intent is to start storing the information for long-term usage, especially as this set of intrusions seems to be persistent against our organizations.

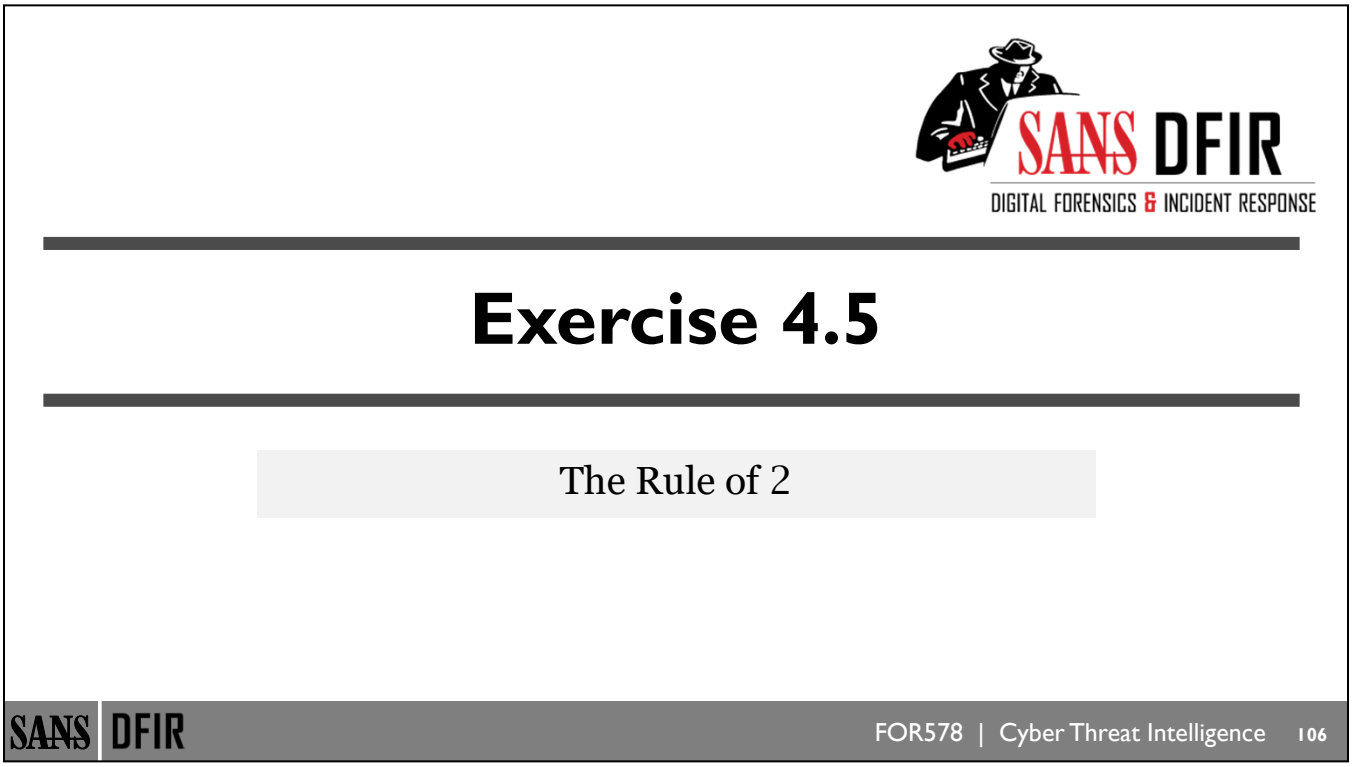

This page intentionally left blank.

This page intentionally left blank.

COURSE RESOURCES AND CONTACT INFORMATION

Here is my lens. You know my methods. —Sherlock Holmes

AUTHOR CONTACT
Robert M. Lee: @robertmlee
RLee@Dragos.com
Rebekah Brown: @PDXbek
pdxbek@gmail.com

SANS INSTITUTE
11200 Rockville Pike, Suite 200
N. Bethesda, MD 20852
301.654.SANS(7267)

DFIR RESOURCES
digital-forensics.sans.org
Twitter: @sansforensics

SANS EMAIL
GENERAL INQUIRIES: info@sans.org
REGISTRATION: registration@sans.org
TUITION: tuition@sans.org
PRESS/PR: press@sans.org