578.5
Dissemination and Attribution

SANS

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

FOR578.5 | Cyber Threat Intelligence

Dissemination and Attribution

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 books *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 United States 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.

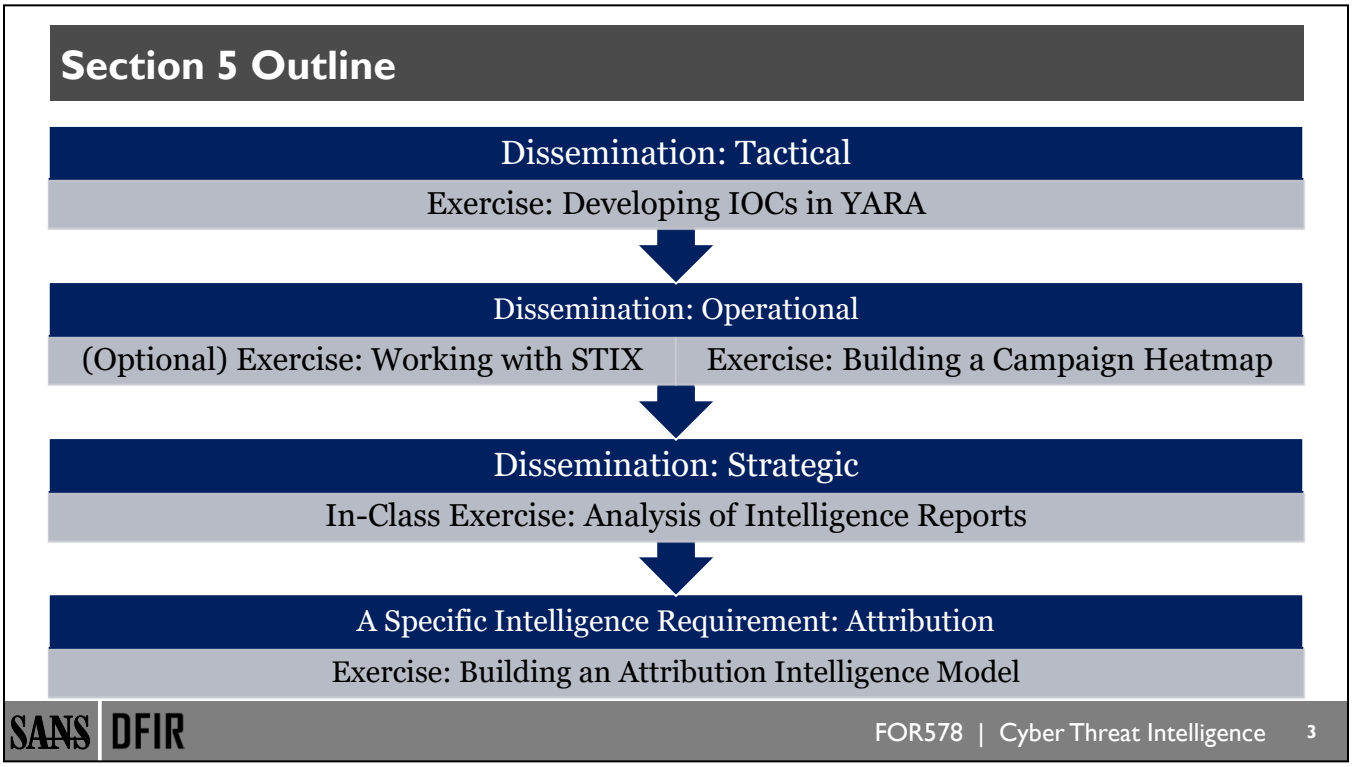

This page intentionally left blank.

Case Study: Axiom

This page intentionally left blank.

PlugX

- Originally a Chinese-based piece of malware that then propagated to other threat groups
- Fairly simple multistage RAT that has C2, file upload, download, keylogging, etc., capabilities
- Been observed in multiple APT campaigns
  - Emerged publicly around 2012
  - Linked to and seen as an evolution of Poison Ivy due to campaign overlap between actors using both
- Campaigns largely focused on government organizations and espionage

PlugX

PlugX has been around since about 2012 and was originally used exclusively by Chinese-based threat groups. It then spread to other groups as well. It has largely been seen as the successor to Poison Ivy; however, both tools are still in operation. The campaigns leveraging PlugX usually target government organizations and utilize the tool as a piece of espionage malware.

Reference:

- https://news.sophos.com/en-us/tag/plugx/

Hikit Malware

- Hikit (aka McRat) is a piece of malware that first came into detection around 2011
- Capabilities include
  - Rootkit functionality
  - Client tools for RAT functionality
  - Kernel driver to monitor traffic
- Does not connect to a C2 server but instead waits for the attacker to connect to it over HTTP or HTTPS
  - Uses a specific HTTP GET request to initiate communication

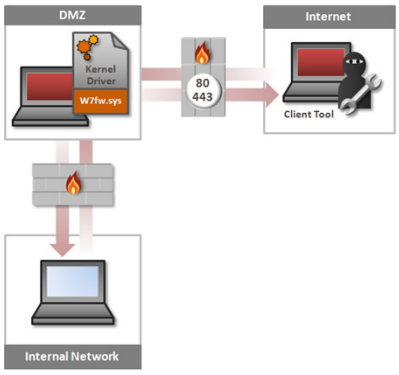

Hikit Malware

The Hikit malware is an interesting piece of malware due to its method of adversary interaction. Instead of connecting directly to a command-and-control server, the malware waits for the adversary to interact with it. This means that clients behind a firewall are more difficult for the adversaries to get initially, so it has been observed that the DMZ systems themselves are usually the initial target. The malware was also seen as having state-related use cases in targeting organizations for espionage purposes.

References:

- https://www.fireeye.com/blog/threat-research/2012/08/hikit-rootkit-advanced-persistent-attack-techniques-part-1.html
- https://community.broadcom.com/symantecenterprise/communities/community-home/librarydocuments/viewdocument?DocumentKey=171b1e21-4dcb-4635-a37b-6c95ea296611&CommunityKey=1ecf5f55-9545-44d6-b0f4-4e4a7f5f5e68&tab=librarydocuments

Hikit Malware and Bit9

- Hikit (aka McRat) was malware linked to the compromise of security vendor Bit9 in 2012
  - The malware stole Bit9 private certs so that it could sign its malicious software to bypass application control solutions
- Compromise began with an internet-facing web server that was hit with an SQL injection attack
  - Adversaries then pivoted into the environment
- The case was interesting because the adversaries wanted to compromise Bit9 customers and needed to bypass their security vendor first
  - Dedication/logistics of adversary was fairly sophisticated

Hikit Malware and Bit9

The most high-profile use of Hikit came in 2012 when it was used in a campaign against Bit9. The security company Bit9 sells an application control solution. Once in the networks of Bit9, the Hikit malware was used by the adversaries to steal certificates so that the malware could be digitally signed to bypass application control technologies in the targeted networks. This showcases that Bit9 was a victim but not the target of the campaign. The victims impacted obviously were profiled before to determine that they used Bit9 and then forced the adversaries to target Bit9 to gain access into the target networks. This is not only a hallmark of a good adversary using long-term logistics and operations planning, but also an interesting case study of the obstacles that a solution like application control can place in an adversary's path. No defensive solution is enough, but it is a good thing if your adversary has to compromise your vendor to even begin the intrusion into your network.

Axiom

- Multi-company team (Novetta, iSight, Bit9, Cisco, etc.) focused on a complex espionage program
  - Attributed to the Chinese Intelligence community
- Key findings included a well-resourced team that:
  - Existed for over six years
  - Targets including governments, NGOs, strategic economic interests, energy organizations, R&D, and infrastructure
  - Leveraged a variety of tools over the course of the campaign, including PlugX, Hikit, GhostRat, Poison Ivy, Hydraq, DeputyDog, and Derusbi

Axiom

The Axiom group was identified as part of a coalition: Operation SMN. The coalition was between companies such as Novetta, iSight, Bit9, Cisco, FireEye, and others that took part in analysis and data sharing of the Axiom group. This campaign highlighted a six-year-plus-long campaign by Chinese-based adversaries. The campaign, attributed to the Chinese intelligence apparatus, targeted governments, Non-Government Organizations (NGOs), and strategic interests of China.

Interestingly, the campaign leveraged various styles of malware, including PlugX and Hikit.

Reference:

- http://www.novetta.com/wp-content/uploads/2014/11/Executive_Summary-Final_1.pdf

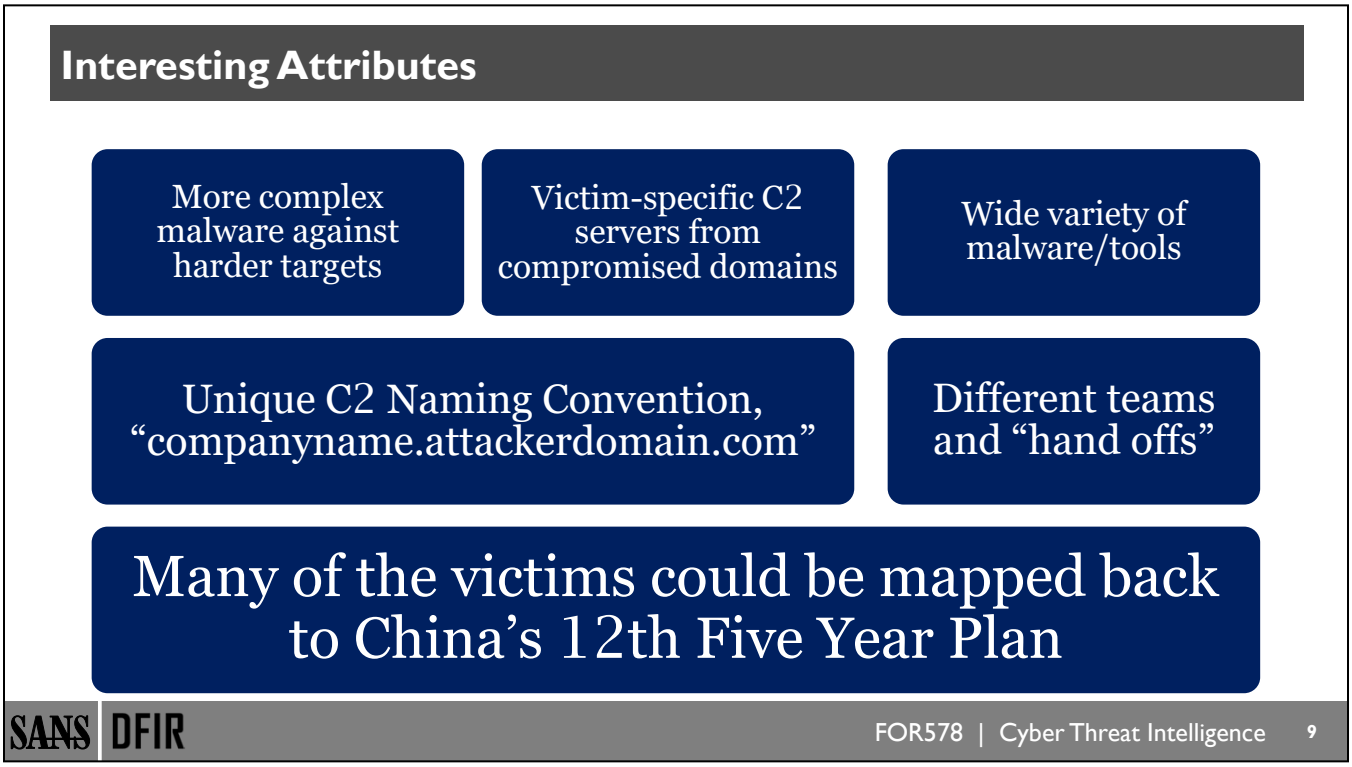

Interesting Attributes

Consider these interesting attributes:

- Numerous C2 domains were named with a similar pattern, such as "companyname.attackerdomain.com"
- Compromised infrastructure to use for specific victims and small organizations (victim-specific C2)
- Utilized more complex malware against more hardened targets and less specialized and capable malware against softer/easier targets
- Many of the victims of the campaign could be mapped back to China's 12th Five Year Plan. This speaks to the geopolitical context of many cyber espionage operations.

One interesting aspect was that the adversary seemed to have different teams. That is, there were teams that used distinct characteristics to get into a system and then seemingly handed off access to another team, thus breaking down operations into development, access, and operations/sustainment.

Reference:

- http://www.novetta.com/wp-content/uploads/2014/11/Executive_Summary-Final_1.pdf

Lessons Learned

For the context of cyber threat intelligence, there were some key lessons learned that we can extract from the Axiom group and the effort to profile them. First, the campaign took place for at least six years, speaking to the amount of historical data needed to analyze the different intrusions and malware samples. Second, the campaign was focused on strategic interests of China, and the geopolitical context of the espionage could be aligned with any different espionage effort—the fact that it took place in "cyber" is just the method of the spying. Organizations should look to see if they fall into the target range of the countries that put out strategic interests that include the industries of their organizations. Third, the uncovering of this massive campaign took multiple security vendors with unique datasets working together. It was not independent and standalone analysis. Fourth, many different intrusions, malware samples, and campaigns were tied back to the same threat group with enough analysis over time.

Dissemination: Tactical

This page intentionally left blank.

Know the Audience

- #1 key to sharing threat intelligence:
  - Know your audience
  - The audience shapes the delivery:
    - Different audiences have different intel needs
    - Different audiences require data in different formats

Pretty pictures and maps on an SOC operations screen are usually more for visitors than the SOC analysts

Know the Audience

As mentioned earlier today, one of the most important parts of sharing threat intelligence is to know your audience. Business executives care more about organizational impact, allocation of company budgets and resources, and return on investments than they do about a software patch or vulnerability related to a new server. They may be related, but the terminology and impacts highlighted depend largely on the audience. It is not your language you need to speak; it is theirs.

Tactical threat intelligence should be presented to those tactical level defenders who will be able to learn from the action or information being shared. The focus is generally more on threat data than finalized threat intelligence, usually shared through indicators.

Image:

- http://threatbutt.com/map/

YARA

- YARA is a pattern-matching tool used to create signatures
- YARA has been compared to the grep command:
  - Extremely flexible including capability to use regular expressions and wildcards
- Capability to add metadata, descriptions, and titles to YARA rules allow the rules to be quickly shared
- YARA is perfect for incident response in that it allows a quick initial triage against acquired data to determine if there is reason to focus on the collected data:
  - Helps to identify the scope of the infection and have confidence that a system is clean
- Threat intelligence analysts will likely see and handle YARA:
  - Not necessary to be a YARA expert but familiarity will help

YARA

YARA is a pattern-matching tool used to create signatures that can identify patterns in data. It is a favorite of malware analysts and is becoming increasingly more popular in the community. It's flexibility and "advanced Grep" type functionality makes it easy to share with other analysts and to search for complex patterns of data inside of other data, whether that is a memory capture or a packet capture. (Some plugins are required for searching packet captures.) To learn more about YARA, you can reference the links here:

References:

- http://www.deependresearch.org/2013/02/yara-resources.html
- http://plusvic.github.io/yara/
- https://yara.readthedocs.io/en/latest/
- https://support.virustotal.com/hc/en-us/articles/360001293377-Retrohunt
- https://gist.github.com/Neo23x0/e3d4e316d7441d9143c7

Sample YARA Rule

```
rule silent_banker : banker
{
    meta:
        description = "This is just an example"
        thread_level = 3
        in_the_wild = true

    strings:
        $a = {6A 40 68 00 30 00 00 6A 14 8D 91}
        $b = {8D 4D B0 2B C1 83 C0 27 99 6A 4E 59 F7 F9}
        $c = "UVODFRYSIHLNWPEJXQZAKCBGMT"

    condition:
        $a or $b or $c
}
```

Reference: YARA GitHub

Sample YARA Rule

This is the typical standard YARA rule. Notice that it has a name at the top, metadata to describe the rule, strings of data it's searching for, and then a condition to determine when the rule should alert. This is a basic format of YARA, although more complex patterns are possible, as will be discussed over the next few slides.

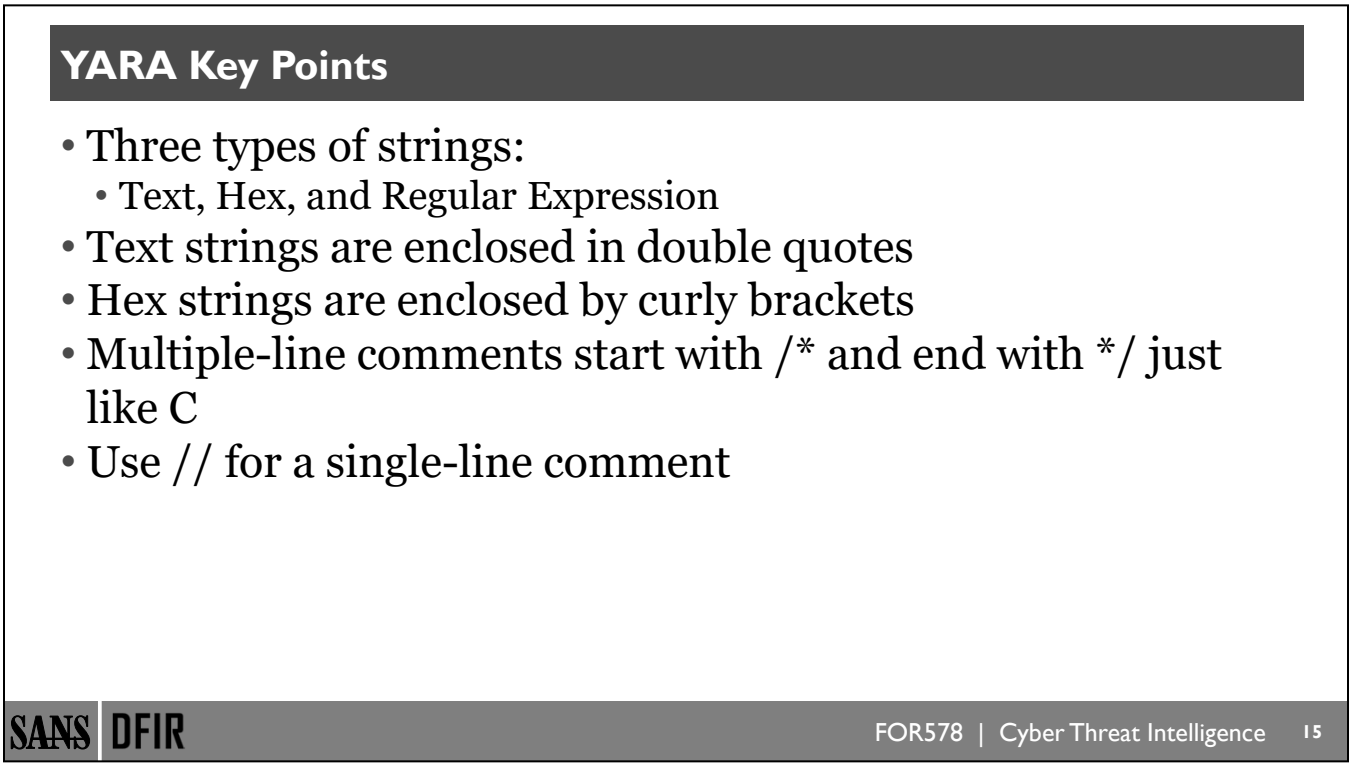

YARA Key Points

YARA is a very easy-to-understand and C programming-like language. Reading the YARA documentation does not take long and can be of great value: https://yara.readthedocs.io/en/v3.4.0/.

There are some key points to keep in mind when writing YARA rules. First, there are three types of strings that you can make rules out of: Text strings, Hex strings, and Regular Expression strings. Text strings are enclosed in double quotes, Hex strings are enclosed in curly brackets and can include spaces between bytes or do without them, and Regular Expression strings are enclosed in either double quotes or curly brackets, depending on the type of data (text or hex) being used.

To include comments, you can enclose a multiple-line comment in a /* */ or in a single-line comment, you can just use //

Multiple Types of Rules

YARA allows you to create multiple types of rules. For example, the "global rule" condition allows you to create a rule that applies to all other rules run. This allows you to tailor large sets of rules with a single rule before you run them against files. As an example, if you only want YARA rules to match against a certain type of file, such as a .exe, a global rule could be made to set a condition declaring that files must match the .exe type. Then all the other rules run would only match against such files. Additionally, private rules can be created. Private rules do not alert and thus seem useless. However, rules can import other rules and build upon them, which means that private rules can be used as a triggering event for other rules. As an example, one YARA rule might look for a suspicious C2 server, but to limit false positives, a private rule is made that looks for other indicators, such as a known malicious process. The C2 YARA rule then could import the malicious process YARA rule and only alert when both rules match. Lastly, rules can take advantage of tags to help analysts filter out rules. For example, an "APT" tag could be made to help analysts focus on only those types of threats faced.

Hex Special Values

- Hex contains question marks for wildcards
- Jumps can tell you how many bytes can exist before the next sequence is seen
- Alternatives allow for an OR-styled Boolean comparison

```
rule JumpExample
{
     strings:
        $hex_string = { F4 23 [4-6] 62 B4 }

     condition:
        $hex_string
}
```

```
rule AlternativesExample1
{
     strings:
        $hex_string = { F4 23 ( 62 B4 | 56 ) 45 }

     condition:
        $hex_string
}
```

More Complex YARA Rules

Reference other rules

```
rule Rule1
{
    strings:
        $a = "dummy1"

    condition:
        $a
}

rule Rule2
{
    strings:
        $a = "dummy2"

    condition:
        $a and Rule1
}
```

Import modules

```
import "pe"

rule test
{
  strings:
    $a = "some string"
  condition:
    $a and pe.entry_point == 0x1000
}
```


More Complex YARA Rules

There are a number of ways to make more complex YARA rules. A few examples will be shown across the next few slides to note options such as declaring file sizes. Here, though, we see the ability to reference other rules inside the condition of a YARA rule. This allows for very tailored rules and for spidering rules out based on different sets of YARA rule families. For example, a script could be created to run one set of YARA rules against a file sample to determine the type of file being interacted with such as an executable. From there, a set of YARA rules could be applied if it was, in fact, an executable file. From there, general rules could be applied across different families of malware, which in turn would cause other YARA rules to be compared against the sample for more specific types of the malware family identified. As an example, a combination of rules could take a file, determine it was an executable, determine it fit into the PlugX malware family, and then match YARA rules to see specifically which sample was being used. This could be done for automation purposes to reduce false positives of simply running all the rules against all the samples encountered.

Importing modules allows you to import portable executable (PE) information, Cuckoo sandbox information, digital hashes, and more. Additionally, there is an easy-to-use module guide to create your own modules for tools or datasets you have. Modules allow YARA rules to move past more simple string matching to having more context such as identifying a digital hash of a string.

Sample YARA Rule: Uncommon File Size

```
rule Suspicious_Size_chrome_exe {
    meta:
        description = "Detects uncommon file size of chrome.exe"
        author = "Florian Roth"
        score = 60
        date = "2015-12-21"
    condition:
        uint16(0) == 0x5a4d
        and filename == "chrome.exe"
        and ( filesize < 500KB or filesize > 1300KB )
}
```

Reference: Florian Roth

Sample YARA Rule: Uncommon File Size

In this example, Florian Roth created a YARA rule to detect suspicious sizes of known files. To create the list of file sizes, he downloaded samples of malicious files from VirusTotal and determined what the normal range in KB was of good files. For example, he identified that chrome.exe is usually between 500–1,300 KB in size but often 10–500 and 1,300 or larger was observed as malicious. This is a great way to make an initial YARA rule that looks for potentially malicious activity. This is not a high-confidence IOC but is useful for hunting for potentially malicious activity.

In the YARA rule, note the "uint16(0) == 0x5a4d." uintXX in YARA designates 8-, 16-, or 32-bit unsigned integers to perform an offset or virtual address from. In this case, the HEX "5a4d" is looking for the "MZ" header at offset 0 that is associated with portable executable files. (Remember that uint uses little endian, so the header may appear reversed as "ZM" when decoding.)

We also see other aspects of YARA here, including the ability to dictate file size and filenames. Here, the YARA rule is looking for "chrome.exe" with a valid MZ header and only smaller than 500KB or larger than 1,300KB.

References:

- https://www.nextron-systems.com/2015/12/22/yara-rules-to-detect-uncommon-system-file-sizes/
- https://yara.readthedocs.io/en/v3.4.0/writingrules.html

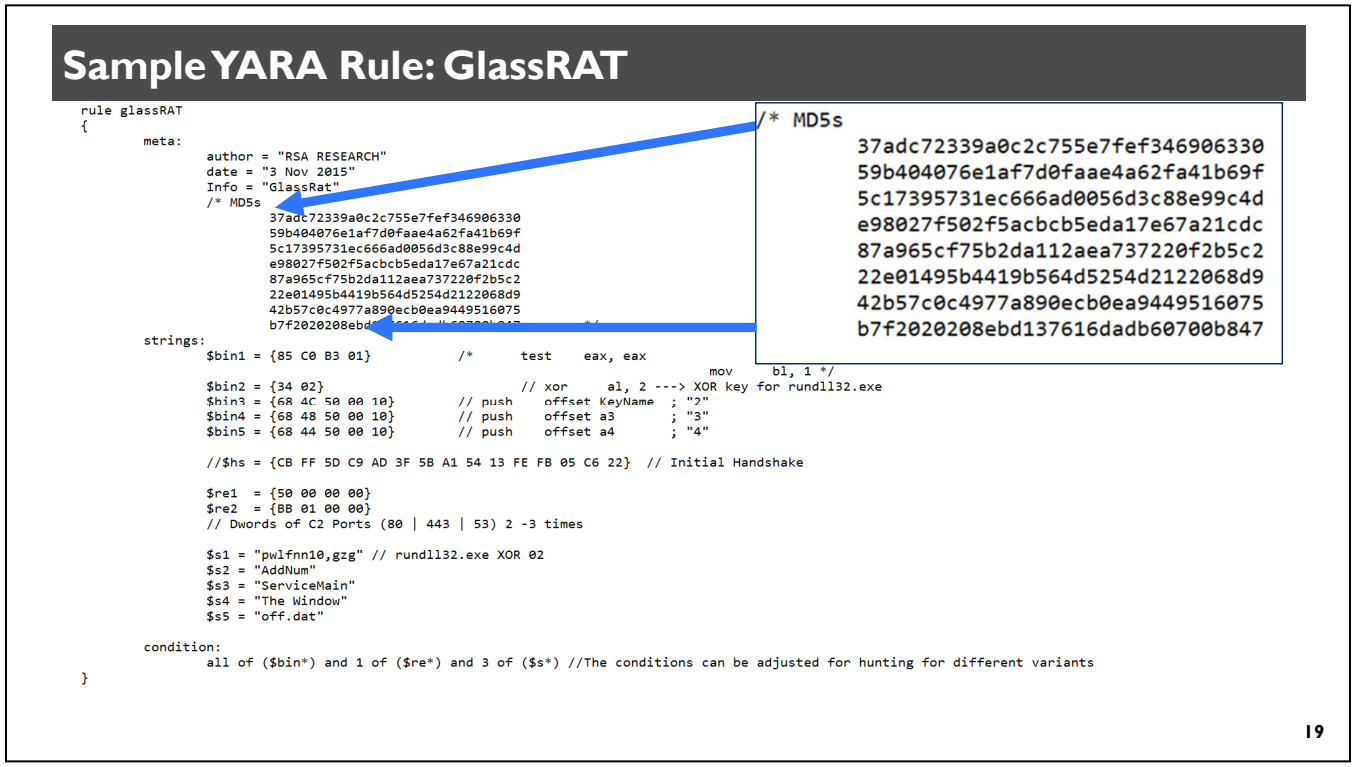

Sample YARA Rule: GlassRAT

This is one of the YARA rules that RSA released for GlassRAT. There are a few important things here to highlight that were done extremely well.

First, in the metadata, the YARA rule specifically gives the MD5 hashes of the samples of GlassRAT analyzed. This helps other analysts know what the YARA rule should and possibly should not (samples that aren't included) work against. Additionally, it gives the samples that analysts can find and analyze themselves or test their other rules against.

Second, the YARA rule has comments (denoted by the "//" for single-line comments) that are for the analysts who review the rules to identify what each string is encompassed of.

Third, the rule segments different types of strings into different groupings based on the variables. As an example, the rule has $bin for like items, $re for like items, and $s for like items. This allows a really nice condition. The condition in this rule denotes that all of the strings in the $bin variable (notice the * for a wildcard that would include all the numbers after $bin), one of the $re strings, and three of the $s strings must be present. This is an excellent example of a tailored and focused rule that still has flexibility.

For easier viewing, here is the content of the YARA rule:

```
rule glassRAT
{
        meta:
                author = "RSA RESEARCH"
                date = "3 Nov 2015"
                Info = "GlassRat"
                /* MD5s
                        37adc72339a0c2c755e7fef346906330
                        59b404076e1af7d0faae4a62fa41b69f
                        5c17395731ec666ad0056d3c88e99c4d
                        e98027f502f5acbcb5eda17e67a21cdc
                        87a965cf75b2da112aea737220f2b5c2
                        22e01495b4419b564d5254d2122068d9
                        42b57c0c4977a890ecb0ea9449516075
                        b7f2020208ebd137616dadb60700b847        */
        strings:
                $bin1 = {85 C0 B3 01}                   /*      test eax, eax
                                                                mov bl, 1 */
                $bin2 = {34 02}                                 // xor  al, 2 ---> XOR key for rundll32.exe
                $bin3 = {68 4C 50 00 10}        // push  offset KeyName  ; "2"
                $bin4 = {68 48 50 00 10}        // push  offset a3       ; "3"
                $bin5 = {68 44 50 00 10}        // push  offset a4       ; "4"

                //$hs = {CB FF 5D C9 AD 3F 5B A1 54 13 FE FB 05 C6 22}  // Initial Handshake ---> can be added or removed for hunting for different variants

                $re1  = {50 00 00 00}
                $re2  = {BB 01 00 00}
                // Dwords of C2 Ports (80 | 443 | 53) 2 -3 times

                $s1 = "pwlfnn10,gzg" // rundll32.exe XOR 02
                $s2 = "AddNum"
                $s3 = "ServiceMain"
                $s4 = "The Window"
                $s5 = "off.dat"

        condition:
                all of ($bin*) and 1 of ($re*) and 3 of ($s*) //The conditions can be adjusted for hunting for different variants
}
```

Sample YARA Rule: Sofacy

```
rule Sofacy_Fybis_ELF_Backdoor_Gen1 {
        meta:
                description = "Detects Sofacy Fysbis Linux Backdoor_Naikon_APT_Sample1"
                author = "Florian Roth"
                reference = "http://researchcenter.paloaltonetworks.com/2016/02/a-look-into-fysbis-sofacys-linux-backdoor/"
                date = "2016-02-13"
                score = 80
                hash1 = "02c7cf55fd5c5809ce2dce56085ba43795f2480423a4256537bfdfda0df85592"
                hash2 = "8bca0031f3b691421cb15f9c6e71ce193355d2d8cf2b190438b6962761d0c6bb"
        strings:
                $x1 = "Your command not writed to pipe" fullword ascii
                $x2 = "Terminal don`t started for executing command" fullword ascii
                $x3 = "Command will have end with \\n" fullword ascii

                $s1 = "WantedBy=mul
                $s2 = "Success
                $s3 = "ls /et
                $s4 = "rm
                $s5 = "ExecStart="
                $s6 = "<table><capt
        condition:
                ( uint16(0) == 0x457f and filesize < 500KB and 1 of ($x*) ) or
                ( 1 of ($x*) and 3 of ($s*) )
}
```

```
condition:
        ( uint16(0) == 0x457f and filesize < 500KB and 1 of ($x*) ) or
        ( 1 of ($x*) and 3 of ($s*) )
```


Sample YARA Rule: Sofacy

Here, we see a more complex condition in a rule for Sofacy written by Florian Roth. In this example, there is a validation that the file is an executable and is smaller than 500Kb, then it is looking for one of the $x variables such as the strings out of the malware, or it's ignoring the file save and PE header and looking for one of the $x variables and three of the $s strings and commands.

```
/*
   This Yara ruleset is under the GNU-GPLv2 license (http://www.gnu.org/licenses/old-licenses/gpl-2.0.html) and open to any user or organization, as long as you use it under this license.

*/
/*
                Yara Rule Set
                Author: Florian Roth
                Date: 2016-02-13
                Identifier: Sofacy Fysbis
*/

rule Sofacy_Fybis_ELF_Backdoor_Gen1 {
                meta:
                                description = "Detects Sofacy Fysbis Linux
Backdoor_Naikon_APT_Sample1"
```

```
author = "Florian Roth"
                reference = "http://researchcenter.paloaltonetworks.com/2016/02/a-look-into-fysbis-sofacys-linux-backdoor/"
                date = "2016-02-13"
                score = 80
                hash1 = "02c7cf55fd5c5809ce2dce56085ba43795f2480423a4256537bfdfda0df85592"
                hash2 = "8bca0031f3b691421cb15f9c6e71ce193355d2d8cf2b190438b6962761d0c6bb"
        strings:
                $x1 = "Your command not writed to pipe" fullword ascii
                $x2 = "Terminal don`t started for executing command" fullword ascii
                $x3 = "Command will have end with \\n" fullword ascii

                $s1 = "WantedBy=multi-user.target' >> /usr/lib/systemd/system/" fullword ascii
                $s2 = "Success execute command or long for waiting executing your command" fullword ascii
                $s3 = "ls /etc | egrep -e\"fedora*|debian*|gentoo*|mandriva*|mandrake*|meego*|redhat*|lsb-*|sun-*|SUSE*|release\"" fullword ascii
                $s4 = "rm -f /usr/lib/systemd/system/" fullword ascii
                $s5 = "ExecStart=" fullword ascii
                $s6 = "<table><caption><font size=4 color=red>TABLE EXECUTE FILES</font></caption>" fullword ascii
        condition:
                ( uint16(0) == 0x457f and filesize < 500KB and 1 of ($x*) ) or
                ( 1 of ($x*) and 3 of ($s*) )
}

rule Sofacy_Fysbis_ELF_Backdoor_Gen2 {
        meta:
                description = "Detects Sofacy Fysbis Linux Backdoor"
                author = "Florian Roth"
                reference = "http://researchcenter.paloaltonetworks.com/2016/02/a-look-into-fysbis-sofacys-linux-backdoor/"
                date = "2016-02-13"
                score = 80
                hash1 = "02c7cf55fd5c5809ce2dce56085ba43795f2480423a4256537bfdfda0df85592"
                hash2 = "8bca0031f3b691421cb15f9c6e71ce193355d2d8cf2b190438b6962761d0c6bb"
                hash3 =
```

```
"fd8b2ea9a2e8a67e4cb3904b49c789d57ed9b1ce5bebfe54fe3d98214d6a0f61"
        strings:
                $s1 = "RemoteShell" ascii
                $s2 = "basic_string::_M_replace_dispatch" fullword ascii
                $s3 = "HttpChannel" ascii
        condition:
                uint16(0) == 0x457f and filesize < 500KB and all of them
}
```

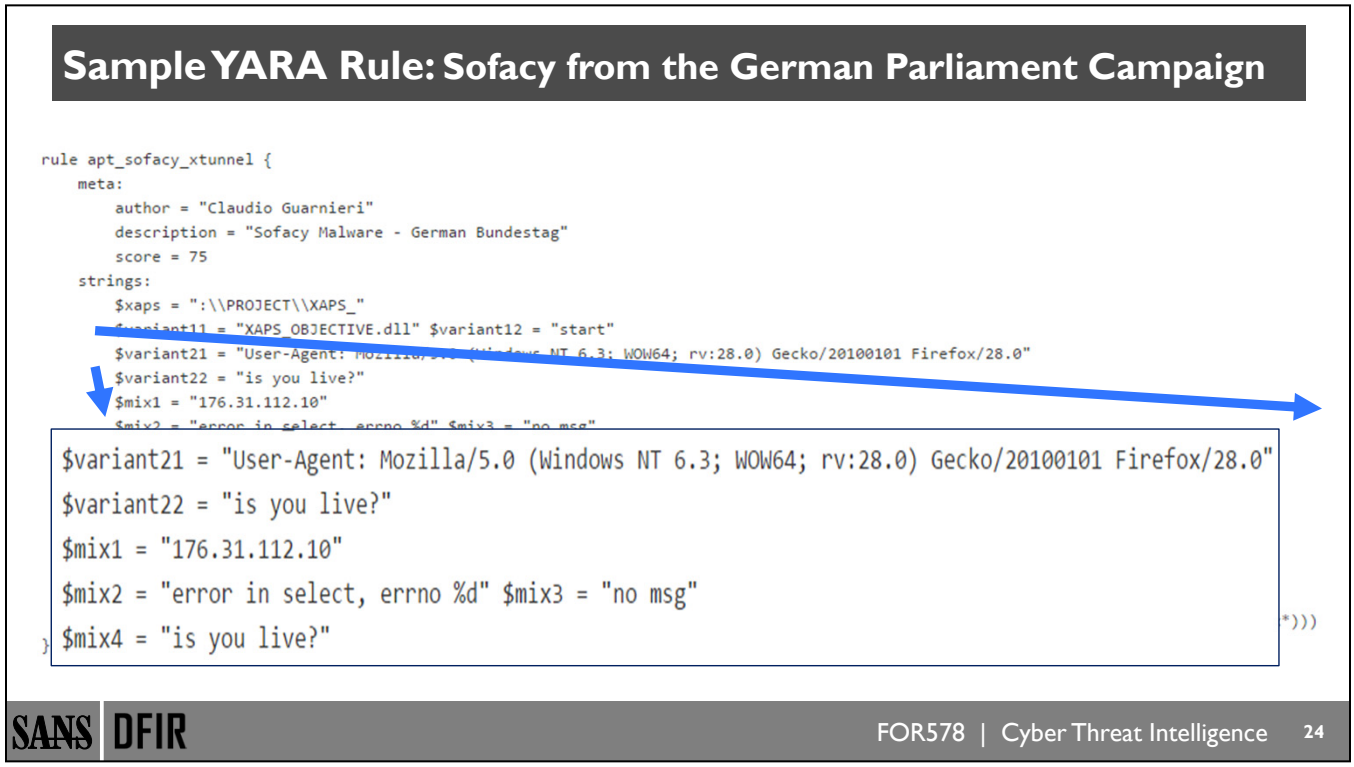

Sample YARA Rule: Sofacy from the German Parliament Campaign

Here, we have a YARA rule written by Claudio Guarnieri that identifies the specific variant of malware Sofacy was using when it targeted the German Government (or Bundestag). Notice that the signature is looking for specific user-agents and strings in the malware, such as "is you live?" Identifying things like broken English or other languages as well as misspelled words and specific structuring of commands can be a great way to identify a piece of malware; adding in specific user-agents and strings around that variant can further help to eliminate false positives.

```
/*
   This Yara ruleset is under the GNU-GPLv2 license (http://www.gnu.org/licenses/old-licenses/gpl-2.0.html) and open to any user or organization, as long as you use it under this license.

*/
rule apt_sofacy_xtunnel {
   meta:
      author = "Claudio Guarnieri"
      description = "Sofacy Malware - German Bundestag"
      score = 75
   strings:
      $xaps = ":\\PROJECT\\XAPS_"
      $variant11 = "XAPS_OBJECTIVE.dll" $variant12 = "start"
      $variant21 = "User-Agent: Mozilla/5.0 (Windows NT 6.3; WOW64; rv:28.0) Gecko/20100101 Firefox/28.0"
```

```
$variant22 = "is you live?"
      $mix1 = "176.31.112.10"
      $mix2 = "error in select, errno %d" $mix3 = "no msg"
      $mix4 = "is you live?"
      $mix5 = "127.0.0.1"
      $mix6 = "err %d"
      $mix7 = "i`m wait"
      $mix8 = "hello"
      $mix9 = "OpenSSL 1.0.1e 11 Feb 2013" $mix10 = "Xtunnel.exe"
   condition:
      ((uint16(0) == 0x5A4D) or (uint16(0) == 0xCFD0)) and (($xaps) or (all of ($variant1*)) or (all of
($variant2*)) or (6 of ($mix*)))
}

rule Sofacy_Bundestag_Winexe {
   meta:
      description = "Winexe tool used by Sofacy group in Bundestag APT"
      author = "Florian Roth"
      reference = "http://dokumente.linksfraktion.de/inhalt/report-orig.pdf"
      date = "2015-06-19"
      hash = "5130f600cd9a9cdc82d4bad938b20cbd2f699aadb76e7f3f1a93602330d9997d"
      score = 70
   strings:
      $s1 = "\\\\.\\pipe\\ahexec" fullword ascii
      $s2 = "implevel" fullword ascii
   condition:
      uint16(0) == 0x5a4d and filesize < 115KB and all of them
}

rule Sofacy_Bundestag_Mal2 {
   meta:
      description = "Sofacy Group Malware Sample 2"
      author = "Florian Roth"
      reference = "http://dokumente.linksfraktion.de/inhalt/report-orig.pdf"
      date = "2015-06-19"
      hash = "566ab945f61be016bfd9e83cc1b64f783b9b8deb891e6d504d3442bc8281b092"
      score = 70
   strings:
      $x1 = "PROJECT\\XAPS_OBJECTIVE_DLL\\" ascii
      $x2 = "XAPS_OBJECTIVE.dll" fullword ascii
```

```
        $s1 = "i`m wait" fullword ascii
    condition:
        uint16(0) == 0x5a4d and ( 1 of ($x*) ) and $s1
}
```

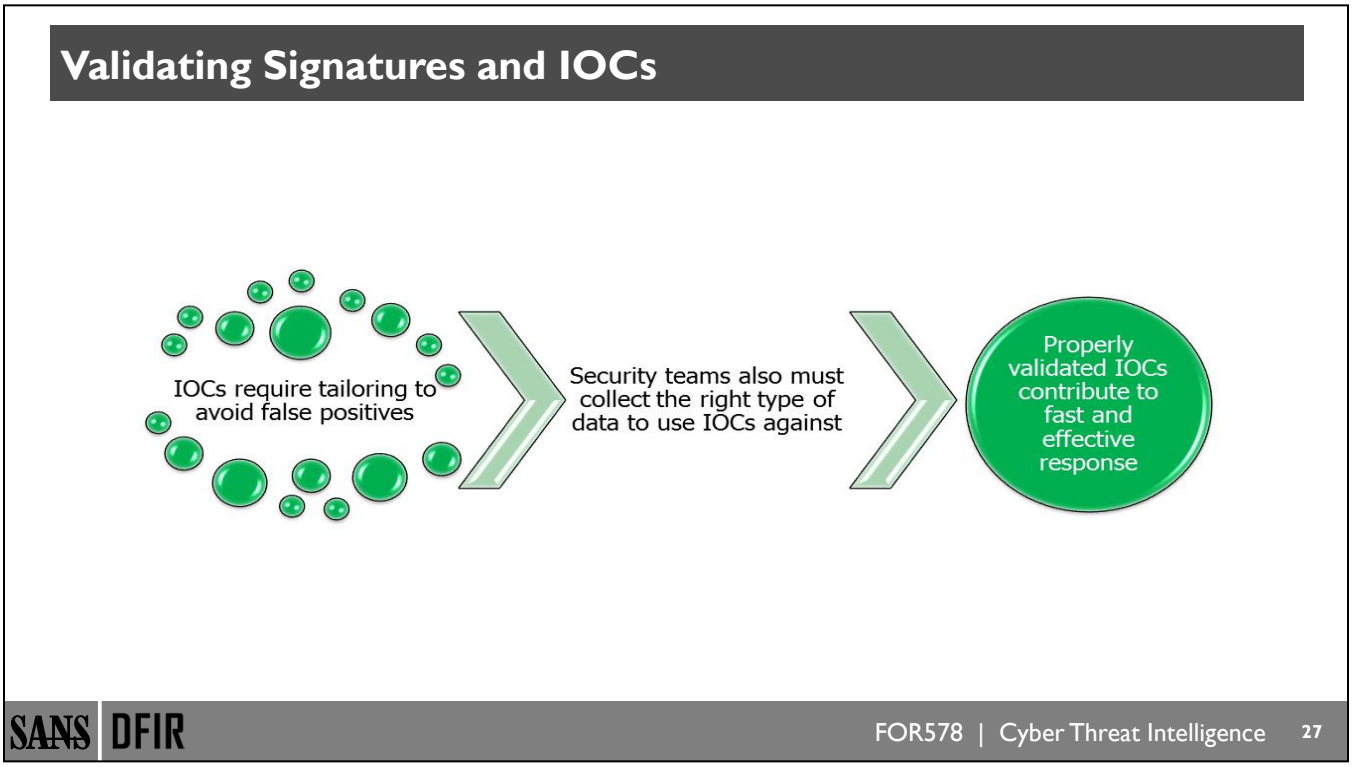

Validating Signatures and IOCs

Signatures and IOCs such as YARA need a lot of tailoring to avoid false positives. A general rule of thumb is to tailor IOCs to eliminate as many false positives as possible prior to uncovering a threat. Once a threat is identified, accept more false positives in an effort to open up the aperture of a rule to find variants of the malware or threat in the environment. False positives prior to finding a threat are one of the costliest aspects of using IOCs and can quickly discourage the security process. Validate IOCs by testing them against digital images in the environment.

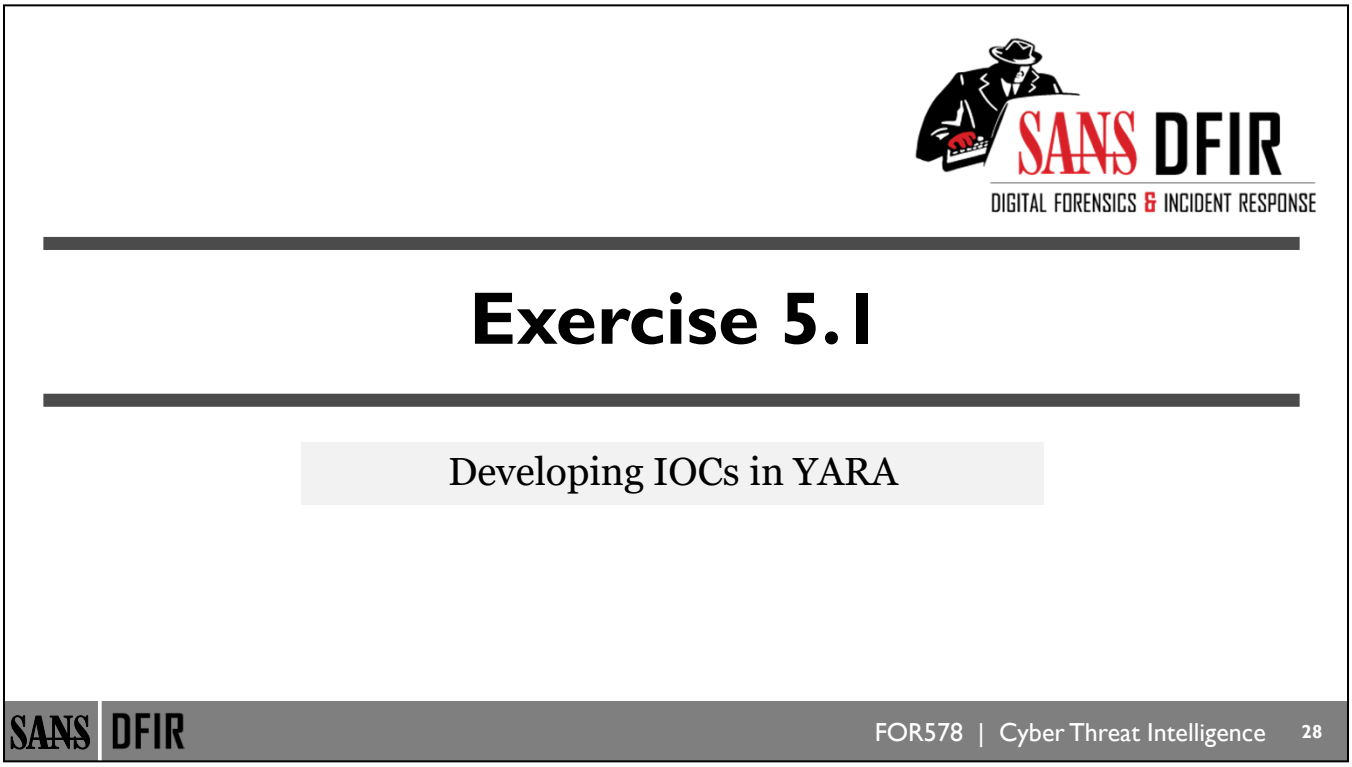

Please refer to your Lab Workbook and complete Exercise 5.1.

Case Study: HackingTeam

This page intentionally left blank.

Case Study: HackingTeam (1)

- Italian security firm HackingTeam specialized in providing surveillance and exploitation services for governments and law enforcement
- On July 5, 2016, over 400GB of data was stolen from the company's servers and posted online

Case Study: HackingTeam (1)

The Italian security firm HackingTeam was known for providing exploits and network access for governments and law enforcement. There were many who questioned the legality and ethics of some of HackingTeam's activities; they were known for providing their services to many countries with questionable human rights records. One of these individuals who went by the name Phineas Fisher hacked into the networks of the security company, stole over 400GB of sensitive information, and released it online. This information included sensitive emails that revealed customers and the source code for HackingTeam's primary tool, their Remote Control Server (RCS).

Reference:

- https://nakedsecurity.sophos.com/2016/04/19/how-hacking-team-got-hacked/

Case Study: HackingTeam (2)

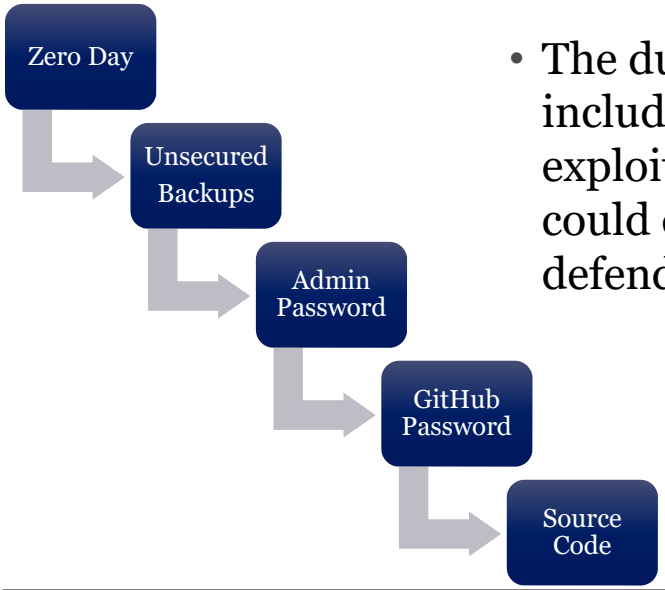

- The dumped information included malware samples, exploits, and target lists that could quickly be leveraged by defenders
  - It was clear that HackingTeam's operations and capabilities have been used by various entities around the world

Case Study: HackingTeam (2)

The hacker known as Phineas Fisher outlines how he was able to hack into the network. The initial access point was through a zero-day exploit that the hacker developed in an embedded system. He claimed that developing the exploit only took two weeks. Once he gained access, he was able to conduct extensive reconnaissance and was able to identify several unencrypted backup servers. One of those servers, an email server, revealed massive amounts of sensitive information on the hacking team, as well as admin passwords that the hacker was able to use to get domain admin, move through the network, and eventually discover the password to the HackingTeam's GitHub repository. Once he had access to the GitHub repository, he was able to pull the code for all of the HackingTeam's sensitive tools.

References:

- https://www.ibtimes.co.uk/hacking-team-hacked-10-things-learned-massive-data-breach-spying-company-1509925
- https://nakedsecurity.sophos.com/2016/04/19/how-hacking-team-got-hacked/

HackingTeam Isn't Alone

- NSO Group is an Israeli-based private company that sells spyware and "lawful intercept" capabilities to "law enforcement" around the world
- Citizen Lab has consistently tracked them and other similar teams whose capabilities and operations end up targeting journalists, human rights advocates, and political dissidents

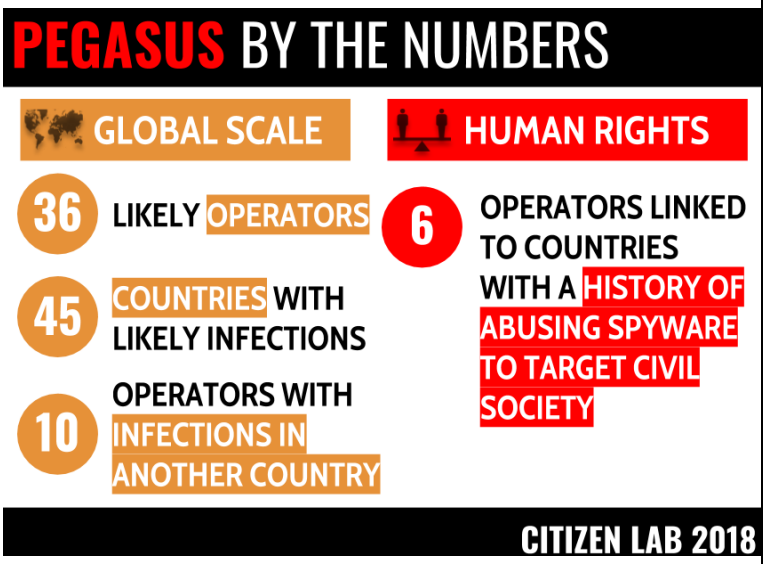

HackingTeam Isn't Alone

HackingTeam isn't the only mercenary group in the world. In fact, there are dozens of such companies publicly known. Each tries to differentiate themselves from the others by stating that they are somehow the moral ones and only sell to "the good guys" which have a varying definition in each part of the world.

The reality is even if the teams really only sell to "law enforcement" and only sell "lawful intercept" capabilities that every country has different laws and views; you start selling to everyone and you quickly realize the overlap makes it where anything and any target is legal.

What is important for our purpose, though, is understanding the impact of this to our understanding of adversaries and their capabilities.

Reference:

- https://citizenlab.ca/2018/09/hide-and-seek-tracking-nso-groups-pegasus-spyware-to-operations-in-45-countries/

HackingTeam's Compromise and Mercenary Group Takeaways

Government agencies worked alongside several security firms, primarily Rook Security, to identify and release IOCs related to HackingTeam's malware and tactics. Rook Security also released a tool to check specifically for HackingTeam activity. The release of this information enabled defenders to not only check to see whether they had been targeted by either HackingTeam or copycat attacks but also to protect themselves by patching the vulnerabilities that were exploited by HackingTeam's tools.

The real takeaway, of course, is that these types of mercenaries make the world a more unsafe place for everyone. Their tools and capabilities get out and others use them. So, on the intel team, it is important for you to be aware of such teams and especially aware of any leaks of their capabilities. You don't have to have a standing intelligence requirement on each of these teams but some general familiarity and ability to focus on them if needed is a smart investment. Further, realize that when we talk about the "adversary", they are very rarely one entity or team.

Reference:

- https://www.securityweek.com/rook-security-unveils-hacking-team-breach-detection-tool

Dissemination: Operational

This page intentionally left blank.

Operational Threat Intelligence

- Operational threat intelligence is the focus for operational level audience members:
  - Those members that serve as the bridge between the strategic and tactical personnel
  - Understand the technical, but look at the bigger picture
- Operational threat intelligence should:
  - Help identify knowledge gaps and foster partner sharing to minimize these gaps
  - Document and understand the evolution of adversary campaigns and threat changes
  - Help structure security teams to match size, training, and subject matter expertise to counter the appropriate threats

Operational Threat Intelligence

Those operating at the operational level of threat intelligence serve a key role in that they are the bridges between the strategic and the tactical. They must understand tactical level threat intelligence to relay it to the strategic decision-makers appropriately, and they must understand the strategic decision-makers' needs and language to translate requirements to the tactical level personnel. They identify knowledge gaps, structure the teams, identify training and requirements for personnel, and identify and initiate threat intelligence sharing between partners and peers.

Communicating About Adversary Operations

- Once we understand Activity Groups, we can design defensive strategies around them as if they are a playbook for the defenders
- As an example, you can create a MITRE ATT&CK navigator for Enterprise
  - For organizations that don't use Enterprise (e.g., ICS) you can create your own such as the one below that lets people click on an Activity Group and maps out the techniques seen
- Defenders can use this for tabletop exercises, red teaming, prevention and detection

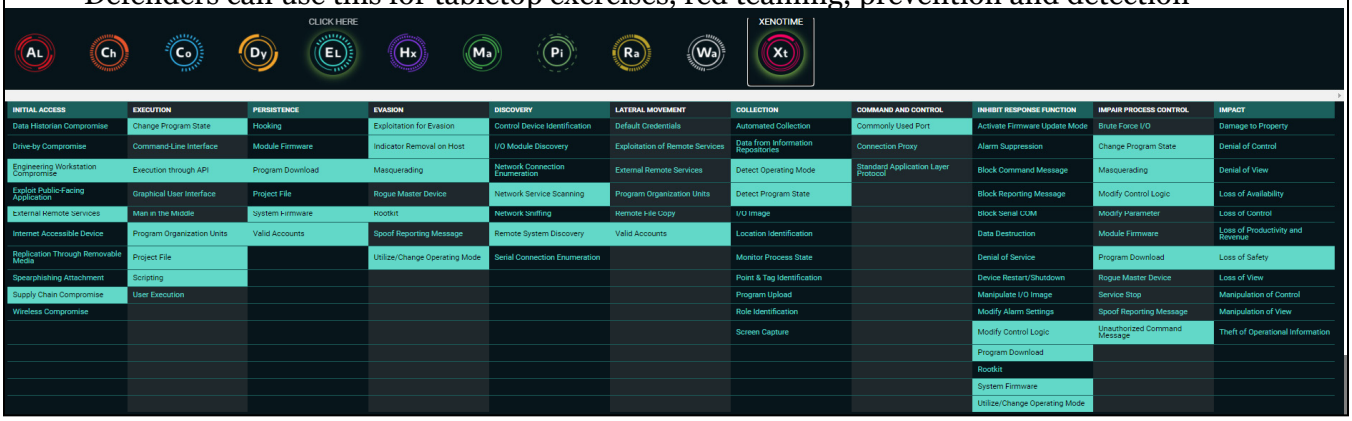

Communicating About Adversary Operations

At the operational level of threat intel, it's incredibly important to be able to communicate about what defenders need; in this case, we're essentially articulating a defensive playbook by selecting an Activity Group (in this case XENOTIME who was responsible for the first ever cyber attack that specifically targeted human life) and the website automatically populates the observed techniques that XENOTIME has been seen using. This helps defenders think through what they will need in prevention, detection, and response-based controls to respond to this style of attack, regardless if it's orchestrated by XENOTIME or not. This can also be leveraged for incident response tabletop exercises, red team and purple teaming, and more.

Reference:

- https://www.dragos.com/mitre-attack-for-ics/

Partners and Collaboration

- The best producers of threat intelligence have great access to collecting data from outside their own networks and sources
- Partnership and collaborations facilitate the best in threat information sharing
- Key sources to consider:
  - Government-private sharing
  - Groups and email distributions
- Collection can also be done without partnering through Open-Source Intelligence gathering

Partners and Collaboration

Establishing partners and collaborating is much easier said than done, but it is vital to doing threat intelligence properly. As discussed previously, there are biases and knowledge gaps each of us face; having partners and collaborating with others (sometimes even competitors) can help overcome these. The big focus for the rest of the section is OSINT collection, because it is more demonstrable in class than how to form a partnership with an organization, but each is equally important. Partnerships generally involve meeting someone, starting the conversation, and getting the proper NDAs in place to start building that relationship. A good informal collaboration opportunity usually exists in the form of malware analysis and threat analysis email distributions and private email groups. To join these groups, you normally have to be sponsored by someone; ask those you work with if they are involved in any, or discuss with your fellow SANS students if you are taking this class in person and attempt to find some of these groups.

National-Level Government Information

- Derived from
  - Criminal investigations
  - Public/private partnerships
  - Foreign intelligence
- Dissemination points include
  - US-CERT (DHS)
  - InfraGard (FBI)
  - NCSC (UK)
  - ACSC (Australia)
  - CSA (Singapore)
  - Etc.

National-Level Government Information

Government information is usually obtained from criminal investigations, public and private partnerships, and conducting intelligence on foreign targets under U.S. Title 50 authorities. Two prime locations in the United States for personnel to get government information is through the U.S. Computer Emergency Response Team (CERT) and through the FBI's InfraGard. Both allow organizations to join and get access to sensitive reports and adversaries about current threats and trends. Criticism of these organizations usually occurs when organizations expect to get all their information from these sources. Instead, these government organizations should be relied upon only for additional sources of information. Private partnerships and internal data collection are needed to have a robust threat intelligence program.

References:

- https://www.infragard.org/
- https://us-cert.cisa.gov/

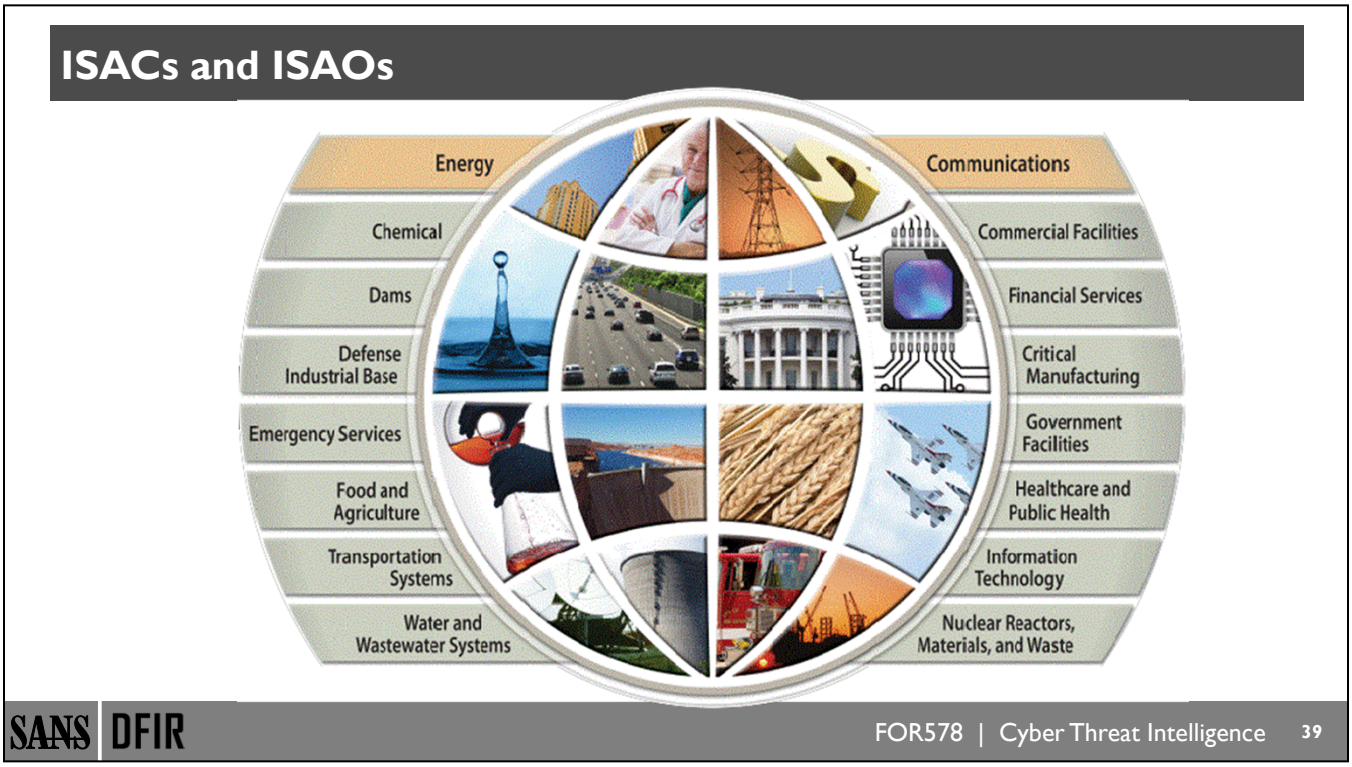

ISACs and ISAOs

Information Sharing and Analysis Centers (ISACs) were an endeavor started in 1998 with Presidential Decision Directive 63 to share threat information among the civilian and government sector. Many industries have ISACs when they are deemed to be Critical Infrastructure. The ISACs are a good way to share information with others in your industry while receiving information about threats out there specific to your sector. Each ISAC has a website that you can access and learn more about them. Operational threat intelligence decision-makers should seek out the ISAC related to their industry (if there is one). If there is not one related to the business operations, it is now possible to seek out an ISAO.

President Obama issued Executive Order 13691 in 2015, which established nongovernmental organizations identified as Information Sharing and Analysis Organizations (ISAOs). ISAOs are an initiative through DHS to encourage private community sharing as well as sharing between private and government sectors. They are organizations that are allowed to form and achieve designation that protects their data while being allowed to share it with the government and receive information from the government that can be useful. ISAOs expand the concept of the ISACs and allow ISAOs to have their information treated as Protected Critical Infrastructure Information, which protects the information from disclosure including through Freedom of Information Act or Sunshine laws. It also means that the information is exempt from regulatory and civil litigation.

However, there are concerns. Analysts should be cognizant of how organizations develop reputations for how they gather and share information. It is advisable to watch ISAOs carefully and base involvement on reputation and the match to the type of threats your organization is looking to learn about.

References:

- https://www.nationalisacs.org/
- https://obamawhitehouse.archives.gov/the-press-office/2015/02/13/executive-order-promoting-private-sector-cybersecurity-information-shari

Additional Resources

- Carnegie Mellon CERT: List of National CERTs
- Cyber Threat Alliance
- International Information Sharing Groups
  - FIRST
- International Associations

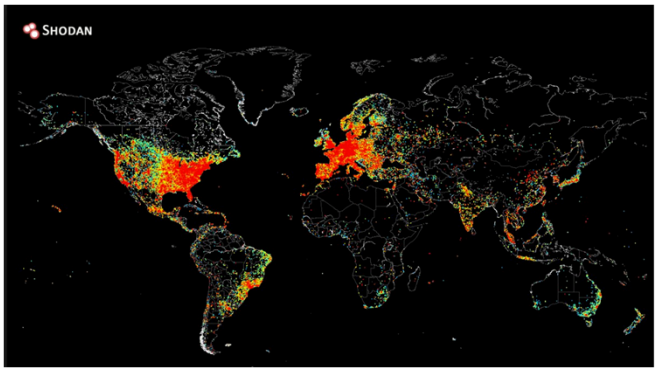

Additional Resources

If you have questions about sharing organizations or other resources in countries that we did not mention, CMU has a list of CIRT contacts around the world that they keep updated. In addition, there are various international information-sharing associations and groups that aim to share information with vetted individuals or organizations around the world.

References:

- https://www.sei.cmu.edu/our-work/cybersecurity-center-development/index.cfm
- https://cyberthreatalliance.org/
- https://www.first.org/

Image:

- Shodan.io

STIX/TAXII

- STIX: Structured Threat Information eXpression
  - Describes threat information
  - CybOX: Cyber Observable eXpression
    - Describes observables (IOCs)
    - Previously distinct, in STIX2, combined with STIX
- TAXII: Trusted Automated eXchange of Indicator Information
  - Transport mechanism for STIX

and open-source solutions (such as Soltra) available to the community. It is important to understand that most organizations wanting to share threat intelligence need to plan for using TAXII. It is entirely okay and often encouraged to create your own standards internal to your organization that works for your people, processes, and tools. However, this should be extensible to TAXII so that conversion can be quick and automated when sharing with other organizations.

References:

- https://oasis-open.github.io/cti-documentation/taxii/intro.html
- https://github.com/TAXIIProject
- http://stixproject.github.io/getting-started/whitepaper/

TAXII Services

The TAXII services are options available in TAXII. Note that this standard is more akin to an RFC for using TAXII. This open ability to use and integrate TAXII in ways that you prefer is one of the reasons organizations were okay with TAXII and one of the reasons it's so difficult to implement. There is commonality in the services but often not much commonality in the specific implementation.

Reference:

- http://taxiiproject.github.io/releases/1.1/TAXII_Services_Specification.pdf

STIX

STIX is a related effort to TAXII. Where TAXII is the specification on how to send threat data, STIX is the standard for how to structure that threat data. STIX1 is XML-based language, and STIX2 is JSON-based. STIX is a combination of IoCs, context, recognized TTPs, any relevant actor attribution, and other information such as suggested courses of action for when STIX data matches an observed threat. The issue is that there is not much definition for how to structure the data other than what is discussed in the standard. That is, the type of data that is fit into each category in STIX is widely varying. STIX does attempt to give options to an analyst; for example, you could define a new piece of malware as a TTP and then see what other indicators out there use the same TTP and begin to uncover and identify campaigns. It's a manual process and is similar to Python scripting, but it needs work to be more widely adopted.

Related efforts include MAEC, CAPEC, and CybOX; each are MITRE-run languages for documenting observable indicators and information about malware, campaigns, and cyber activity.

References:

- https://github.com/STIXProject
- http://stixproject.github.io/documentation/idioms/campaign-v-actors/
- http://stixproject.github.io/data-model/1.1.1/
- http://maecproject.github.io/about-maec/

TAXII Implementations

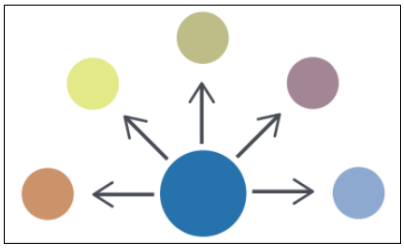

- Source/Subscriber
  - Source determines what is available to the Subscribers
  - Subscribers can Pull data but not Push (can request info but not submit)

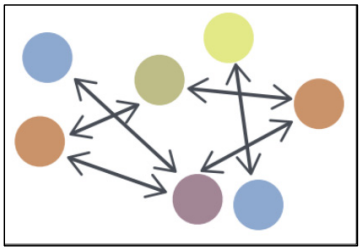

- Peer to Peer
  - Multiple organizations can produce data and multiple organizations can consume data

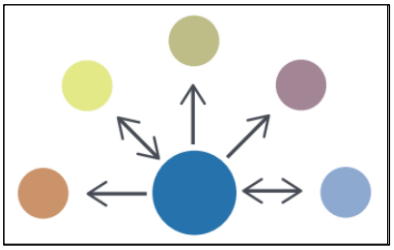

- Hub and Spoke
  - Hub acts as the clearinghouse for all information
  - Subscribers can Pull data and Push data

TAXII Implementations

TAXII can be set up in one of three formats. First, the Source and Subscriber model means that there is one Source that various Subscribers can pull data from. They cannot submit threat information to this Source but can freely pull information from the Source at any time. The Peer-to-Peer model means that Subscribers can act as Sources and Subscribers to multiple organizations. In this model, some organizations will be only a Source, some will be only Subscribers, and some will act as both. In the Hub and Spoke setup, there will be a "clearinghouse" of sorts—such as an ISAC or the US-CERT—which can receive data but is the only Source that data can come from. That is, everyone submits their data to the Source, who then validates the data and pushes it out in the original form or a new form to the Subscribers.

References:

- http://taxiiproject.github.io/community/
- http://taxii.mitre.org/about/documents/Introduction_to_TAXII_White_Paper_May_2014.pdf

STIX 2.1 Objects

STIX 2.*

- In response to critiques of STIX 1.0, STIX 2.0 and 2.1 was put under OASIS
  - OASIS is a technical committee taking input and governing STIX 2.1
- STIX 2.1 attempted to be more flexible and simplistic as well and leveraging a graph-based model approach
- Changes from STIX 1.0 include:
  - One standard (CybOX now in STIX)
  - JSON instead of XML
  - Indicator pattern language with KC phases

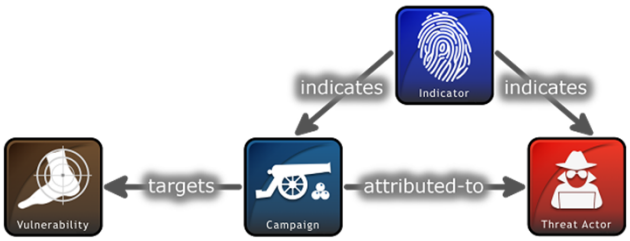

STIX 2.*

STIX 2.0, and subsequently 2.1, were made in response to criticisms of STIX 1.0, and the team behind it has been pretty proactive in improving the standard. Some considerable changes include a graph-based approach and a structured language that supported classification by kill chain phases. Additionally, STIX 2.* uses JSON instead of XML and no longer requires you to use CybOx and STIX but just STIX.

If your organization is still using STIX 1.0, it should be an easy transition to STIX 2.* if your organization or sharing group is looking to stand up a STIX server. It is our recommendation to use STIX 2.* and keep up with the latest version from OASIS.

Reference:

- https://oasis-open.github.io/cti-documentation/

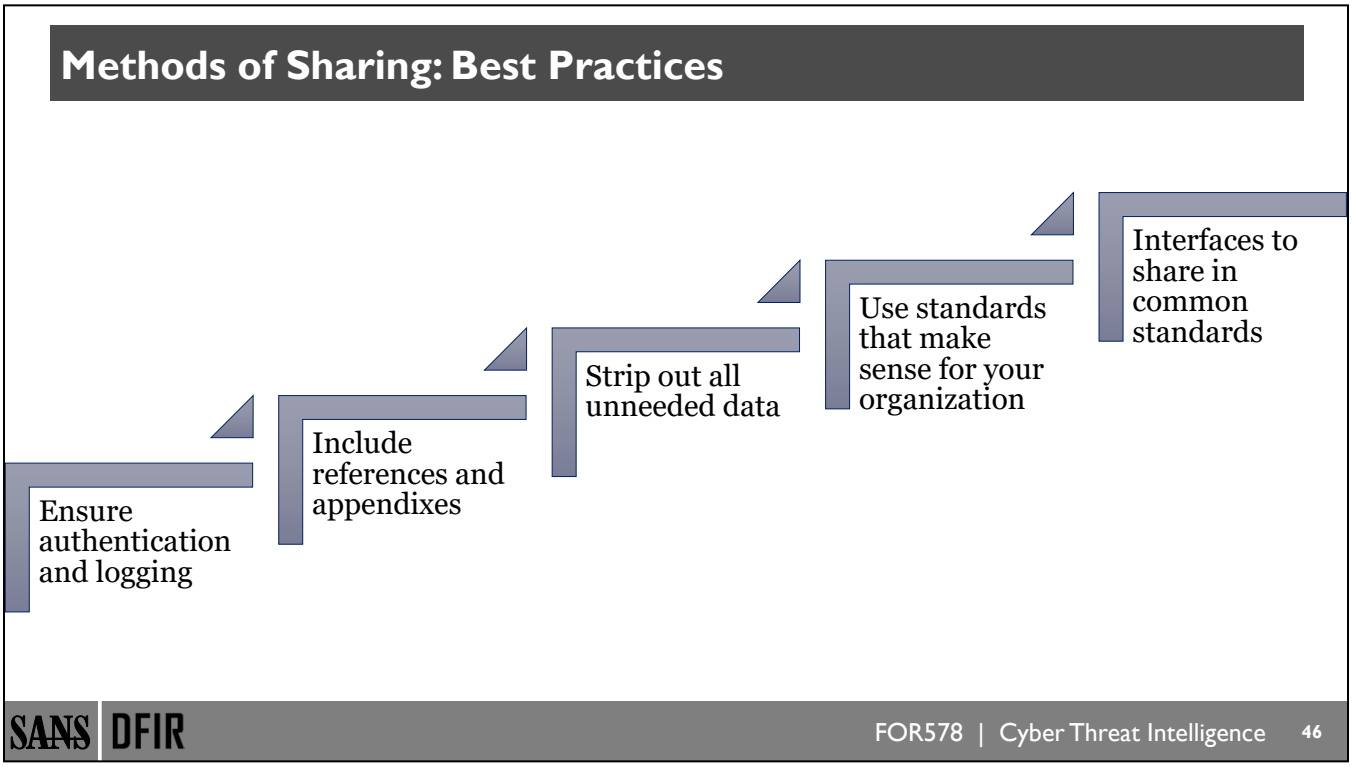

Methods of Sharing: Best Practices

When sharing your information on or off your team, make sure you have validation and authentication for the personnel, even those external to your organization. In addition, threat intelligence reports should have linkable appendixes that have IOCs that can be stored outside the reports but easily accessed. You need to store the IOCs out of the reports so that you can apply analytics to your databases to search for common trends and links between IOCs. Also, have them in such a form so that you can share the IOCs more easily when you might not want to share the actual threat intelligence reports, which may contain sensitive data. Be sure to strip all data you do not need as well. For example, if you have process information for the systems that you're gathering information from, but it's not used for threat intelligence, do not store that data in your threat intelligence database or report. You want to separate data appropriately. (This can also be a regulation issue in various industries; always make sure you are in compliance with your actions.) Lastly, try to use recognized standards so that your data is useful to others and so you can learn from well-established processes.

- Ensure authentication:
  - Even for internal users, ensure a form of authentication such as a "minimum" of unique user/pass.
- Threat intelligence reports should have references or an appendix for the IoCs on the observed threat:
  - Store the IoCs outside of the report (linked in the report) and establish API access to the IoCs for use in various tools.
- Strip all unneeded data:
  - PII, unneeded process data, and so on should be removed; if it does not support understanding the threat or building defenses, remove it.
- Use recognized standards:
  - Do not repeat the process; this ensures that you can share the data easily internally but also that you can share the data externally in various formats and with authorized entities.

Exercise 5.2 Introduction

- This exercise will introduce how to work with STIX and TAXII by using a script to pull down data from the MITRE ATT&CK TAXII server
- The exercise will then go over producing a visualization to help map data sources to techniques based on the data pulled back from the TAXII server

Exercise 5.2 Introduction

The focus of this lab is to gain exposure to STIX and TAXII as well as analyze a MITRE ATT&CK dataset to map data sources to techniques.

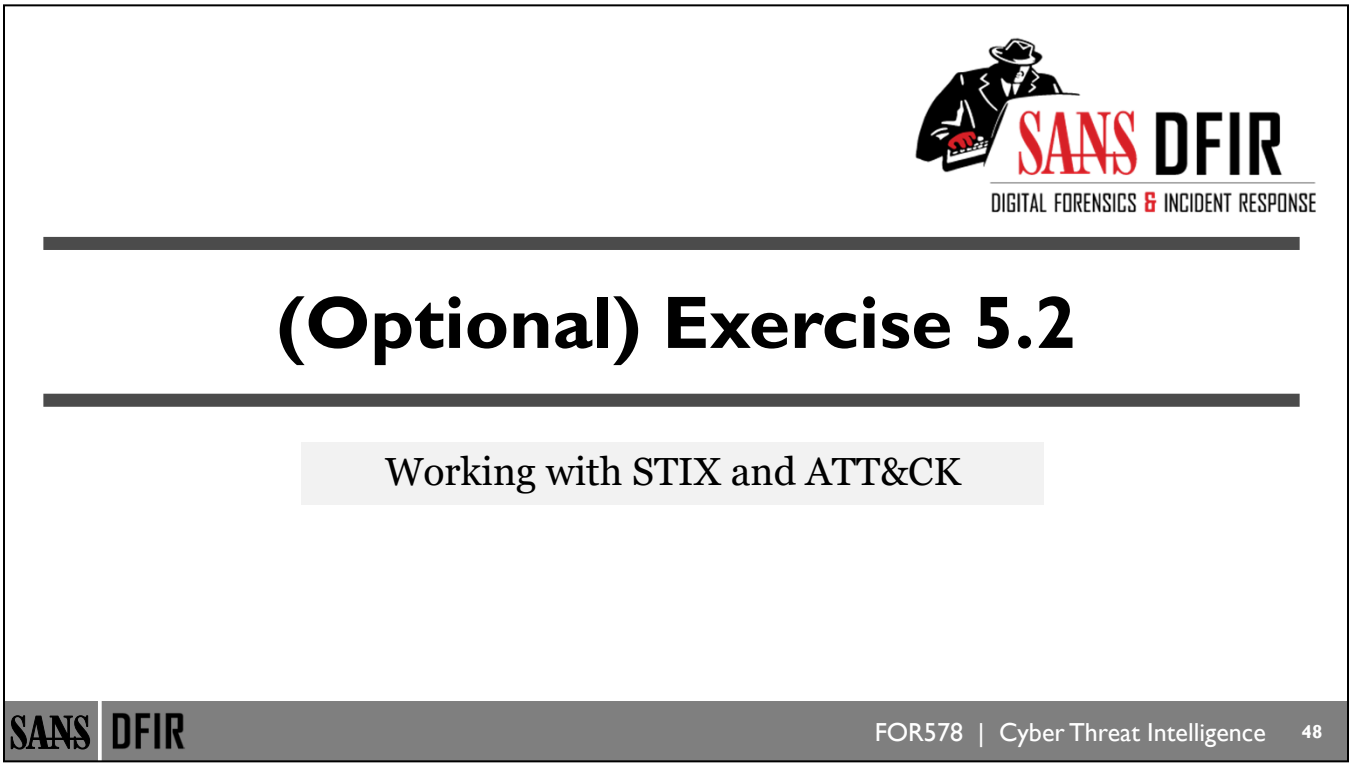

Depending on the timing in the class, this lab may be moved to an optional lab for completion after class if you want.

Also, thank you to Katie Nickels for this lab; the original idea I had for the lab was a lot worse until Katie updated it. You rock, Katie.

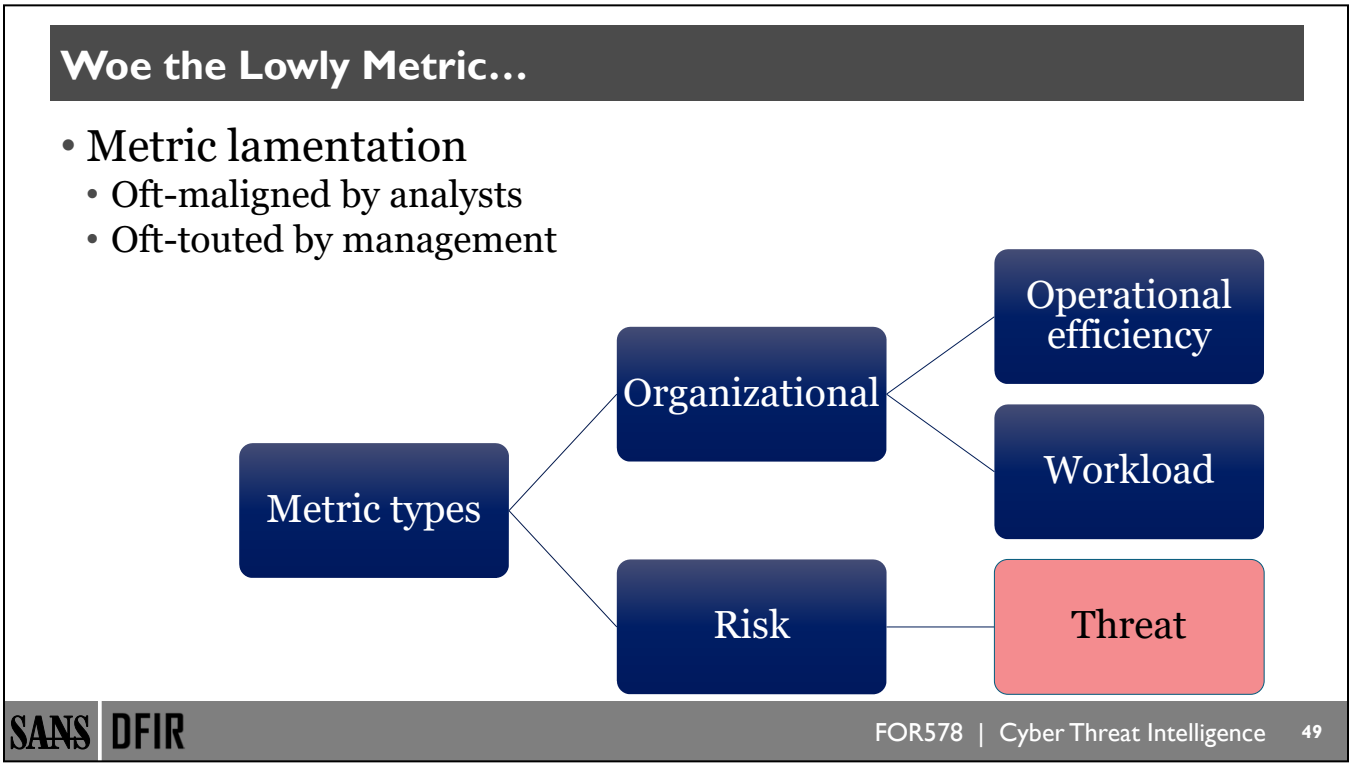

Woe the Lowly Metric…

Metrics are often maligned by analysts as worthless and a time-sink, while simultaneously being touted by executives and policymakers.

They can be divided by organizational metrics and risk metrics.

- Organizational metrics may be subdivided into operational efficiency and workload metrics.
- Risk may be subdivided into a number of other metrics; the one that we will focus on for CTI is threat-oriented metrics.

There are a number of issues that underlie problematic metrics. They include:

- Inconsistent terminology, such as the use of the term "attack" in metrics ("we were attacked X times")
- The significance or weight of the metrics is unclear or ambiguous
- The measures themselves are incorrectly interpreted as a quantification of some situation—such as network sweeps are not attacks
- The tendency to try to quantify everything, known in the field of economics as "Physics Envy." This includes attempting to assign numerical values to non-numerical criteria; such as, if the adversary is capable of A, B, and C, we say their sophistication level is "1."
- The measures themselves are subject to interpretation (nondeterministic), or the assignment of criteria to numerical values is subjective.
- The metrics do not map to nor suggest follow-up actions that will influence future measurements of the same metric.
- The metrics are defined by management or non-SME analysts.

Why You Should Embrace Metrics

"Metrics are useless"
"… then you just don't get it."

- Opportunity for clear, concise communication of message
- Can be interpreted by wide audience
- Represents large amount of data
- Visual data representation naturally more compelling
- Target metrics useful by management AND analysts

Why You Should Embrace Metrics

Many analysts eschew metrics as useless. The reality is that the utility of metrics is the degree to which they represent meaningful data and inform a subsequent action. Metrics distill a large amount of information down into a clear, concise message and provide an opportunity to create a message that is digestible by a broad audience through the use of visual techniques that the human brain is physiologically predisposed to interpret. These representations can make a message more compelling.

If you are being asked to provide metrics to management that you feel aren't useful, develop metrics that are at least useful to analysts. Hopefully, in time, your management will also see the value of those metrics and embrace them. At worst, you and your colleagues will be promoted and eventually ask for them as management yourselves. ☺

In this section, we will discuss some metrics that experienced CTI analysts have found useful over the years.

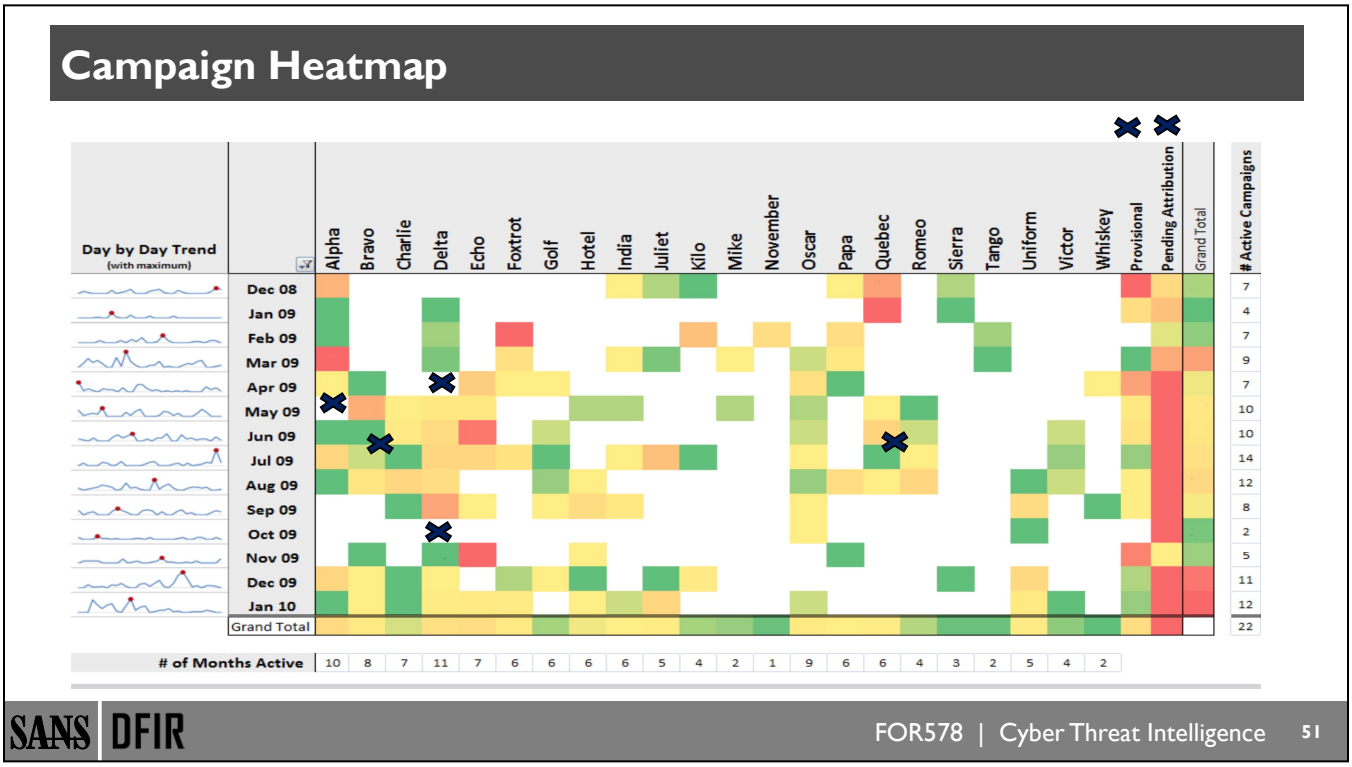

Campaign Heatmap

Vertical columns represent different campaigns tracked by a CTI organization. The rows represent distinct intrusion attempts in a single month, with the color indicating relative activity level (red being high, green being low). The far-right column indicates the total number of distinct campaigns executing operations in a month, while the bottom-most row represents the number of months shown in which each given campaign is active. You could do this same analysis for any clustering of your choice, whether it is activity groups, threat groups, or anything else.

"Provisional" refers to intrusions that correlate to other intrusions, but no currently-named campaign. "Pending Attribution" represents the number of intrusion attempts bearing characteristics of APT activity that do not correlate to other intrusion attempts. These columns represent intelligence gaps or opportunities for greater definition of campaigns.

X's mark potential areas for exploration, such as:

- Unattributed intrusions ("Pending"/"Provisional")
- Campaigns that tend to be active every month that are observed to be inactive in one month (suggesting perhaps one of the "Pending Attribution" intrusions in the same month might actually attribute to that campaign)
- Multiple campaigns that seem to operate in the same months as one another with a high level of consistency

Credit: Lockheed Martin CIRT

Organizational Heat Maps

- Use tools such as the MITRE ATT&CK Navigator to articulate where you have mitigations and detections
- E.g., mitigation + detection = green

Alternatively:
Lighter color = 1 detection
Mid color = 2 detections
Darker color = 3+ detections

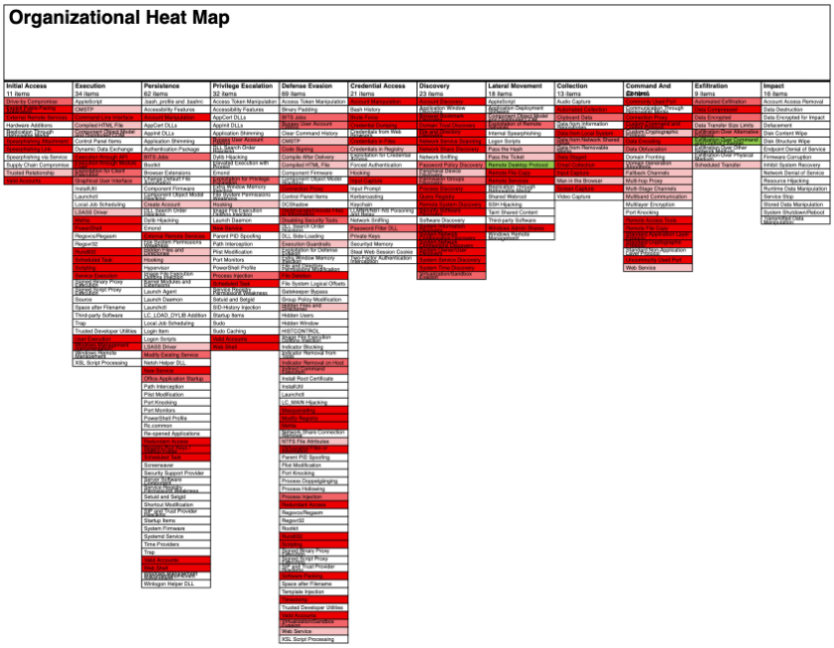

SANS DFIR

Organizational Heat Maps

The concept of using heat maps to articulate defender strategies isn't new, but MITRE ATT&CK's model being a commonly accepted viewpoint has allowed broader communication of this mechanism. In this image from FireEye, they note that they use the different colors to note when they have strong vs. weak mitigation or detections against various techniques and in context of different adversaries.

You could do this per adversary or for your overall team.

Further, you do not have to combine mitigations and detections; in most cases, it makes sense to separate them. Instead, we'd suggest using a light color or different color altogether such as Red for areas you have 1 solid behavioral detection (remember, when you're detecting tactics and techniques, you're using things like threat behaviors not necessarily indicators), Yellow/Orange for 2 detections, and Green for 3+ detections.

Then as your team matures and you start filling the heat map reset it where Red = 2, Yellow/Orange = 3, and Green = 4+. Constantly move the bar forward and you'll see the progression of your team's maturity and defensive capabilities over the years.

The point is not to play bingo with MITRE ATT&CK though; this should be a mechanism to communicate about your broader strategy, not just throwing things to change colors on the "scorecard."

References:

https://www.fireeye.com/blog/products-and-services/2020/01/operationalizing-cti-hunt-for-defend-against-iranian-cyber-threats.html

https://mitre-attack.github.io/attack-navigator/

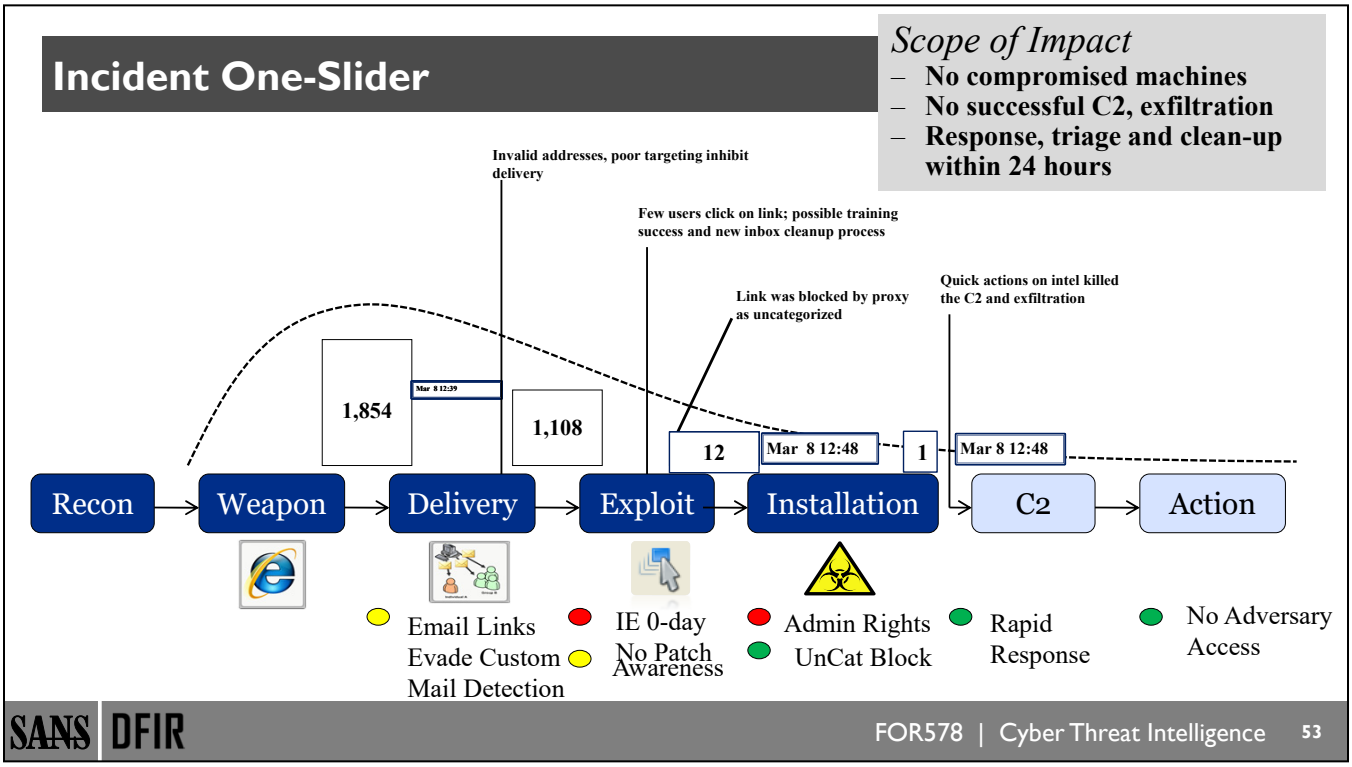

Incident One-Slider

This is an example of how the kill chain can be used to describe the progress and success of an intrusion from an adversary's perspective, as well as the applicable mitigations and their cumulative effect on the kill chain.

Here, we see a phishing intrusion which ended up being unsuccessful that targeted multiple users. The vertical call-out text describes the infrastructure design, adversary mistake, or tactical mitigation that caused the target reduction between each two kill chain phases. This slide is a great opportunity to call out successes and challenges to leadership in an easy-to-understand way.

Credit: Lockheed Martin CIRT

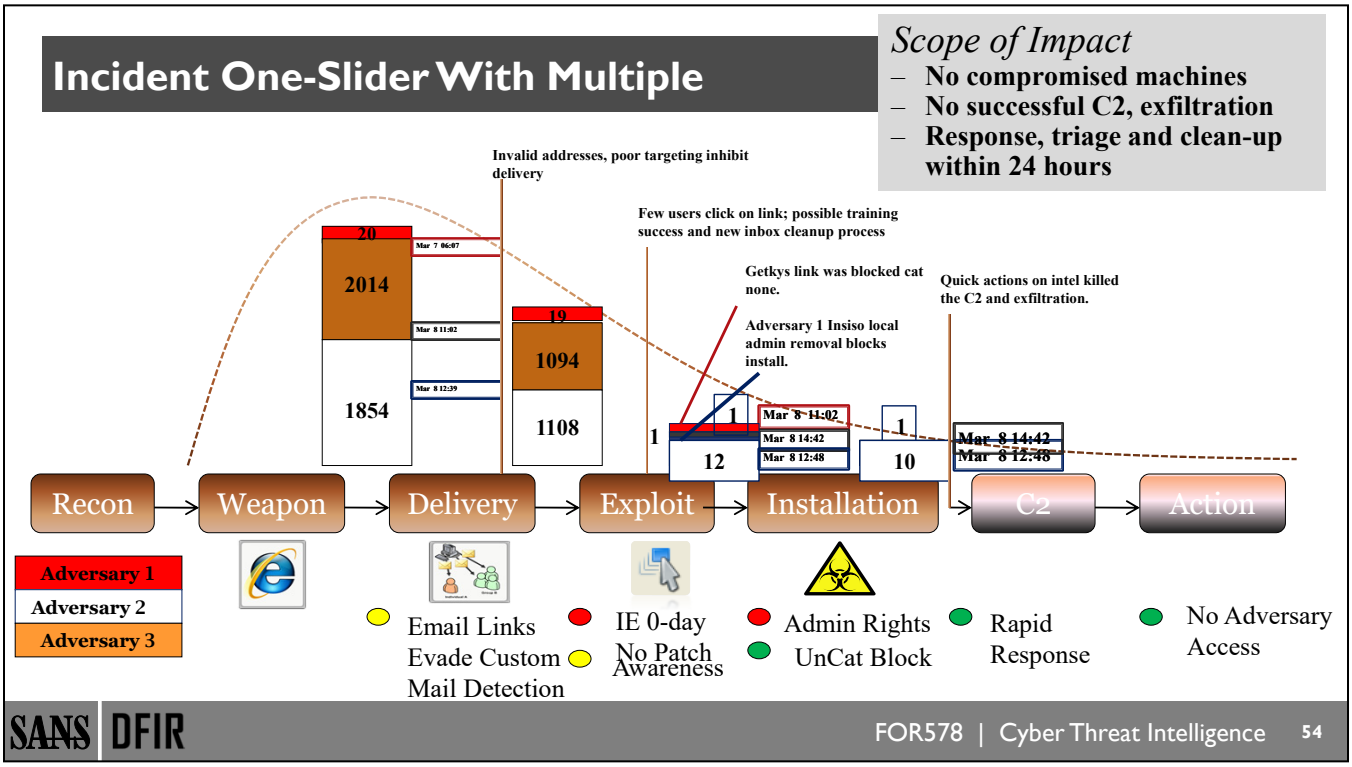

Incident One-Slider With Multiple

As we view three different adversaries together in the same way, we can draw some unique conclusions.

As an example, looking at Adversary 2 (1854 intrusions), we'd likely focus on them and Adversary 3 (2014) because they have the most intrusions. But in reality, Adversary 2 and Adversary 1 (20) have the most intrusions that actually get the furthest in the kill chain; Adversary 3 is well taken care of by our security controls.

Looking at Adversary 2 and Adversary 3, though, we have the open-ended question: Did we really stop them or did we lose visibility? We could then have an intelligence-driven hunt take us into the Actions on Objective phase inspired by those two adversaries to determine if we're already compromised.

Credit: Lockheed Martin CIRT

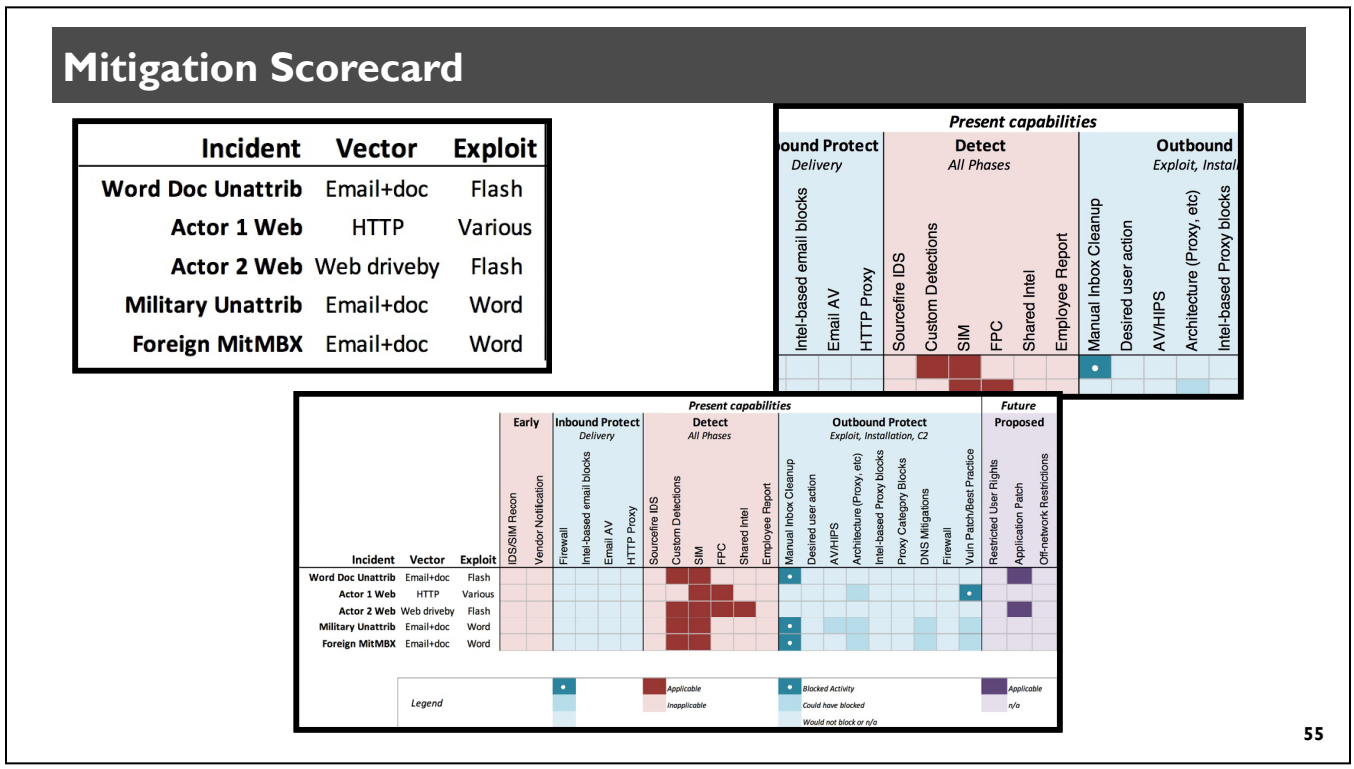

Mitigation Scorecard

The mitigation scorecard is one way to measure the utility of passive and mitigating courses of action. It maps specific incidents to the capabilities of network defenders, organized loosely by kill chain phase, providing a high-level visual of threat to courses of action mappings.

Columns represent capabilities that map to passive courses of action. Dark shading indicates the technology was applicable to the incident in a given row. The first group of columns is “early warning” or “over-the-horizon” capabilities, which may provide advanced knowledge of an impending intrusion. The second set of columns applies to all other phases of the kill chain after recon.

Other shaded columns contain mitigating capabilities or architectural decisions. Lightly-shaded cells represent capabilities that *would have* mitigated the activity based on their current configuration, had it not been mitigated by something else. Dark cells with a white dot represent the capability or architectural decision that, in reality, mitigated the intrusion. These columns are organized by “Inbound” and “Outbound,” which roughly applies to Delivery and Weaponization, and Exploit through C2, respectively.

Columns on the right represent proposed changes, new technologies, or corporate initiatives, and their applicability to each row, or intrusion attempt. Those cells that are shaded dark were applicable to the intrusion at the beginning of the row.

Credit: Lockheed Martin CIRT

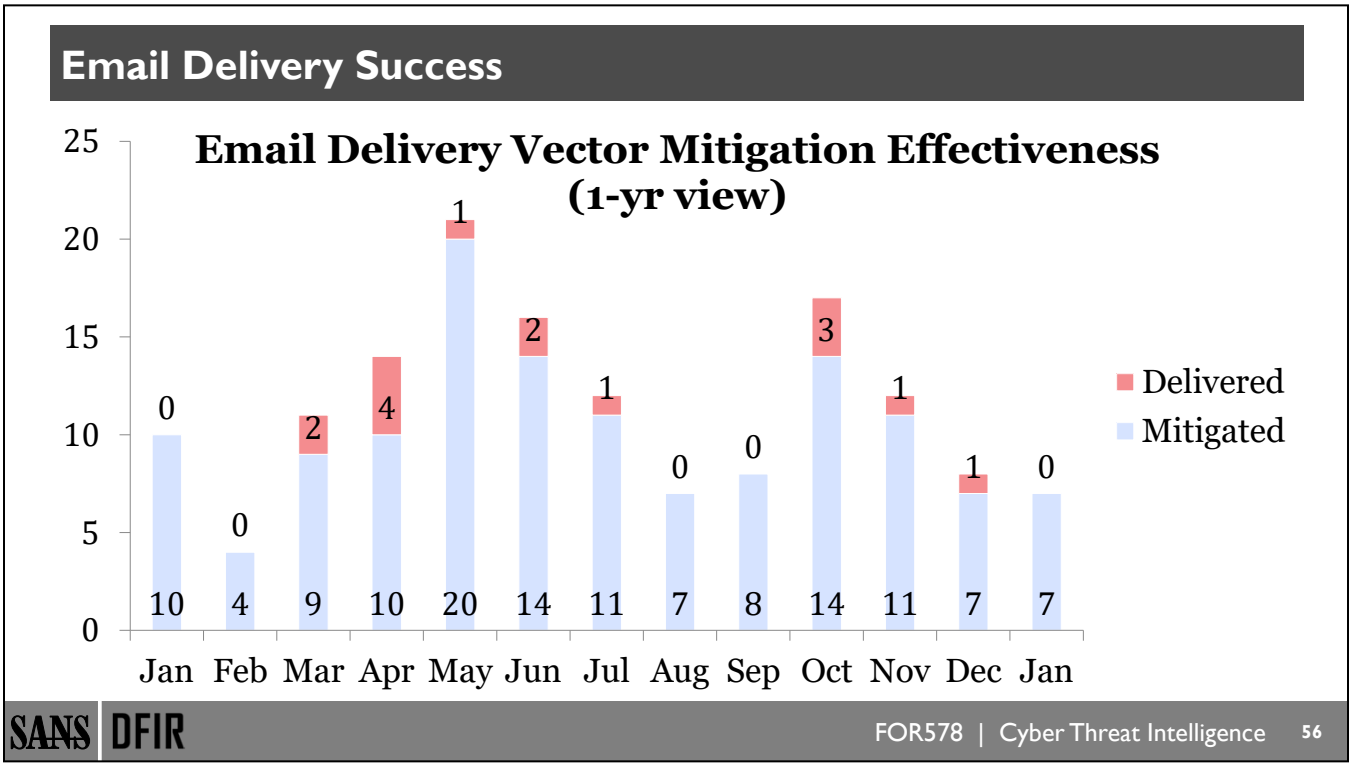

Email Delivery Success

This metric is relatively simple: It illustrates the overall success of adversaries using email delivery vectors over the course of a year. It is helpful in determining not only overall threat activity, but also will diagnose an underlying condition related to log availability and reliability.

Ideally, an independent and fully functional CIRT or CTI team will have reduced these measurements to near zero. This means there are no external dependencies to the success of your organization in defending itself. Of course, we have all learned that external intelligence provides crucial detail on the operations of adversaries beyond your organization's ability to detect and respond, or as a leading indicator of the nature of future intrusions.

Do note the small numbers here—in the 0–20 range. These are what are sometimes referred to as "waves," or an identical intrusion attempt spanning multiple days, users, or organizations. These attempts generally support a single objective and, while they may be counted on a per-target basis, are most usefully reflected in metrics as a single effort, as they are here.

Credit: Lockheed Martin CIRT

Analytical Completeness

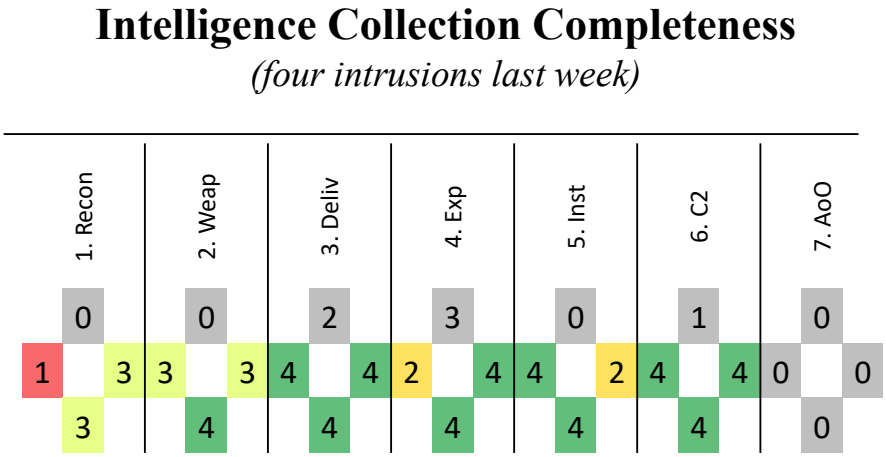

SANS DFIR FOR578 | Cyber Threat Intelligence 57

Analytical Completeness

This analytical completeness metric illustrates, over the four intrusion attempts against an organization from the previous week, the completeness of intelligence collection at each phase of the Kill Chain and Diamond Model. The numbers in each Diamond vertex indicate the number of intrusions for which *some* intelligence was collected. Highlighted vertices indicate those for which intelligence is normally collected, but for numerous intrusions appear to be missing.

Credit: Lockheed Martin CIRT

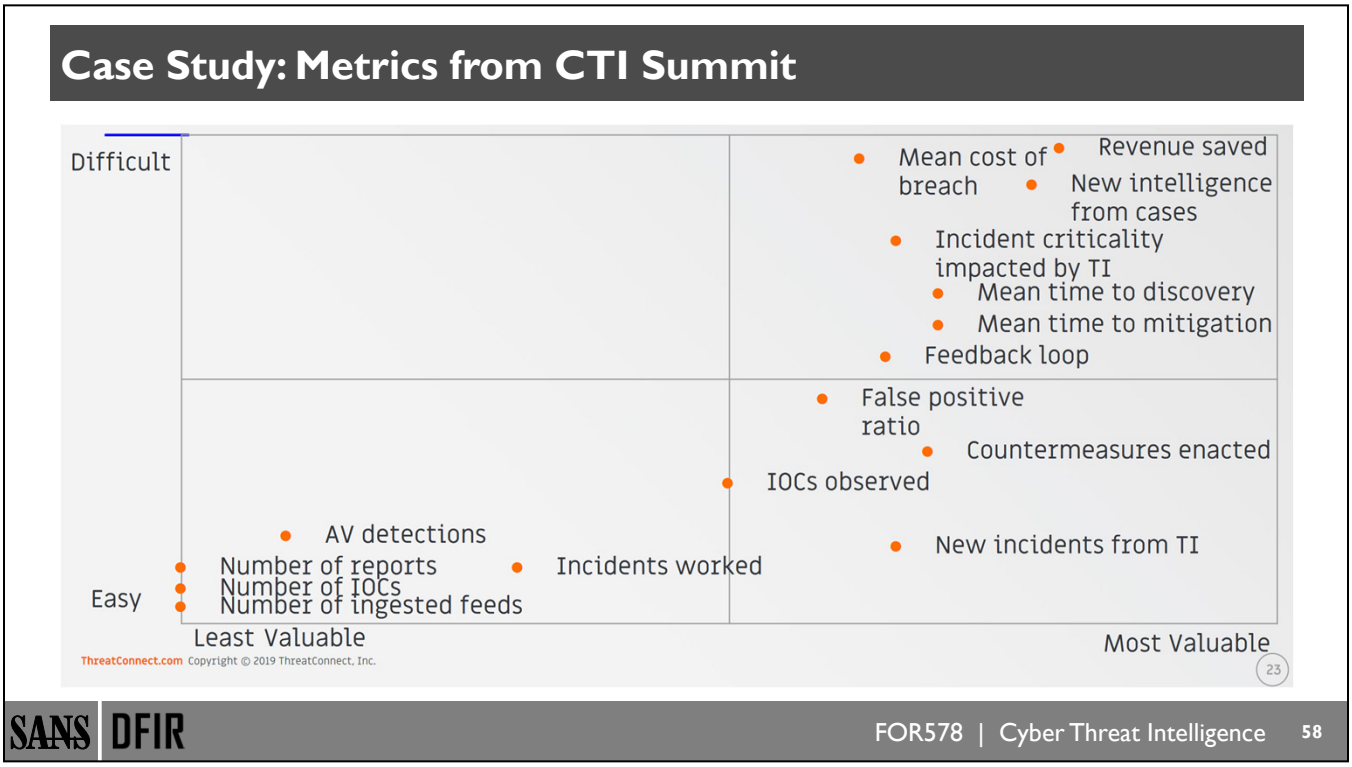

Case Study: Metrics from CTI Summit

This chart comes from an excellent presentation by Toni Gidwani and Marika Chauvin while they were at ThreatConnect and presented at the SANS CTI Summit. The concept behind this chart is that many of the most valuable metrics are also the toughest to identify. These were the metrics that worked best for them and their customers. For them, mapping to revenue saved, reduced mean time to discover adversaries, or reduced the mean cost of a breach were all very impactful. Your organization may value different metrics. By striving for meaningful metrics like these, you will be able to better justify the value of your CTI team. If you do not create metrics yourself, your leadership may create them for you and reduce your team's work to metrics like number of IOCs that don't properly explain your value. Take the challenge of finding good metrics into your own hands and create valuable ones for your team.

Remember, you can over-metric your team very quickly. Don't try to track everything but instead figure out the best metrics for you. Above all, satisfying intelligence requirements is what we measure ourselves on. These other metrics are just great data points on top of that, and the point of metrics is to help make us better—you should have some that are better than others.

Reference:

- https://www.first.org/resources/papers/london2019/1130-How-to-Get-Promoted-Gidwani.pdf

Exercise 5.3: Gaining Historical Perspective

- Move DERPYHOOVES into long-term tracking
  - Revisit historical intrusions
  - (Re)assign attribution as appropriate
- Threat metric for leadership: Campaign Heatmap
  - Must be rebuilt to trend campaigns w/new attribution
- Analysis of newly attributed intrusions reveals important patterns of activity
  - Historical duration of threat
  - Annual patterns of activity

Exercise 5.3: Gaining Historical Perspective

As we've achieved our intelligence requirements and understand the Activity Groups targeting Edison International, it's now a good idea to get into a place where we can track these threats long term. One effective mechanism to do this and to keep high level insights is a Campaign Heatmap. We're going to build this with a focus on RAINBOWDASH, PINKIEPIE, and DERPYHOOVES as well as some other offenders that other analysts on our team are tracking.

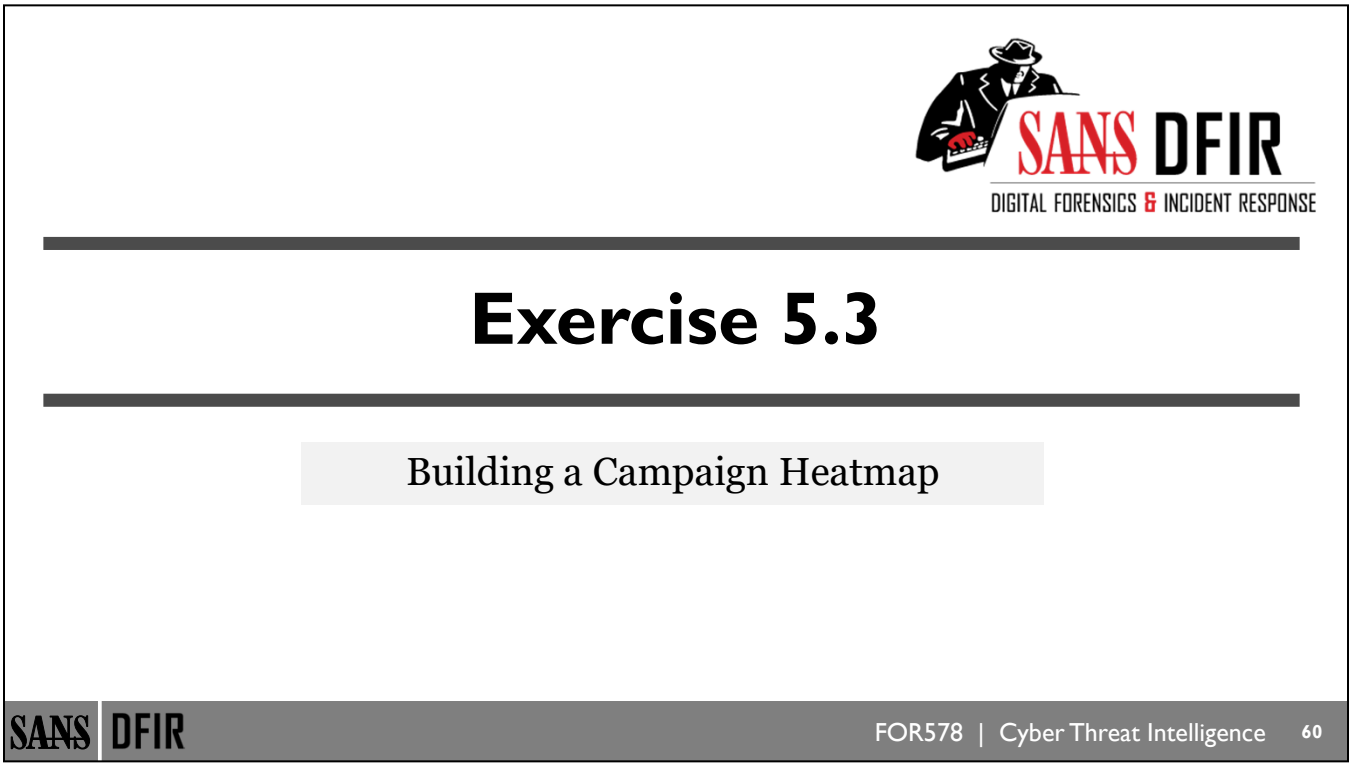

This page intentionally left blank.

Dissemination: Strategic

This page intentionally left blank.

Strategic Threat Intelligence

- Strategic threat intelligence is generally presented to executives with a focus on policy outcomes
- Threat intelligence presented at this level should be the most polished and complete intelligence product possible:
  - Generally, an audience that needs a complete picture
  - Direct attribution tends to be more valued
  - Tactical level defenders empathize with errors better
- Sample reasons to share strategic threat intelligence:
  - Global threat analysis and trends
  - National security and foreign policy
  - Security threats that impact business operations
  - Trade agreements or merger and acquisitions

Strategic Threat Intelligence

Strategic threat intelligence denotes the audience members as being the strategic level decision-makers; this grouping is sometimes mistaken to represent the type of data presented. It is okay to present technical information to strategic decision-makers, but the focus is usually not the technical information; it is what the technical information together means and what decisions need to be made. At the national level, all the technical data and analysis from the Operational and Tactical level are woven together, cross-analyzed, and produced into professional intelligence reports to reveal global threat analyses, trends, and implications for international organizations such as foreign policy or economic considerations. Attribution for national level threats is often important at the Strategic level as executives attempt to use various methods of pressuring appropriate actions from actors or choosing to take countermeasures.

Image:

Herb Kelleher quote

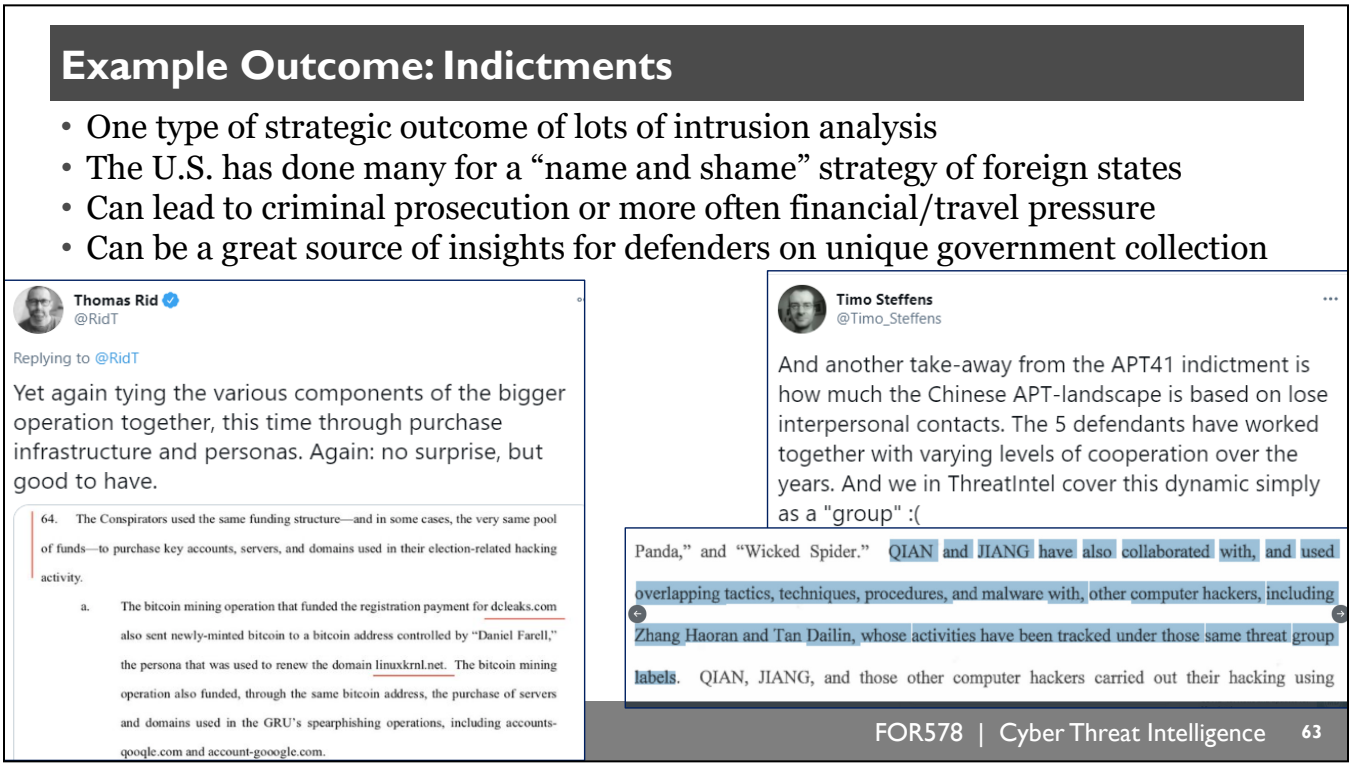

Example Outcome: Indictments

As part of a state’s strategy to hold other states or criminals accountable, they can use indictments. These indictments often take place from organizations like the U.S. Treasury or U.S. Department of Justice and carry with them penalties on organizations, freeze their assets, make it difficult for them to travel worldwide, etc.

They have varying degrees of effectiveness, but when used as part of an overall strategy as a single tool, they can be highly effective. There are concerns to be aware of, though. I’ve written extensively on the blow back effect of putting other military and intelligence agency’s people, especially those in uniforms, on wild west styled posters as if that is useful. What we perceive to be illegal is no more illegal than what those states perceive our government’s operations to be. The idea that we are the “good guys” is a very silly notion.

Regardless, CTI analysts should look to indictments as an example of the outcome of a lot of intrusion analysis to have a strategic goal. They also form a great source of collection for insights that governments have validated or contributed through their own unique collection.

References:

- https://twitter.com/Timo_Steffens/status/1306297038674558977
- https://twitter.com/RidT/status/1017915611501023232

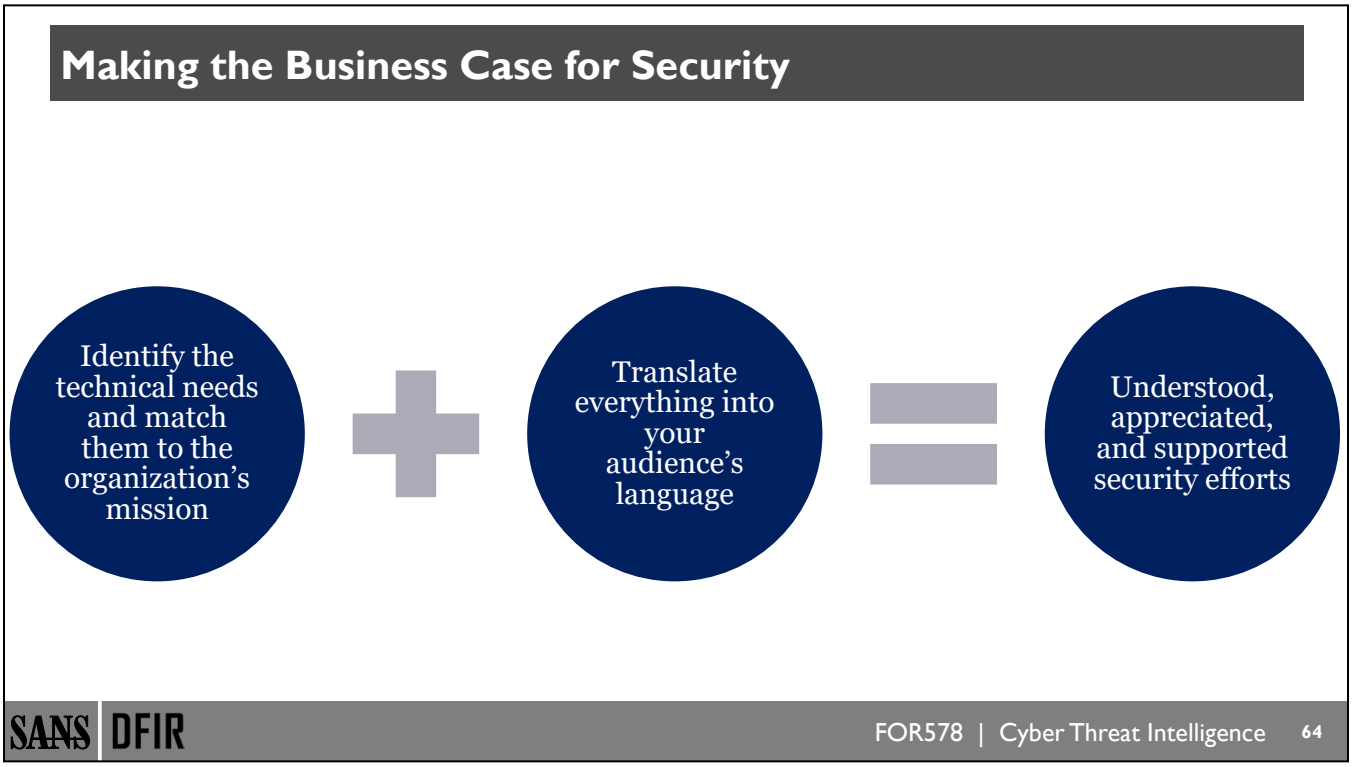

Making the Business Case for Security

To share cyber threat intelligence with the strategic level players requires analysts to understand and identify the technical needs, requirements, and considerations, and map it to the organization's mission. Understanding the impact to the organization and translating it into the audience's language will lead to understood, appreciated, and supported security efforts. Ultimately, security starts from the top, and the cyber threat intelligence analysts are in a great position to help drive security for the entire organization. Analysts should always avoid talking in overly technical ways or highlighting things that they think are technically cool but may not make sense or apply to the organization.

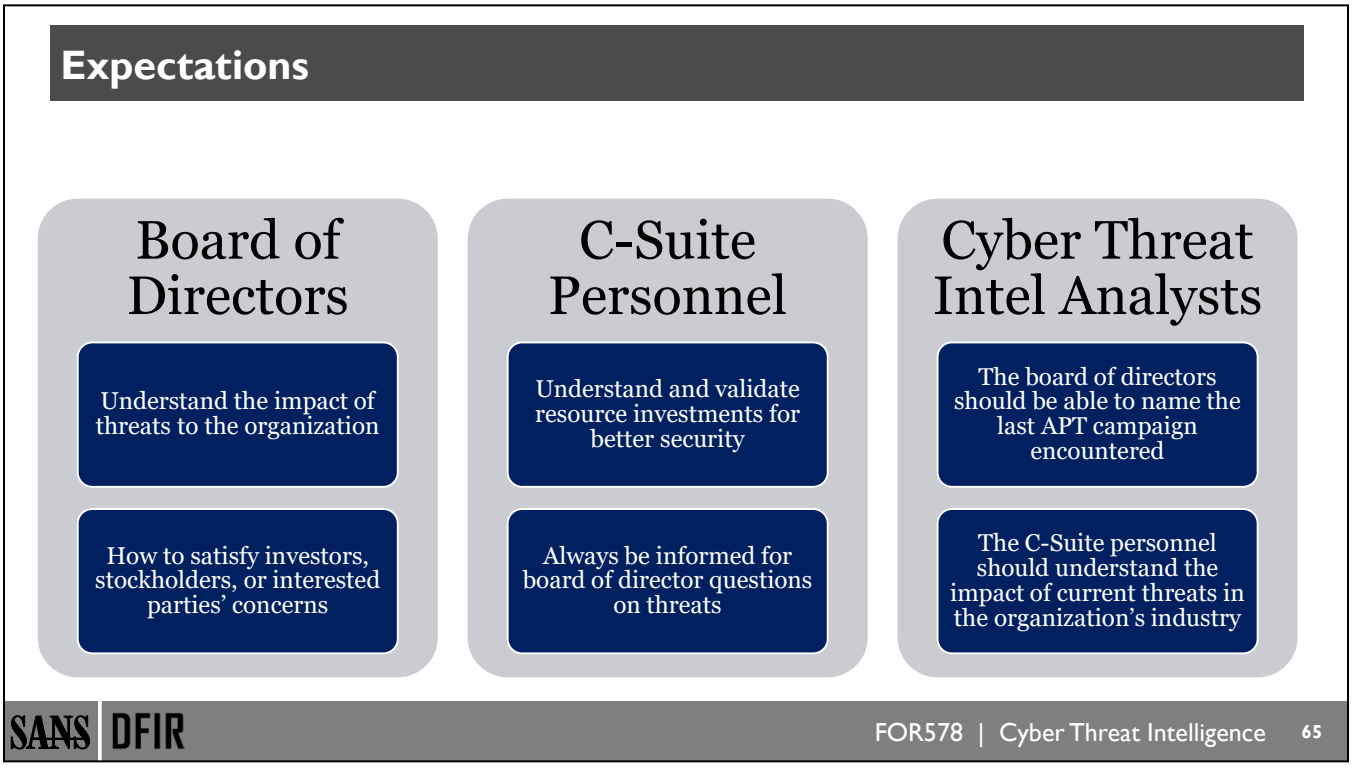

Expectations

Everyone has expectations. The board of directors of a company are stewards of the investments made into the organization and expect the security of an organization to be able to keep them aware of threats that are impacting the organization and how to satisfy the concerns of those interested and invested in the company. This is an ideal scenario, as many board of director members may not be aware of cybersecurity at all—this is where the cyber threat intelligence analyst must work through the C-suite personnel to remedy this situation.

The C-suite personnel want to know how to best spend their resources and validate their investments for better security at the organization. There are a number of competing interests and groups for limited resources: If money is given to security, the C-suite wants to know why and what the return is. Additionally, the C-suite should always be prepared for questions by the board.

Cyber threat intelligence analysts should expect, and help foster, that the board of directors can name the last APT campaign encountered that impacted the company. Additionally, the C-suite personnel should understand the current threats in the industry the organization exists in. For example, if the organization is a financial organization, then the C-suite should understand ongoing financial threats.

Cyber threat intelligence analysts may complain that their C-suite or board of directors does not care. It is your job to ensure—in their language—that they understand why they care. The C-suite personnel especially should never hear about a threat to their industry from any other source first other than the cyber threat intelligence team.

Lessons from the Field: Shoe Company and Anti-Hype

- A Fortune 500 shoe company has an impressive threat intelligence team made up of professionals from the private and government sector
- One of the chief outputs of the threat intelligence team is an email that goes out first thing every morning to the executives
- The email simply states:
  - What was in the news regarding cyber threats
  - All the hype that was put out and how it doesn't relate to the company
  - Relevant information the executives should know and focus on

Lessons from the Field: Shoe Company and Anti-Hype

Yes, a shoe company has one of the more impressive threat intelligence teams in the industry. Its understanding of what impacts its organization, such as the news and its hype that worries executives, and countering that, even through a simple email in the morning, shows a mastery of knowing itself and knowing the real threats. Threat intelligence teams practicing at the strategic level must be aware of the organization and its goals and must be creative in accomplishing them. This saves the entire security team, not just the threat intel analysts, time because the executives are more informed and are less likely to have time-sensitive needs to understand things it has heard. More important, the executives can learn about what is and is not relevant, which helps their buy-in to the security process. Security, especially related to threat intelligence, driven from the top is the most effective.

Reports/Narrative-Form Intelligence

- Reports should combine various sources of threat intelligence to present an overall and easily consumable narrative:
  - Should highlight the most important information that the intended audience should care about
- Reports are often the most lasting form of threat intelligence:
  - Should be meticulously created, edited, proofed, and cross-examined
  - Ensure completeness and professionalism
- The finished intelligence report may be the only view into the threat intelligence and security program:
  - Represents the work of many professionals
  - May be the only view an executive ever has into those teams
  - Likely to be shared throughout the organization

Reports/Narrative-Form Intelligence

Reports are often the longest-lasting form of intelligence. Where a technical indicator may expire or be phased out, there is intelligence reporting from the 1950s still present in a classified system somewhere. Likewise, intelligence reports at a company will be around for a long time and unfortunately may eventually be leaked or become public in some other way. Ensure that reports are proofed, edited, cross-analyzed, critically examined, and held to the highest of standards. They represent all the work of those at various levels of intelligence creation and defense.

Reference:

- https://www.sans.org/webcasts/top-10-writing-mistakes-cybersecurity-avoid-110220

Observation Versus Interpretation

- Analysts must clearly separate:
  - Facts, observables, evidence, from
  - Interpretations, conclusions
- This must be done:
  - In the analyst's mind
  - In the analyst's writing
- Accomplish this through consistent presentation of each:
  - Findings, then analysis
  - Analysis, then findings

Ambiguous writing begets ambiguous thinking and vice versa.

Observation Versus Interpretation

One of the most important aspects of analysis is the separation of the objective from the subjective—the what I see and the what I think—observation and interpretation. These two constructs must be clearly delineated in the mind of the analyst and also in the communication by the analyst. In both thought and writing, make this distinction unambiguous. The best way to do this is to develop rigor around how you communicate your findings; how you communicate influences at a deep cognitive level how you think, just as how you think influences how you communicate. By insisting on rigor in your writing that distinguishes between these two things, you can build the clarity of thought necessary to be a good analyst.

One way of doing this is to always order your analysis and your evidence (or your observation and your interpretation) the same way.

For example, we assess with high confidence that this intrusion is attributable to APT1. We make this assessment based on the following observations that are consistent with APT1 key indicators:

- The use of email address XXX in the delivery vector
- The use of Trojan YYY in the C2 phase
- Persistence in the installation phase using the registry key ZZZ

Or do the opposite:

We observed the following characteristics in this intrusion:

- The use of email address XXX in the delivery vector
- The use of Trojan YYY in the C2 phase
- Persistence in the installation phase using the registry key ZZZ

Based on the alignment of this evidence with key indicators, we assess, with high confidence, that this intrusion is attributable to APT1.

If you aren't thinking about these things distinctly, you won't write about them distinctly, and you will create ambiguity in the reader, making your work less compelling and, possibly, leading to misinterpretation.

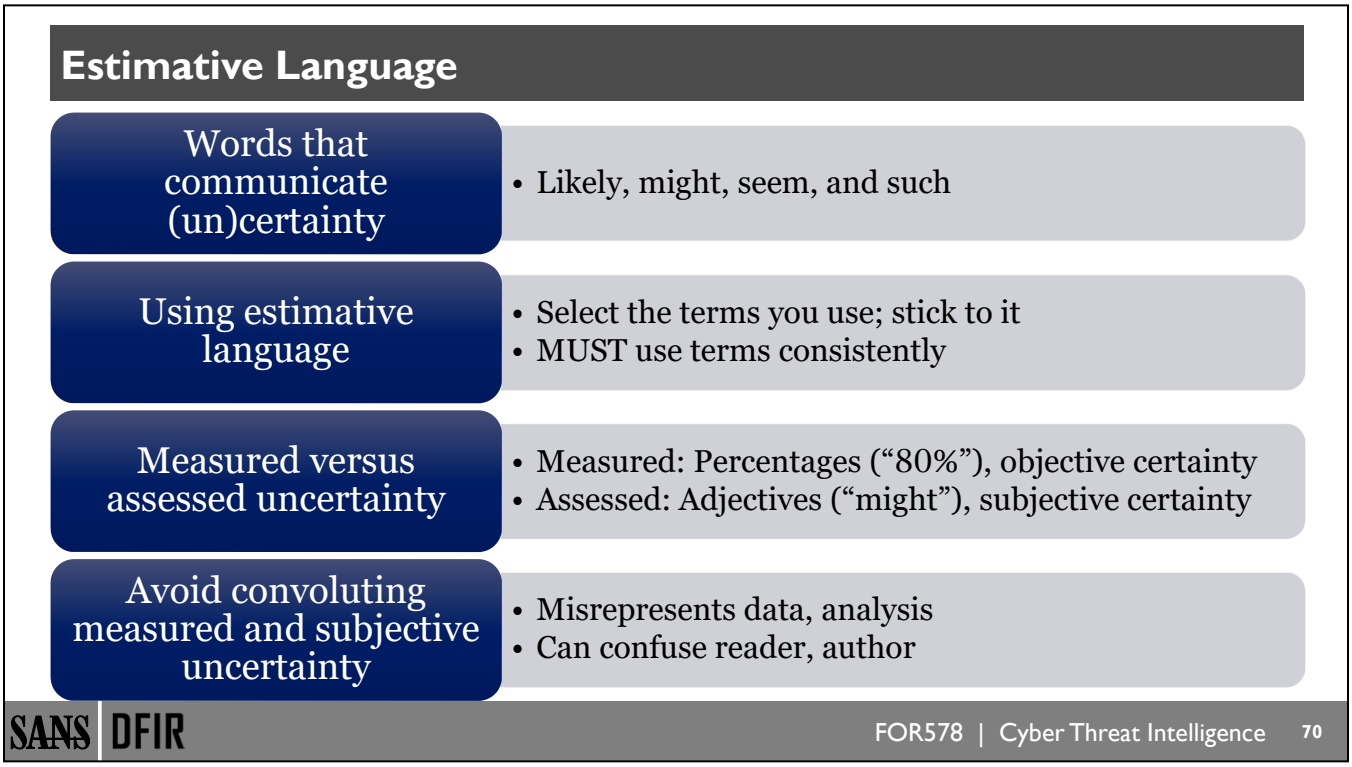

Estimative Language

Intelligence analysis is far more akin to qualitative scientific studies than those of a quantitative nature. In many cases, you compare the results of logical reasoning, not measurements and calculations. The means in which you communicate uncertainty is **estimative language**. Words such as "likely," "unlikely," "seems," and "might" are words expressing a degree of certainty. There are hundreds of such words. To be clear to the reader and in your thinking, it is critical that you select a set of uncertainty terms for your analysis and always use them consistently.

In the domain of cyber threat intelligence analysis, because we're often talking about the behavior of deterministic finite-state systems (computers), there is also a huge amount of measurable, objective data against which statistical methods can be applied.

Each of these analyses produces findings. The first is subjective and the second is objective. Both are important, and neither necessarily carries more weight than the other when discussing adversaries at an abstract level. However, they are fundamentally different and should be treated as such. One way to do this is to separate the estimative language used for each.

The nature of objective measurement and calculation lends itself naturally to numerical expression of certainty. For example, because adversary X uses weaponized PDF documents in 90% of past identified intrusions, I can say there is a 90% likelihood the next intrusion will contain a PDF (all other things being equal). Ninety percent is a specific, derived, numerical value based on objective measurement.

Conclusions that are the result of interpretation rather than calculation do not have the luxury of such precision. These assessments are subjective in nature, and uncertainty should be presented as such.

For example, “Based on the behavior of the adversary while on the victim network as compared with other intrusions attributed to adversary X, it seems likely that this intrusion is also attributable to adversary X.” “Likely” is a subjective estimative term. It clearly communicates the nature of the analysis from which it came.

The important differences between objective and subjective uncertainty and the frequency with which cyber threat intel analysts deal with both make it misleading to provide your conclusions in numerical form (“90% certain”) when there is no numerical basis for the evidence used. Even when rules are used to ensure analysts will provide a numerical value in a repeatable fashion (say, with the use of a common decision tree), the formulation of that decision tree is itself subjective, meaning that common subjective analysis is misrepresented as a measurable and objective quantity. Using the proper estimative language for the analysis conducted makes the nature of the analysis clear.

Estimative Scales

- Probability/uncertainty always falls on a scale:
  - Objective uncertainty is numerical (0–100%)
  - Subjective uncertainty is linguistic
- Understand the scale of terms you use
- Consider how facts changing would necessitate different estimative words

Remote **Unlikely** **Even Chance** **Probably, Likely** **Almost Certainly**

Estimative Scales

Be aware that all estimative language falls on a scale. In probability or objective uncertainty, this scale is typically from 0−100%. It's more ambiguous for subjective uncertainty. This can make people uncomfortable because it is open to interpretation, but that is the essence of this sort of analysis. The ambiguity of language is befitting a subjective analytical assessment.

That said, there is a relative scale of terms. When you use a term such as "likely," consider what other terms you could use if you felt more or less certain of your conclusion—if the evidence were more reliable or there was scarce evidence to support the finding, and so on. This can help you gauge which term is appropriate relative to what estimative language you would use were the situation different.

ALWAYS REMEMBER

- ASSESSMENTS ARE NOT FACTS
- ASSESSMENTS ≠ FACTS
- ASSESSMENTS != FACTS
- ASSESSMENTS <> FACTS
- ASSESSMENTS
  - ARE
    - NOT FACTS

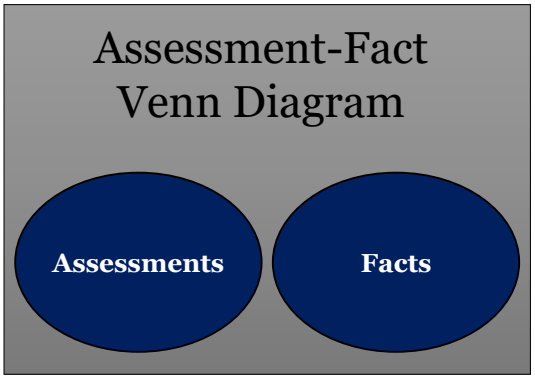

Always Remember

Assessments are not facts. This is important enough that we made a whole slide for this one statement. Assessments are not facts.

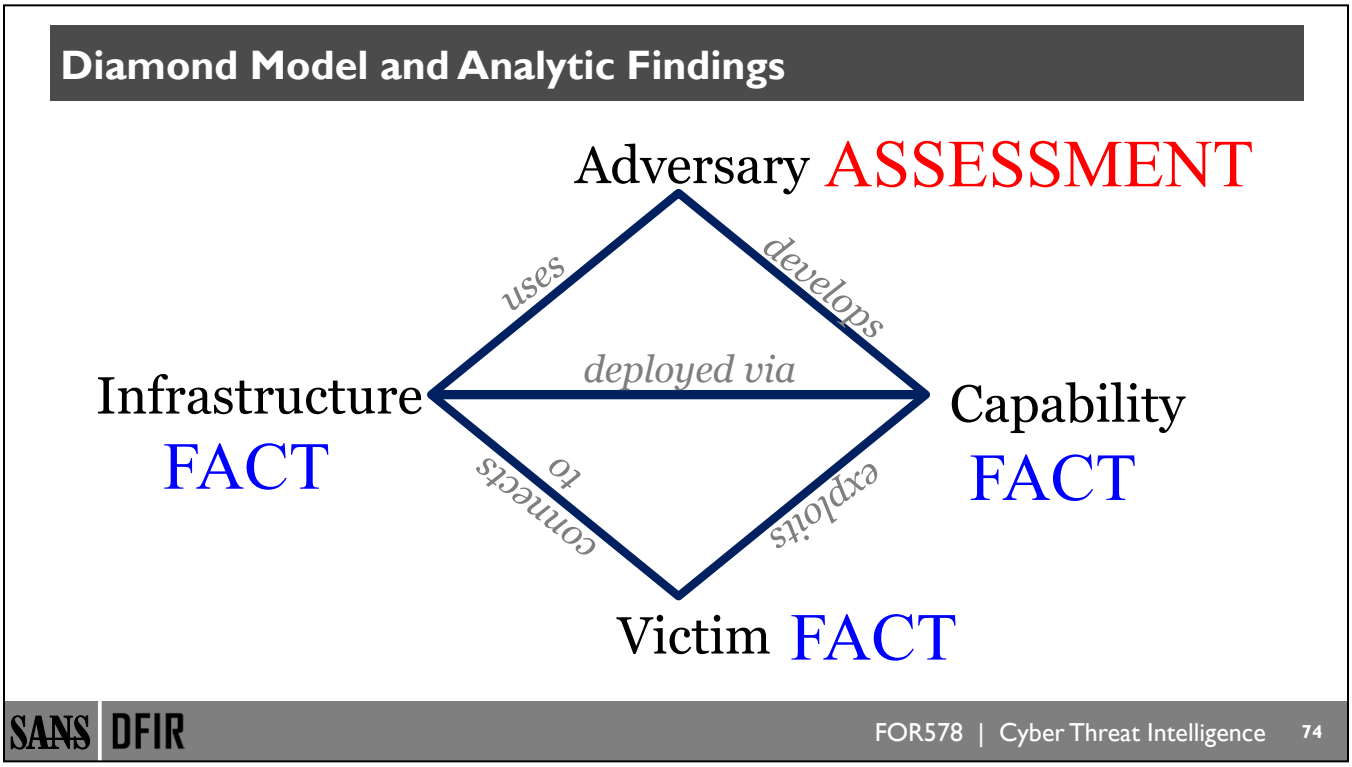

Diamond Model and Analytic Findings

In revisiting the Diamond Model, the infrastructure intelligence is factual in nature. There is little question about whether a packet came from a particular IP address. Similarly, the demonstrated capabilities are based in fact. The adversary used a particular Remote Access Trojan (RAT) or exfiltrated data using a specific method. Finally, the victims are also factual in nature. An organization that experiences a compromise isn't an assessment, although the fact that an organization would be potentially targeted would be. This brings us to the adversary. Due to the nature of cyber space, in many instances, it is difficult to know who the individual or even organizationally who is responsible for conducting these computer network operations. In communicating your analytic findings, for attribution, it is important to keep this in mind. In general, we say that three sides of the diamond (infrastructure, capability, and victim) provide information that enables us to make an attribution assessment. We now focus on the fidelity that certain types of indicators provide versus others in determining the proper language to use in our assessments.

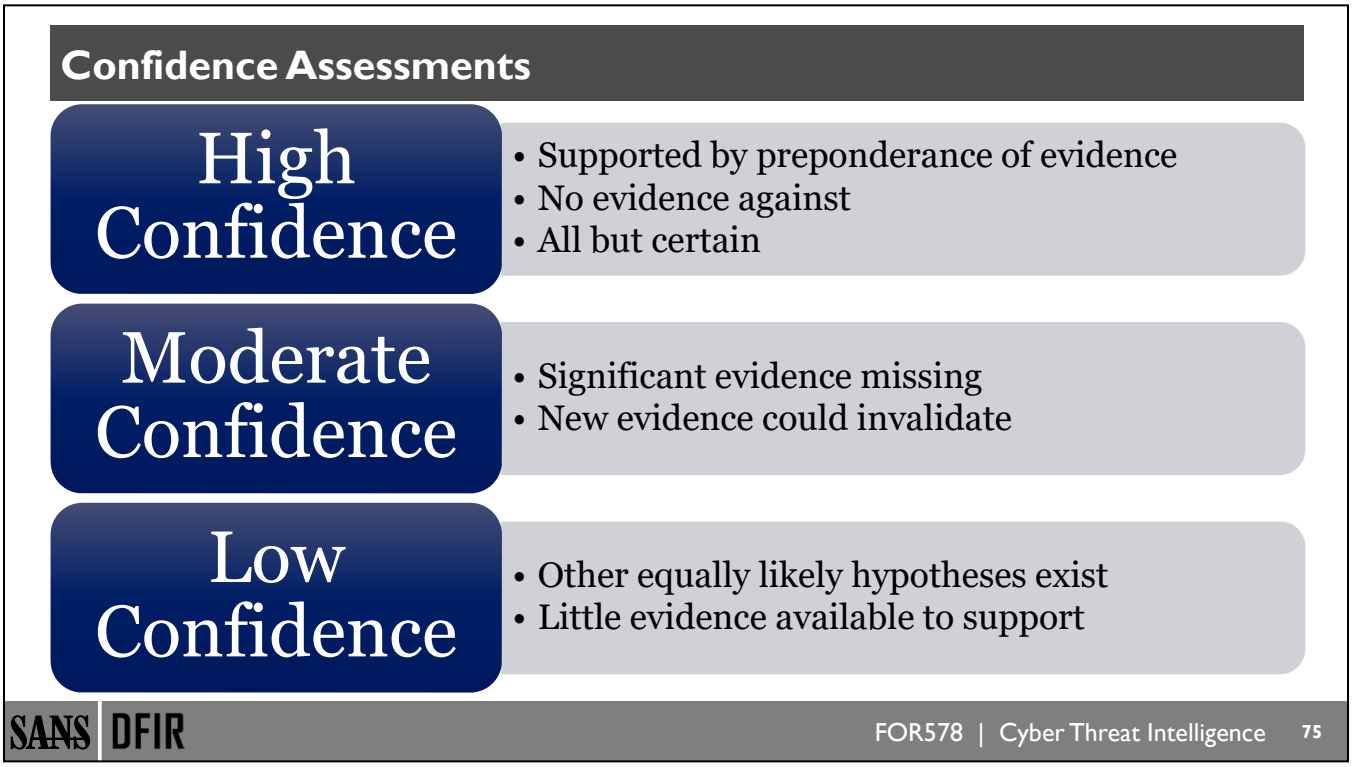

Confidence Assessments

One example of estimative language is confidence assessments. Confidence assessments are critical in communicating subjectively just how strongly you feel evidence supports your interpretation.

- High confidence assessments are those supported by a preponderance of evidence for which no or little evidence contradicting the conclusion is available. The likelihood of the assessment being true is all but certain.
- Moderate confidence assessments may be true but have significant evidentiary gaps, or some meaningful evidence exists that could make the assessment invalid.
- Low confidence assessments are those for which other valid hypotheses or explanations may exist, little evidence is available, or significant evidence against the assessment may exist.

Thumb Rules for Attribution Confidence

Although yours might look different, here is one set of rules of thumb analysts can use to qualify their assessments in a consistent manner.

- **Low confidence**: Intrusions that correlate to others in a campaign in only one phase of the Kill Chain, regardless of the number of indicators within that phase, are low confidence correlations.
- **Moderate confidence**: Intrusions whose distinct indicators align to other intrusions already correlated to a campaign in more than one phase of the Kill Chain, or correlate to a single phase of the Kill Chain with a generalized TTP alignment in other phases (though specific indicators may not align), are moderate confidence correlations. Another common set of criteria for moderate confidence attribution is the identification of two sides of the Diamond Model in any single stage, that is, *three* vertices.
- **High confidence**: Intrusions whose distinct indicators align to other intrusions in a campaign in two or more phases of the Kill Chain, as well as a generalized TTP alignment, are high confidence correlations.

Constructing Assessments

- Can be viewed as an equation

Assessment =
confidence + analysis + evidence + sources

- We assess with <insert confidence> that <insert assessment> because of <insert evidence> <insert sources>

Constructing Assessments

When we build our assessments, you can view the elements that should go into the assessment as variables in an equation. If we provide values for each of the variables, we should "solve the equation." There may be some variation to this, but in general, following this structure makes for the most complete assessments for CTI products.

By incorporating the indicator-tiering concept into this process, we can look at the indicator level of most confidence (that is, low, moderate, or high) and use that as a starting point for confidence level. For instance, if we have a tier 1 indicator such as an actor-registered C2 domain, we can assess with high confidence based on that indicator that the particular adversary was responsible. Similarly, if we have only multiple low-level indicators, it may be a judgment call as to whether we assess with low confidence or based on the preponderance of evidence that we move to a moderate level of confidence. Unfortunately, the latter process isn't quite as straightforward. We rely on our subject matter expertise, the amount of supporting data points, and possibly the results of an ACH analysis discussed on Day 1 to get to a better decision.

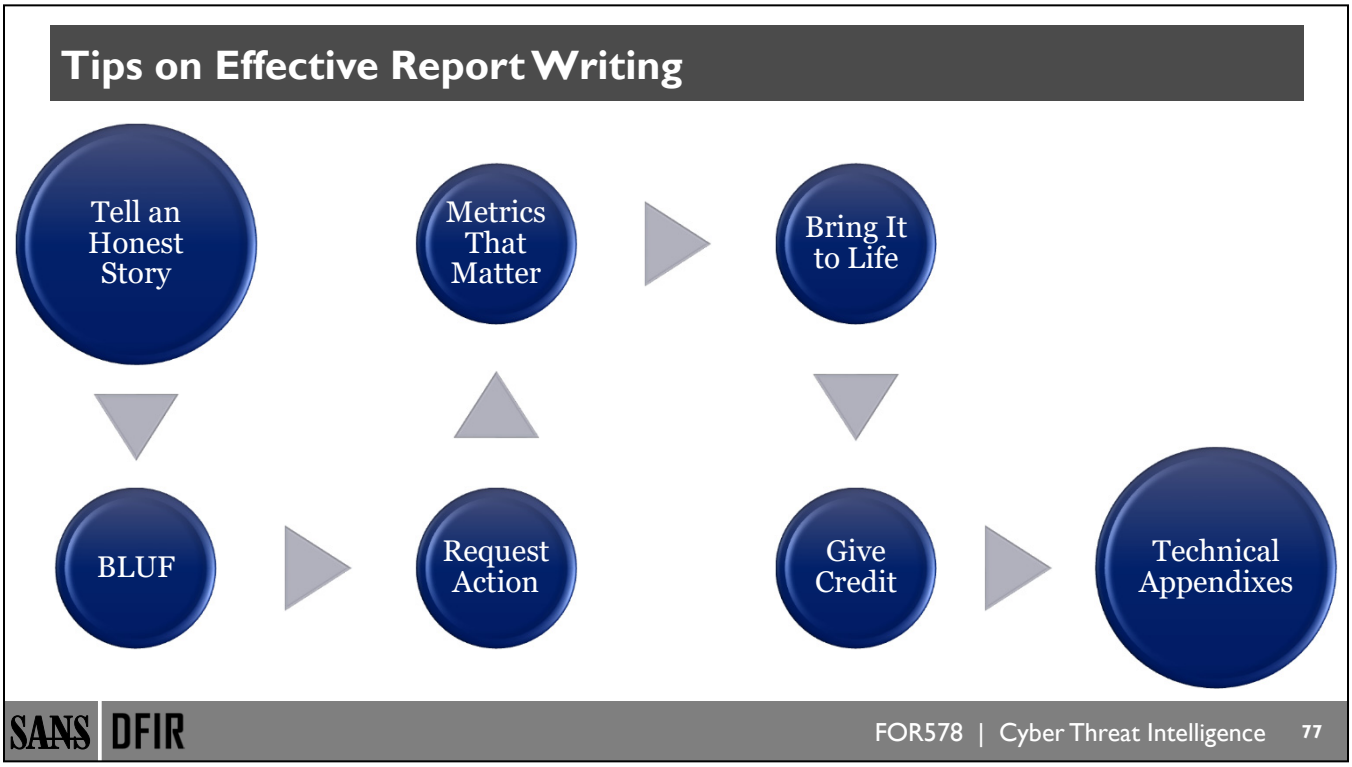

Tips on Effective Report Writing

Report writing is a skill that is developed over time, but certain tips should always be considered. First, the purpose of the report is to tell an honest story. Whatever the organization needs out of the intelligence must be told but in an honest, non-alarmist, way. The report should always include an Executive Summary or BLUF (Bottom Line Up Front) for executives or readers who have only 5–10 seconds to read the report. Ensure they take away the right things. Also, be sure to have some sort of requests for action or let readers know this is only for their information and situational awareness. If you want something done, though, make it clear.

Any pictures, technical data, or metrics are only there to highlight points and add to the story. They should not be overly relied upon or overvalued. The metrics have to matter and be consistent with the story, whereas the most important information must be pulled up to the start of the report so that the passing individual with 5 seconds to read the report takes away the appropriate key points. In addition, some folks fail to give credit to the teams involved with various aspects of reporting. Success is not limited and should be shared when possible; the credit given to one team is likely to be repaid back over time in a positive culture if nothing else. In addition, anyone writing threat intelligence reports should consider Sergio's "15 Things Wrong with Today's Threat Intelligence Reporting" a required reading (http://www.activeresponse.org/15-things-wrong-with-todays-threat-intelligence-reporting/).

Always link to the technical data or include it in an appendix. Technical folks will likely get the report from an executive at one point or another and need to know where to go get the data. It is a common courtesy aspect of writing threat intelligence reports. It also ensures others can validate the report.

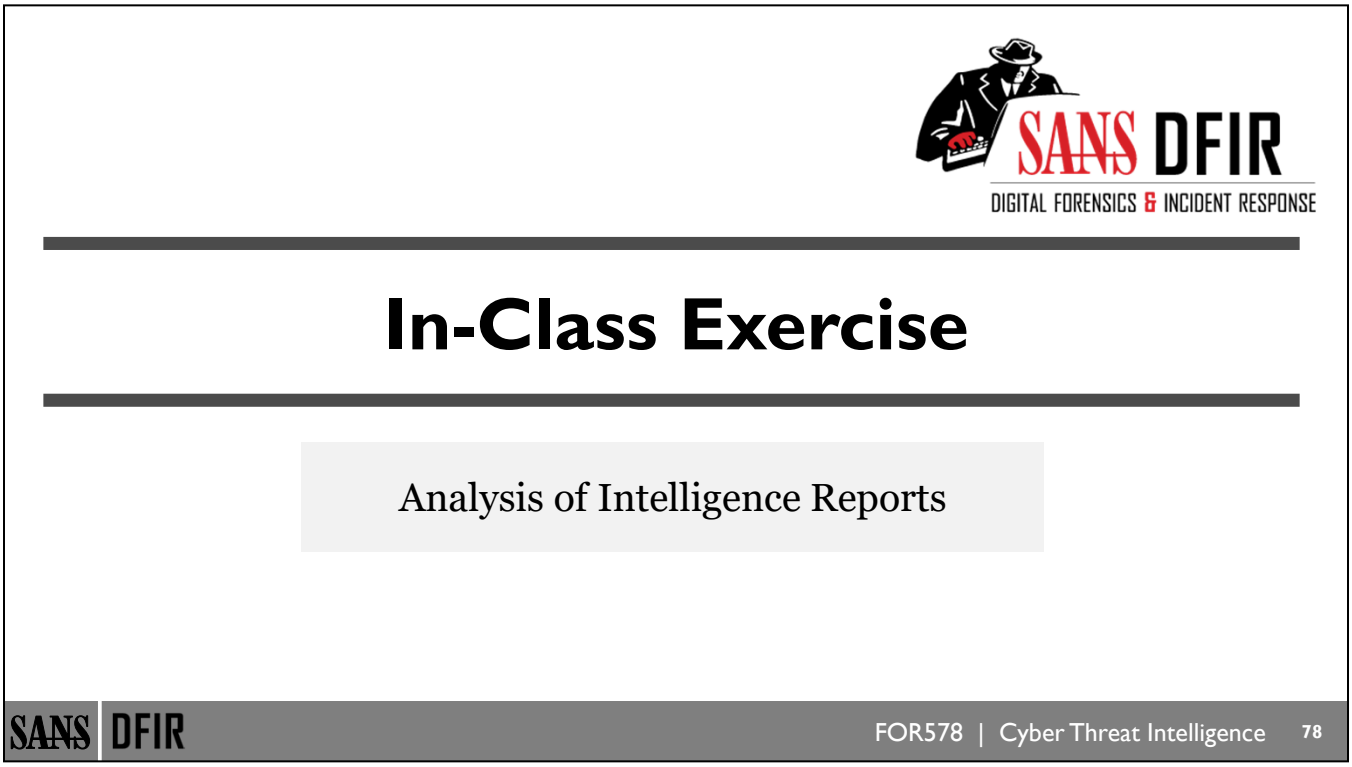

Refer to the Threat Intelligence Reports folder in your course media file under Supplemental Materials. The purpose of this thought exercise is not to deeply analyze the technical data presented in each report. The purpose is to look at the report writing and the style at which each approached. From this, you can learn the pros and cons of each report, which reveal lessons learned for other analysts. That is, what did these vendors do correctly or poorly (in your opinion) with their report writing styles?

There are few better ways to get better at writing reports than to critically evaluate other reports. Evaluating especially good and especially bad reports can expedite report writing and communication skills of analysts.

Proofpoint's North Korea Bitten by Bitcoin Bug

- Spend a few minutes reading Proofpoint's Report
- Identify things you like and do not like

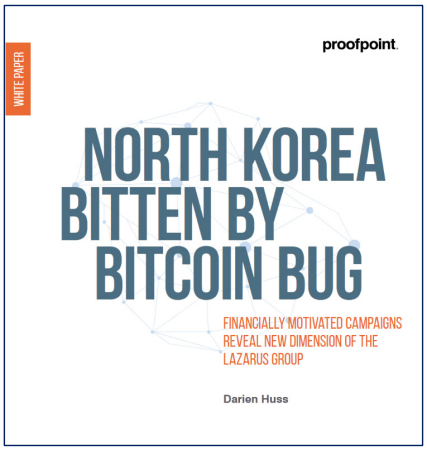

Proofpoint's Report

Spend a few minutes skimming the report (ProofPoint_us-wp-north-korea.pdf). You do not need to read every word; simply skim for things you like and do not like.

Reference:

- https://app.box.com/s/xez1hl78xz2l55mqe5cqvlwb5ytckhxf

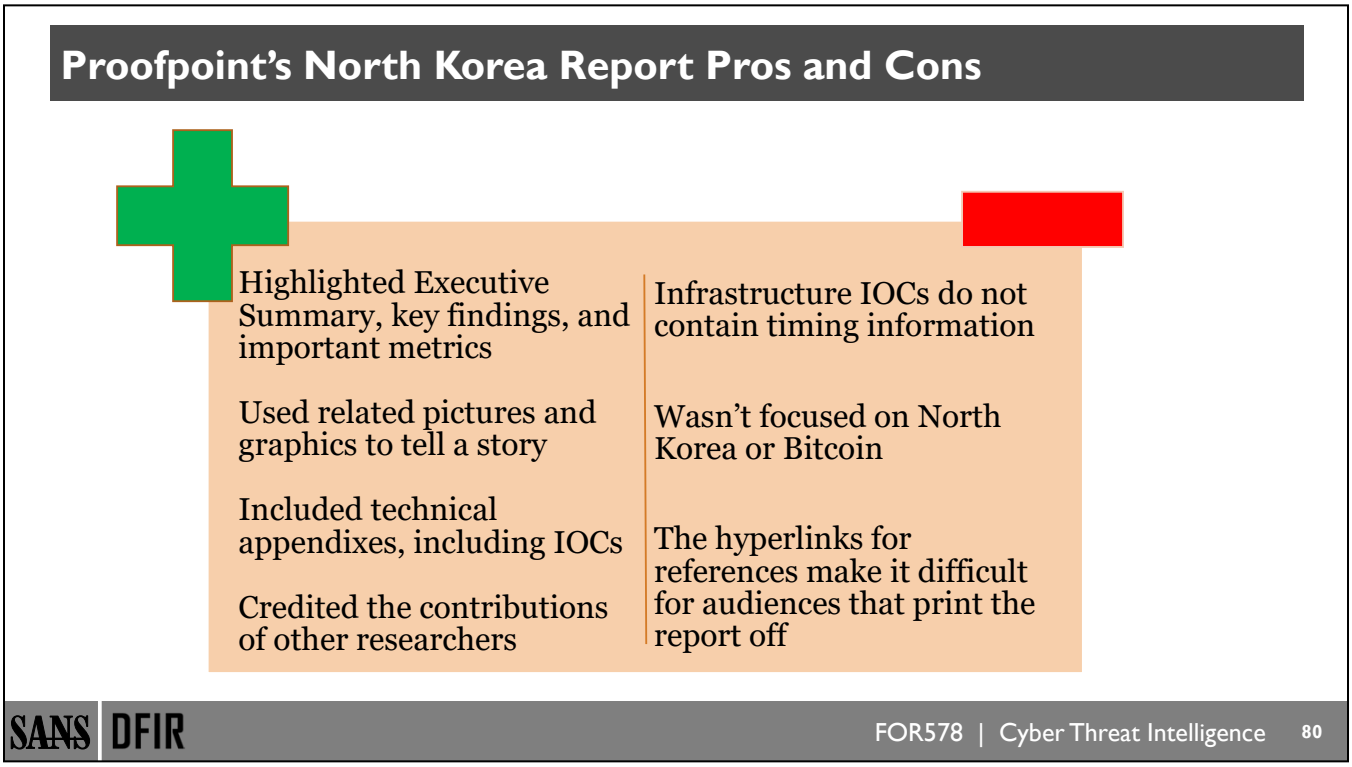

Proofpoint's North Korea Report Pros and Cons

This report is written as just that, a report. It's not the traditional style of an intelligence report but it utilizes threat intelligence and an understanding of the adversary to present a well-structured document forward. Right up front, we see a good understanding of the table of contents with an executive summary.

The report is fairly technical in nature, so it would be for practitioners, but contains sufficient usable information and a look at the various components of the adversary's capabilities. Additionally, the screenshots are well cropped and sorted for folks to use. The campaign timeline is also a nice touch to show an understanding that this is a long period of understanding of the adversary.

Importantly, the researcher gives credit to the contributions of others and includes IOCs as an appendix after the conclusion.

The IOCs that are external infrastructure, such as IP addresses, should come with timing information such as when they were observed to be malicious to help make the IOCs more useful. Additionally, the bolded words are hyperlinks. Many reports are still printed off to be read; it would be much better to include the links as references in an appendix instead of hyperlinking in the text.

Norse's Iran CIB

- Spend a few minutes reading Norse's Iran CIB Report
- Identify things you like and do not like

Norse's Iran CIB

Spend a few minutes reading Norse's Iran CIB report (Norse_JIB_IRAN_011_JANUARY_27_2015.pdf).

Reference:

- https://www.aei.org/

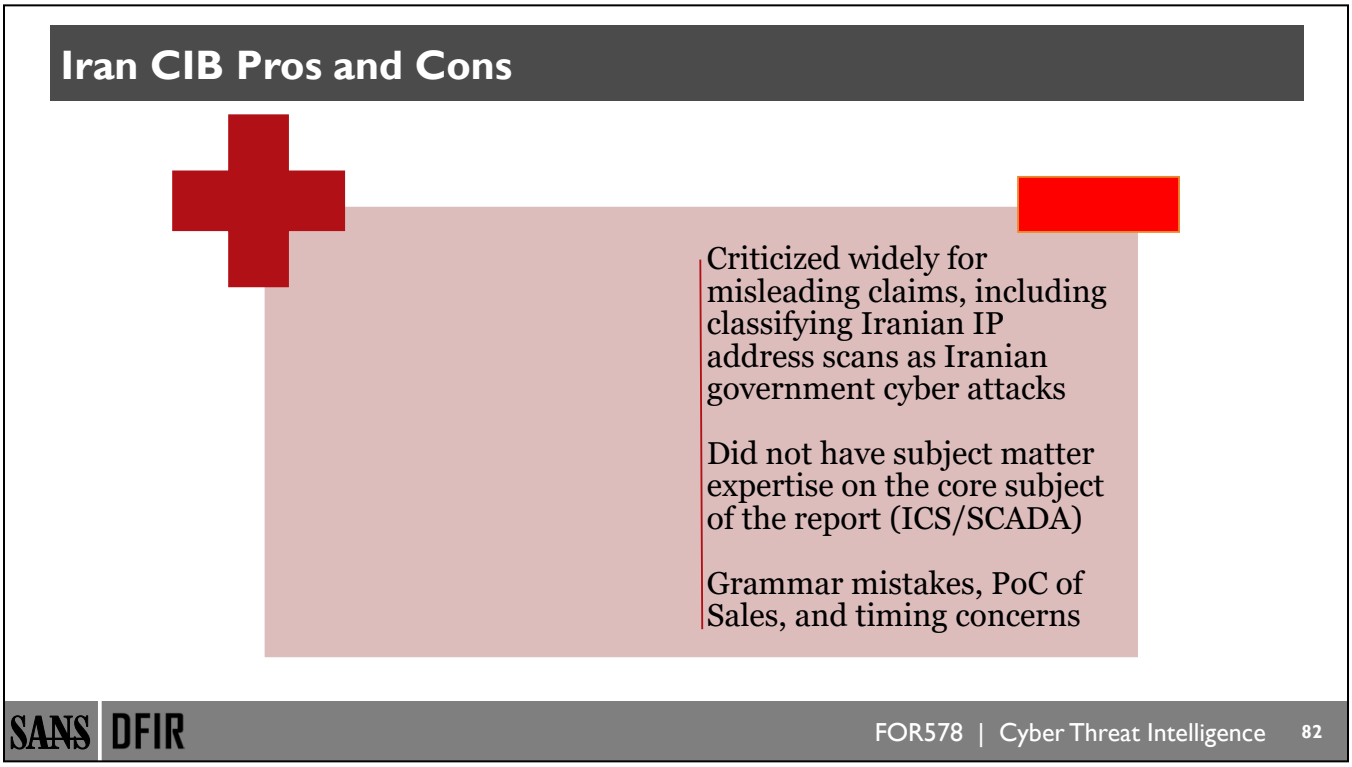

Iran CIB Pros and Cons

Background

This threat intelligence report had a bit of controversy surrounding it, which is good to put into context for the purpose of evaluating it. This also drives the point home that while we are analyzing threat intel reports, we need to view the full picture as analysts. This report was released in February 2015, as a TLP:GREEN report from Norse; its team called together a group of senior DoD and DHS members to brief the report to them and warn of an Iranian cyber offensive. The report met criticism primarily because the report did not make it clear that the attribution done was based on IP addresses alone and that the attacks on critical infrastructure were scans against IP addresses not associated with real infrastructure. (Each scan against a TCP or UDP port related to ICS was considered an attack.) Following the criticism, Norse worked with the conservative think tank AEI to release a more academic version of the data with policy recommendations against the (at the time) Iranian nuclear negotiations. The purpose of adding this report to this exercise is not to shame Norse or AEI. It is a perfect case study, though, to analyze the crossover of an intelligence report into academia and into policy recommendations and where there are opportunities to learn from what occurred. The critique of those actions is good context for this exercise but are outside the scope of this exercise; therefore, the following references are for those interested. For this exercise, focus on the merit of the report alone.

The Report's Pros and Cons

Norse's Iranian report did grab a lot of public attention about cyber threats. Specifically, Iran has been a concern for many national leaders and cybersecurity personnel for years now while it is usually not seen as a top-tier player. In this way, Norse forced some discussion about the threat and what others were observing or not observing. For the cons, though, this report gained a lot of criticism for the claims made. Norse did not do a good job of educating the audience that the message of "Iranian government is attacking critical infrastructure" was an analysis of Iranian IP addresses scanning or performing reconnaissance against unregistered IP addresses.

Norse could have done better if someone familiar with Iranian intelligence operations (or at least an intelligence analyst) had written the report with someone who had ICS expertise. (For those unfamiliar with ICS, some of the systems listed did not make sense in the analysis, including the use of the DNP3 protocol as "an ICS" and the discussion of events such as web defacement in relation to ICS where it wouldn't matter or be related.) The authors of this report had neither, and it showed in some of the analysis. Another misstep is that the report has multiple mistakes and came at a sensitive time in Iranian nuclear negotiations. Why is this such an issue, though? Grammar. The chart that shows "Jan 15 – Dec 15," capitalizing words differently throughout the report, using hashtags and acronyms where they didn't make sense, and so on—why is the focus on that? Because these oversights make it appear that the report did not go through a full intelligence life cycle and that it appears to be rushed. Intelligence can never look rushed because it devalues the intelligence. A full intelligence life cycle should always have personnel catching the small mistakes as they read over and cross-analyze the report.

This is not meant to bash Norse, though. Every company makes missteps; it's important to learn from them.

References:

- https://www.nytimes.com/2015/04/16/world/middleeast/iran-is-raising-sophistication-and-frequency-of-cyberattacks-study-says.html
- https://www.csmonitor.com/World/Passcode/Passcode-Voices/2015/0417/Opinion-Security-firm-s-Iran-report-mostly-hype
- https://www.aei.org/research-products/report/growing-cyberthreat-from-iran/
- https://www.thedailybeast.com/the-overhyping-of-irans-cyberarmy

Kaspersky's Equation Group (Optional)

- Spend a few minutes reading Kaspersky's EG Report
- Identify things you like and do not like

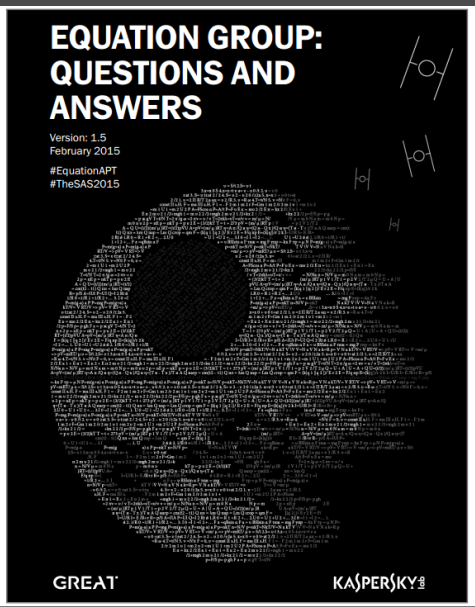

Kaspersky's Equation Group

Optionally, spend a few minutes reading Kaspersky's Equation Group (EG) report (Equation_group_questions_and_answers.pdf).

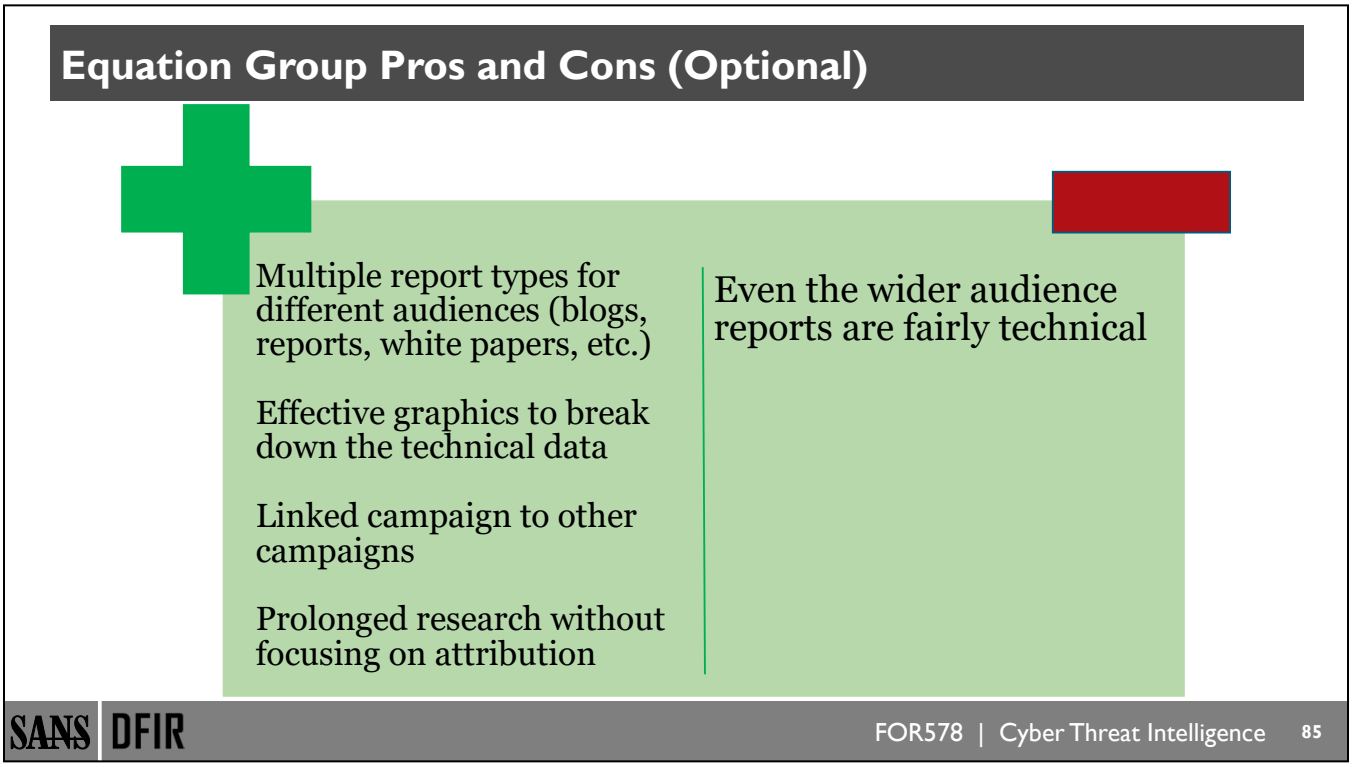

Equation Group Pros and Cons

The Kaspersky Lab folks create some of the most technical reports in the business. To accommodate for their highly technical nature, they released this report in multiple forms, including blogs, reports, and white papers, to reach a wider audience. They effectively used graphics to break down the technical information, and they linked the various campaigns they've observed together. This brought context to the reader and helped enforce that what was observed was not simply an intrusion but a true adversary campaign. This also helped show that Kaspersky's team performed research over a long period of time and gathered hundreds of intrusions to gather a strong picture of what was going on. It also did not try to focus on attribution more than the technical material. A con is that even the more technical Kaspersky reports are often too technical for the normal audience. It tries to accommodate for this in news articles and blog posts, but it could do a better job for executive leaders.

Case Study:
APT10 and Cloud Hopper

This page intentionally left blank.

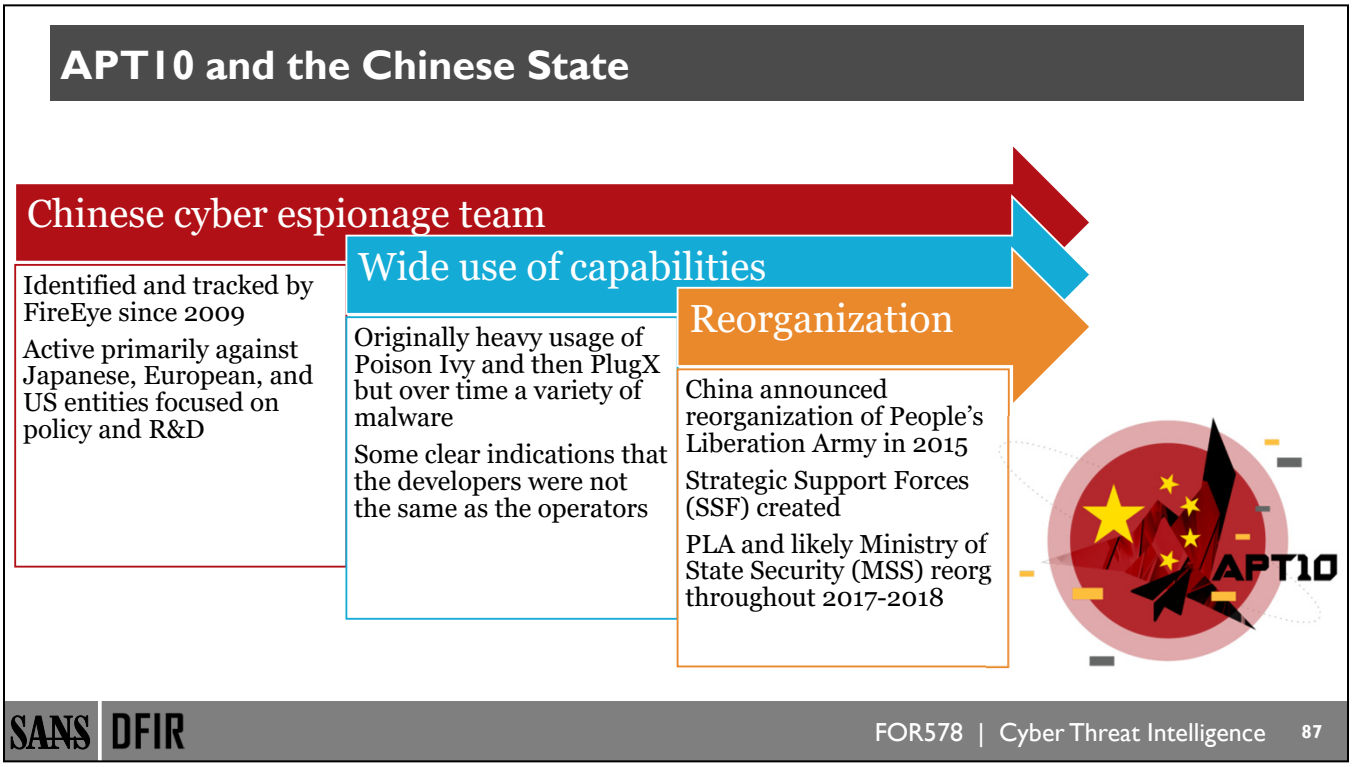

APT10 and the Chinese State

APT10 is a threat discovered and tracked by FireEye that is of Chinese state origin and has been operating at least since 2009. APT10 historically had taken significant advantage of the Poison Ivy malware and eventually moved to PlugX and other now common families of malware. Interestingly, there was a clear progression in the malware capabilities to be more user friendly, indicating a clear difference in the development team of these capabilities and the operators who leveraged them.

APT10's targeting is fairly wide with organizations around the world but took a particular focus on Western defense companies, European organizations, and a large focus on East Asian and Japanese manufacturing and policy focused organizations. Some of their more well-known breaches include the left of advanced aerospace research and development intellectual property for the advancement of China's national security goals.

The overall APT10 family uses a wide variety of capabilities, has a wide variety of victims, and has a wide variety of TTPs. FireEye and team have done a fantastic job tracking APT10 over the years; but it's worth pointing out that if a defender wanted to protect against "APT10" there's so much data, insights, TTPs, etc., that it would be particularly difficult. We will explore this later in the class around different ways to track groups to be able to slice/dice the insights for more rapid consumption tailored to defenders' needs. This is not in any way intended as a slight against FireEye but just meant to expose different examples of what you can do.

References:

- https://www.fireeye.com/current-threats/apt-groups.html
- https://www.fireeye.com/content/dam/fireeye-www/summit/cds-2019/presentations/cds19-executive-s08-achievement-unlocked.pdf

APT10 and the US Government

- In 2018, the US Department of Justice (DOJ) and Federal Bureau of Investigation (FBI) indicted two individuals and the broader APT10 group
- FBI's indictment alleged activity from 2006-2018, including intellectual property theft across aviation, manufacturing, oil and gas, computer R&D, maritime, and personal details of 100k Navy personnel

APT10 and the US Government

The work of FireEye and other firms contributed to the US DOJ and FBI's indictment of two individuals associated with APT10 and a broader call out of APT10. The indictment details are as follows:

Details:

On December 17, 2018, a grand jury in the United States District Court for the Southern District of New York indicted ZHU HUA, aka "Afwar," aka "CVNX," aka "Alayos," aka "Godkiller," and ZHANG SHILONG, aka "Baobeilong," aka "Zhang Jianguo," aka "Atreexp," two members of a hacking group operating in China known in the cybersecurity community as Advanced Persistent Threat 10 (the "APT 10 Group"), with conspiracy to commit computer intrusion, conspiracy to commit wire fraud, and aggravated identity theft. The defendants worked for Huaying Haitai Science and Technology Development Company located in Tianjin, China, and they acted in association with the Chinese Ministry of State Security's Tianjin State Security Bureau.

As alleged in the Indictment, from at least 2006 through 2018, the defendants conducted extensive campaigns of global intrusions into computer systems aiming to steal, among other data, intellectual property and confidential business and technological information from more than at least 45 commercial and defense technology companies in at least a dozen states, managed service providers ("MSP"), which are companies that remotely manage the information technology infrastructure of businesses and governments around the world, and U.S. government agencies. The victim companies targeted by ZHU HUA and ZHANG SHILONG were involved in a diverse array of commercial activity, industries, and technologies, including aviation, space and satellite technology, manufacturing technology, oil and gas exploration, production technology, communications technology, computer processor technology, and maritime technology. In addition, for example, the APT 10 Group's campaign compromised the data of an MSP and certain of its clients located in at least 12 countries including Brazil, Canada, Finland, France, Germany, India, Japan, Sweden, Switzerland, the United Arab Emirates, the United Kingdom, and the United States. The APT 10 group also compromised computer systems containing information regarding the United States Department of the Navy and stole the personally identifiable information of more than 100,000 Navy personnel.

The full indictment may be read here: https://www.justice.gov/opa/press-release/file/1121706/download

Reference: https://www.fbi.gov/wanted/cyber/apt-10-group

Indictments for Attribution: APT10

- Government intelligence can support attribution such as the group (APT10), person (Zhu Hua), persona ("Afwar"), state (China), Organization (Huaying Haitai), Customer (MSS Tianjin SSB), and role (infrastructure team)

2. From at least in or about 2006 up to and including in or about 2018, members of the APT10 Group, including ZHU HUA, a/k/a "Afwar," a/k/a "CVNX," a/k/a "Alayos," a/k/a "Godkiller," and ZHANG SHILONG, a/k/a "Baobeilong," a/k/a "Zhang Jianguo," a/k/a "Atreexp," the defendants, conducted extensive campaigns of global intrusions into computer systems. The defendants worked for Huaying Haitai Science and Technology Development Company ("Huaying Haitai") in Tianjin, China, and acted in association with the Chinese Ministry of State Security's Tianjin State Security Bureau.

13. ZHANG SHILONG, a/k/a "Baobeilong, [illegible] Jianguo," a/k/a "Atreexp," the defendant, who worked for Huaying Haitai, registered malicious domains and hacking infrastructure used in connection with the APT10 Group's intrusion campaigns.

Indictments for Attribution: APT10

Government indictments, such as this one from the US DOJ on APT10, can serve the basis of more confidently assessing attribution at various roles. For example, here we not only see the group attribution, APT10, but also the name of the individual, their role in the operations, the company they worked for, the government agency they supported, and the state the government agency operated on behalf.

Even government attribution is not infallible. However, these indictments and similar documents from governments provide views given their intelligence capabilities. The private sector usually has much more depth and insight into cyber threat intelligence because they get to work the intrusions and incidents firsthand; but in the world of classic intelligence operations, governments around the world still dominate. When most Western governments perform attribution, it is often based on numerous types of collection such as SIGINT, intrusion analysis, and HUMINT.

I have seen indictments mess up technical details, but the overall themes are usually very accurate. These can support our own attribution efforts.

Indictments for TTP Discovery: APT10

The MSP Theft Campaign

In furtherance of the MSP Theft Campaign, Zhu, Zhang, and their co-conspirators in the APT10 Group engaged in the following criminal conduct:

- First, after the APT10 Group gained unauthorized access into the computers of an MSP, the APT10 Group installed multiple variants of malware on MSP computers around the world. To avoid antivirus detection, the malware was installed using malicious files that masqueraded as legitimate files associated with the victim computer's operating system. Such malware enabled members of the APT10 Group to monitor victims' computers remotely and steal user credentials.
- Second, after stealing administrative credentials from computers of an MSP, the APT10 Group used those stolen credentials to connect to other systems within an MSP and its clients' networks. This enabled the APT10 Group to move laterally through an MSP's network and its clients' networks and to compromise victim computers that were not yet infected with malware.
- Third, after identifying data of interest on a compromised computer and packaging it for exfiltration using encrypted archives, the APT10 Group used stolen credentials to move the data of an MSP client to one or more other compromised computers of the MSP or its other clients' networks before exfiltrating the data to other computers controlled by the APT10 Group.

- Throughout the indictments various TTPs are discussed that the indicted used
- Some are publicly known while some may be from unique govt collection

Indictments for TTP Discovery: APT10

Beyond attribution, the indictments and government documents can be very useful for uncovering TTPs and victimology. Most of these will be public, especially when referencing public groups such as APT10; again, the private sector companies have unique insight there. However, through government collection, there may be unique insights into TTPs or timelines given their collection as well. This is a rare opportunity to validate your findings and add to them.

Reference:

- https://medium.com/katies-five-cents/cyber-indictments-and-threat-intel-why-you-should-care-6336a14bb527

Indictments for IOC Discovery: APT10

- Sometimes, indictments including references to documents, filenames, malware, and other descriptive details that can be used directly for IOCs for historical sweeping and pivoting purposes

China under the control of the APT10 Group. That email, which was sent to employees of another victim company ("Victim-2") involved in helicopter manufacturing, had the subject line "C17 Antenna problems," a malicious Microsoft Word attachment named "12-204 Side Load Testing.doc," and stated the following: "Please see the attached the files." When the attachment named "12-204 Side Load Testing.doc" was opened, malware was installed on the computer of Victim-2. By using these spear phishing methods, the conspirators intended to and did cause the recipients of the emails to open the attachments without arousing suspicion as to the source of the email or its attachments.

Indictments for IOC Discovery: APT10

Many indictments will include specific references of emails, files, malware, communication that the adversary took part in, etc., that can be used to create IOCs. Using these IOCs as detections for future intrusions is very unlikely to be successful post indictment and, if anything, it presents an opportunity for other adversaries to spoof the indicted. However, for historical review and log discovery and system sweeping, these indicators can be very useful. A simply YARA rule could be useful in sweeping system images and old forensic artifacts from older cases to see if there are new pivot points and insights that can be revealed.

Cloud Hopper

- 2017 report by PWC UK and BAE Systems on activity they tracked starting in 2016; stated as a campaign highly correlated with APT10
- The focus of "Operation Cloud Hopper" was IT managed service providers (MSP)
- Targeting MSPs to get access to their clients
  - The MSPs were the victims but not the targets

- The report identified some key takeaways:
  - Victimology and ongoing heavy focus on MSPs
  - Shift toward bespoke malware and open-source tools
  - Largely dynamic-DNS domains and their assessment on an increasing ops tempo
  - Analysis of the "human fingerprints" of the malware to further the attribution

- It is a good example of a report to a broad community

Cloud Hopper

PWC UK and BAE Systems released a report on "Operation Could Hopper" in 2017 that was a great example of a well-done report.

There's a number of things about the report we'll highlight over the next few slides as examples to consider. Some of the best cyber threat intelligence reports are short, to the point, reports satisfying one and maybe at max two different intelligence requirements. This is so that the user can consume the information they need as quickly as possible. Words have a cost to your reader. However, when you're communicating to a broad audience, you don't know their intelligence requirements and largely have to make decisions based on your understanding of the wider community need. Therefore, the reports we see publicly are often longer with a lot more background information and "showing the work" that goes into the reports. You can easily make mistakes when doing this type of large public report, but the Operation Cloud Hopper one really stood up to scrutiny.

One of the things to note is that PWC UK doesn't truly know the definition of APT10. That's FireEye's group and though they surely coordinated and shared notes, the reality is PWC UK is not in a good position to say: "this activity is APT10." Fortunately, they didn't do that. Instead, they chose to say that the activity is "almost certainly" APT10 and chose to focus on their point of view and collection as a specific operation. It's unlikely every single intrusion they analyzed was APT10, but the broader picture and focus was delivered in such a way that the point and takeaways were the same.

Reference:

- https://www.pwc.co.uk/cyber-security/pdf/cloud-hopper-report-final-v4.pdf

Observations for CTI Analysts: Communicating Broadly

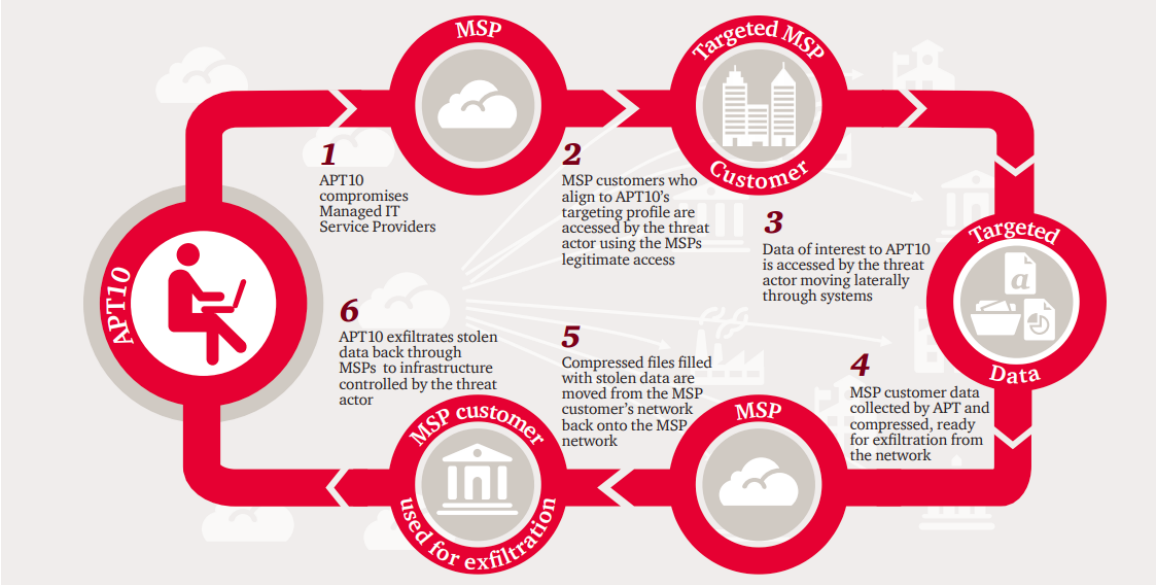

Observations for CTI Analysts: Communicating Broadly

While the report is speaking to a primarily technical audience, you do not need to be a malware reverse engineer to take away an understanding of what is going on here. The PWC UK and BAE team go to great lengths to break things down as simply as possible but also communicate through visuals. Here is a kill-chain-like view showing the efforts of the adversary against their victims and targets.

One graphic like this requires domains, malware samples, an understanding of victims, open-source intelligence (OSINT), exfiltration collection and analysis, and a deep understanding and clustering of seemingly disparate intrusions.

Observations for CTI Analysts: Human Fingerprints

- Choices by the APT10 actors such as work hours and workdays where domains are registered, and malware is compiled, is not a "smoking gun" but a good data point
- Registration details reveal behavioral choices by the individuals as well as a common name server that can connect various infrastructure

Table 3: Known APT10 registration details showing a common name server

Domain	Registrant email	Name Server	Contact Name	Contact Street
belowto[.]com	robertorivera@india.com	ns1.ititch.com	Roberto Rivera	904 Peck Street Manchester, NH 03103
ccfchrist[.]com	wenonatmcmurray@india.com	ns1.ititch.com	Wenona McMurray	824 Ocala Street Winter Park, FL 32789
cloud-maste[.]com	meganfdelgado@india.com	ns1.ititch.com	Megan Delgado	3328 Sigley Road Burlingame, KS 66413
poulsenv[.]com	abellonav.poulsen@yandex.com	ns1.ititch.com	Abellona Poulsen	2187 Findley Avenue Carrington, ND 58421
unhamj[.]com	juanitardunham@india.com	ns1.ititch.com	Juanita Dunham	745 Melody Lane Richmond, VA 23219
wthelpdesk[.]com	armandovalcala@india.com	ns1.ititch.com	Armando Alcala	608 Irish Lane Madison, WI 53718

Figure 7: Operational times of APT10 in UTC+8

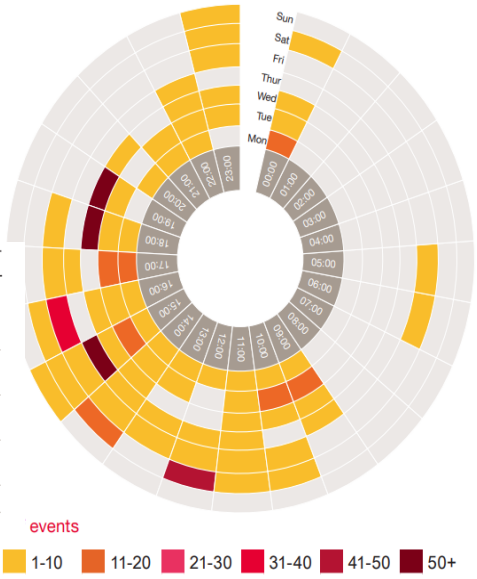

SANS DFIR

Observations for CTI Analysts: Human Fingerprints

The PWC team analyzed their intrusions to look for the "Human fingerprints" (choices and actions made by the adversary) such as when malware was compiled, when domains were registered, and when they could see interactive operations in their intrusions. Performing this analysis, they were able to determine that the times the adversary was most active was during work hours in the Chinese day and primarily during the workweek with activity very rarely ever taking place on Saturday and even less on Sunday.

Time zones, work hours, workdays, etc., can seem like an easily spoofed thing from the adversary. Why not just come in early or work weekends? Because people have lives. Even your adversary has day care, hobbies, and strives for some semblance of a work life balance. The idea of coming in on off days or switching to night shift just to screw with defenders in case they actually get detected is not something that shoots to the top of the adversary's concern when they are trying to get their mission accomplished. This isn't to say it doesn't happen, but this type of analysis is more useful than often perceived because of the human element of operations.

Another useful type of analysis that the PWC UK team did was looking at the APT10 registration details for domains and emails and there were a number of behavioral commonalities in name choices, contact streets, and registration emails. Additionally, they all shared a name server showing a solid pivot point to link these domains together. You should not only be looking for pivot points that can become obvious but the behavioral choices that may not be something a computer, pivot table, or algorithm is going to as easily pick up on as the analyst mind.

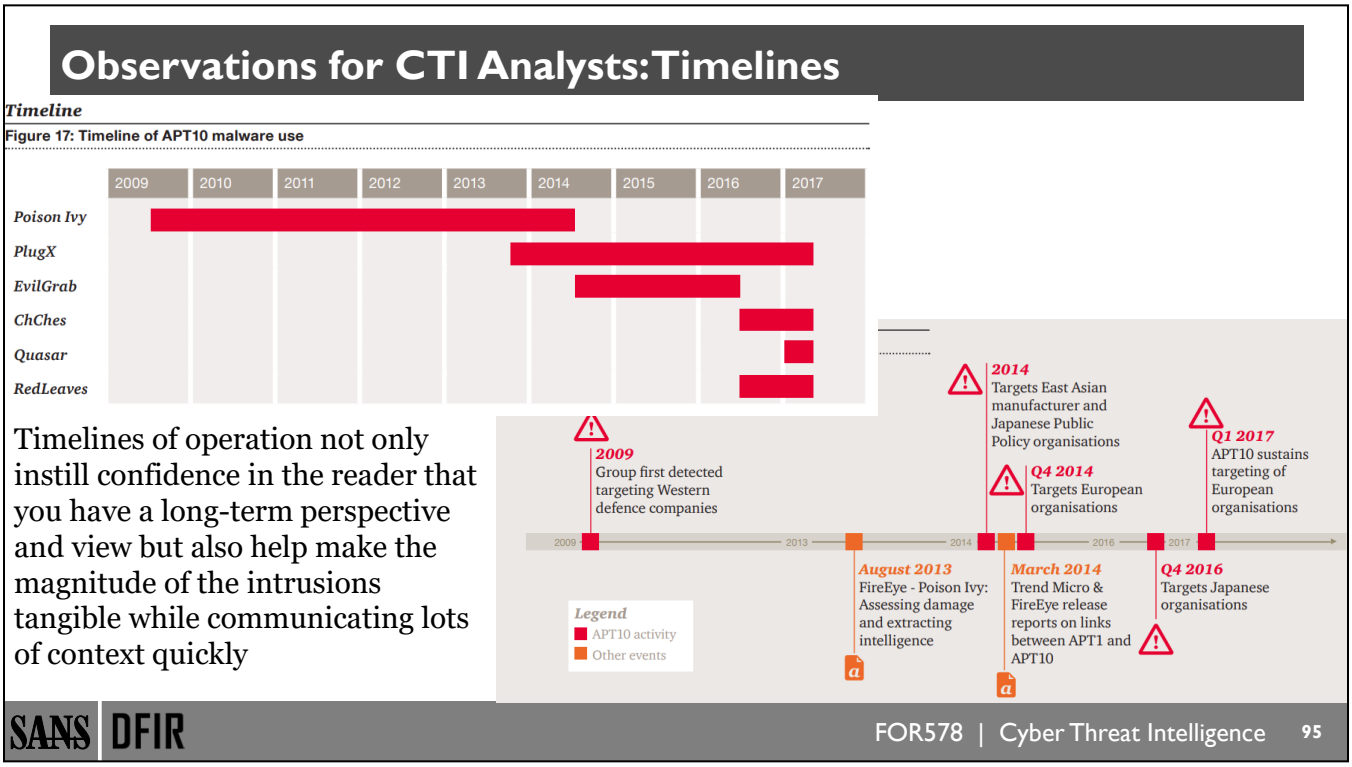

Observations for CTI Analysts: Timelines

Another great thing to do in CTI reports is to use timelines if you have the data to support them. As an example, the first image shows a timeline of APT10 malware usage. As you view this, you get an immediate perspective that this team has been tracking and collecting on this adversary for a long time. You also take away very quickly as a technical analyst the malware choices across time that may correlate with your own intrusions and observations helping to guide a lot of complex topics in a brief and abstract way.

Further, as we see the timeline of the APT10 operations in general, there's not only confidence in the tracking of this adversary but also credit given to the teams that have been involved in some of the major reporting, including FireEye and Trend Micro. That fosters collaboration but also further gives confidence to the reader that this view isn't some isolated claim, but a body of knowledge represented by a larger community of experts.

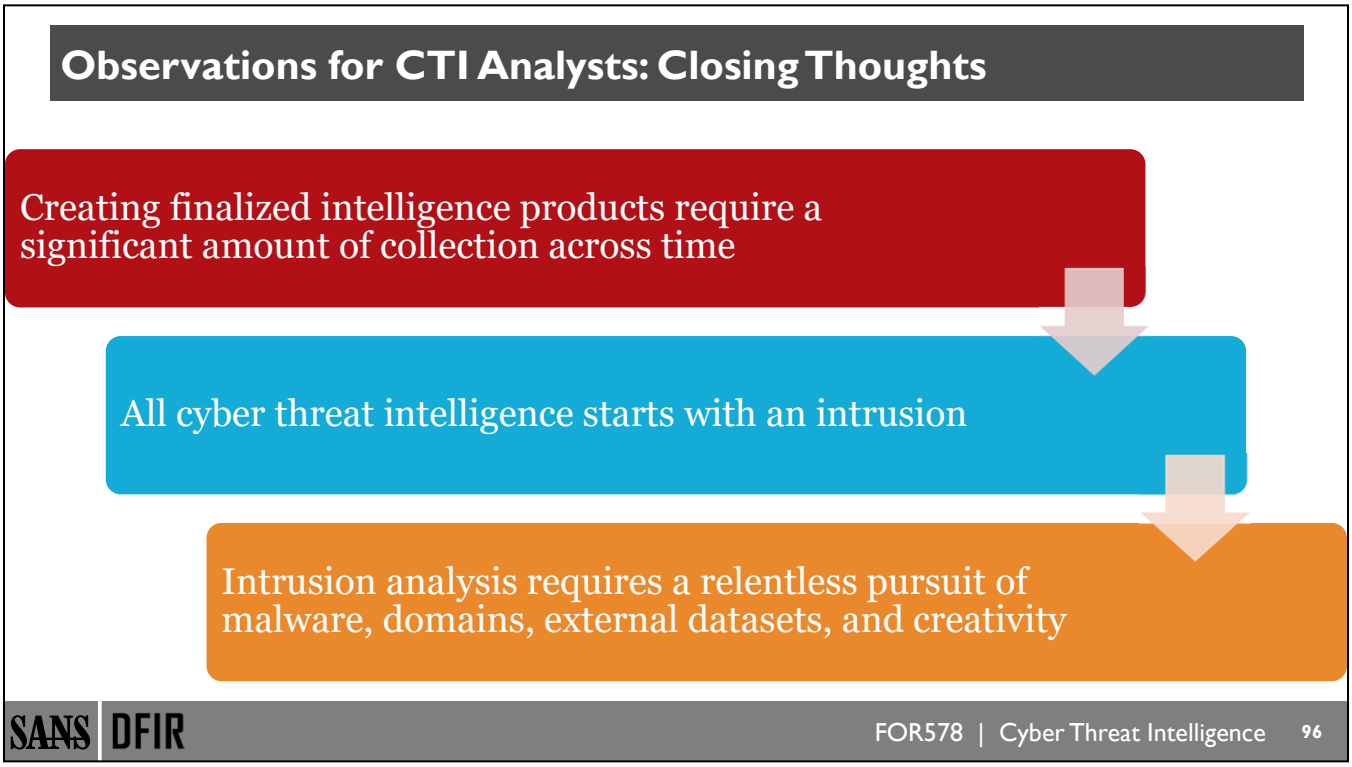

Observations for CTI Analysts: Closing Thoughts

Cloud Hopper was a report on a specific set of intrusions, clustered to reveal an operation of a specific group, APT10, tracked across more than a 10-year period. To achieve that type of report requires a significant focus and analysis on each intrusion that the analysts can get ahold of; a relentless collection of data. You cannot simply get to that type of finalized product with any level of reasonable confidence just through OPSEC errors by the adversary or a single really good incident response case. There's a lot of different ways to do CTI, but some of the best products come as the result of long-term analysis over disparate and unique collection sources.

A Specific Intelligence Requirement: Attribution

Because Sometimes It Matters

This page intentionally left blank.

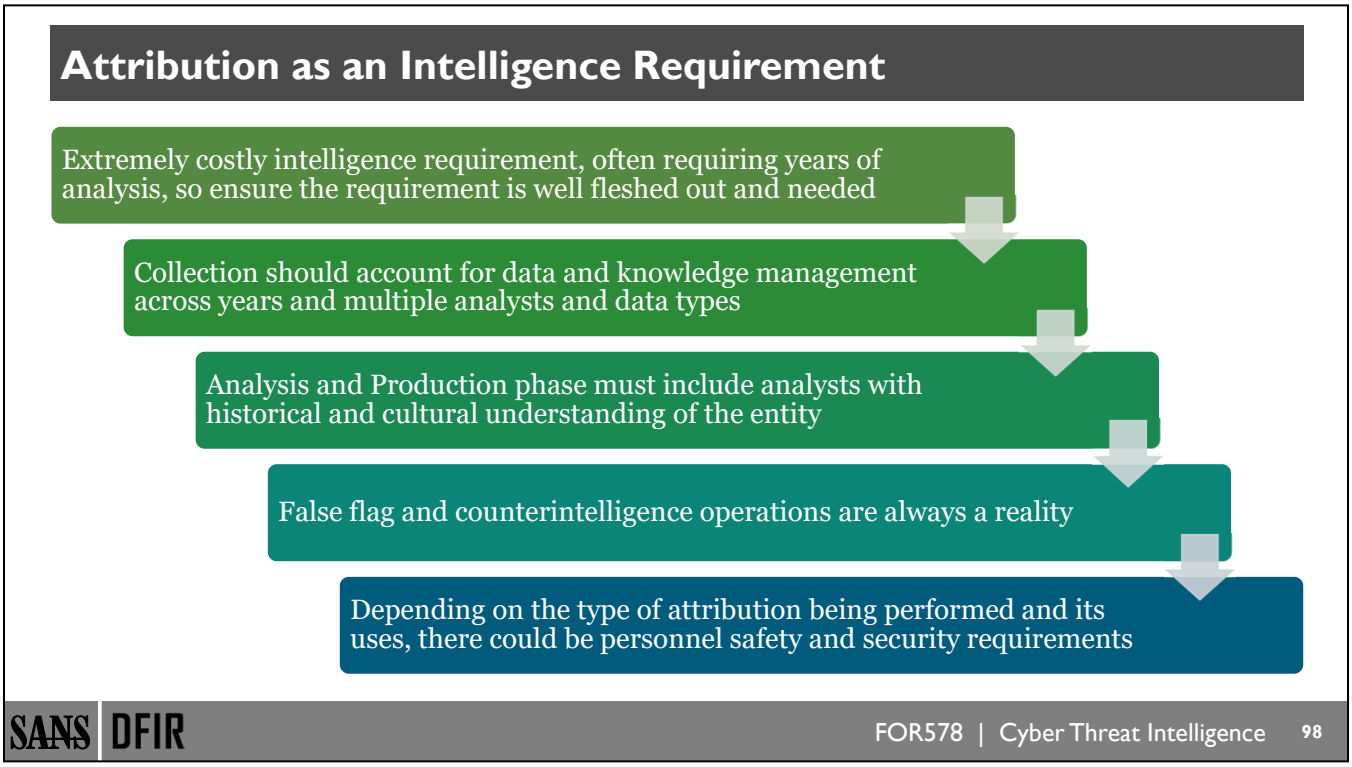

Attribution as an Intelligence Requirement

Attribution is an extremely costly intelligence requirement. It often takes years of analysis to do correctly or just tons of dumb luck; even when you think you have tons of dumb luck you cannot rule out risks such as false flag operations and counterintelligence efforts. It should constantly be reevaluated, and gaining confidence in attribution is a very difficult thing to do; the rule of thumb is you have no more than a "Low" confidence score in attribution if you are dealing with any one type of data (intrusion analysis as an example no matter how good really only gets you a low score, but adding in classic intelligence such as HUMINT or video surveillance can help achieve a Moderate and potentially High confidence score). Right now, the security community quickly applies moderate and high confidence scores inappropriately on attribution. Some firms are better than others, but the community is riddled with failed attribution that entities do not walk back.

Attribution as a requirement is going to drive a stringent collection program that needs years of data and structuring, so the storage and knowledge management of this data is key. Additionally, the analysis and production will need to include analysts deeply familiar with that state or entity you are trying to attribute.

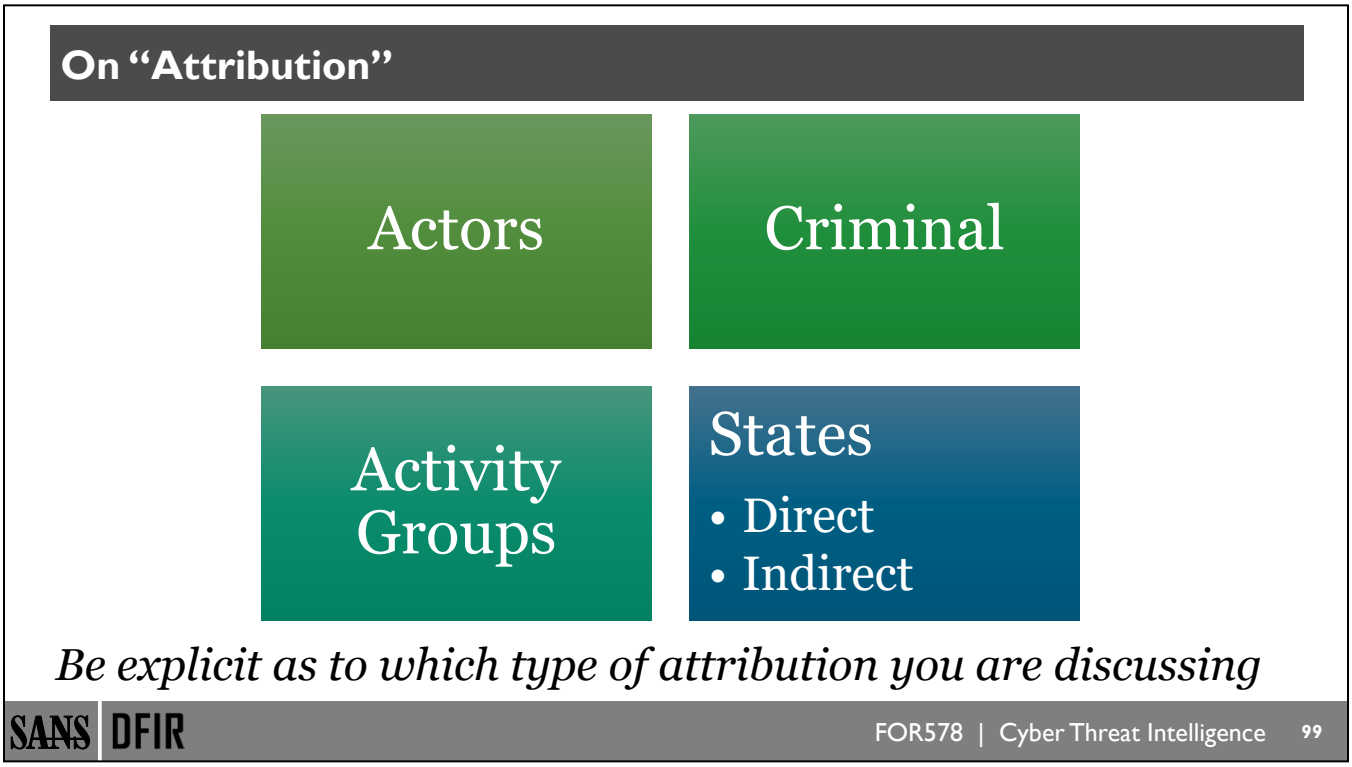

On "Attribution"

When we use the word "attribution," we may use it to refer to a number of things that aren't necessarily the same. We may use it to attribute some intrusion or activity to a:

- Activity Group, or a common set of shared attributes across an intrusion, a behavioral profile
- Actor, or the actual individual who was executing the actions that led to the observations made
- Criminal attribution to individuals that may be prosecuted
- Or a state in whose interests the actions were executed. This might be in the form of:
  - Direct attribution, where the actors were operating on the explicit instructions of a government, or
  - Indirect state attribution, where actors were operating in the specific interests of a government but perhaps without the explicit instructions of that government.

It's important to note that in the last case; in many cases, it may not matter whether the government was responsible for the activity, complicit, or willfully ignorant. Whether that matters depends on the ends supported by your analysis.

When speaking of attribution, it's always important to be explicit as to which type of attribution you mean.

I will use the phrase "true attribution" to talk about actually identifying the person or team behind the intrusion, but that is just a shorthand for purposes of the class and not meant to define a class of attribution more broadly in the community.

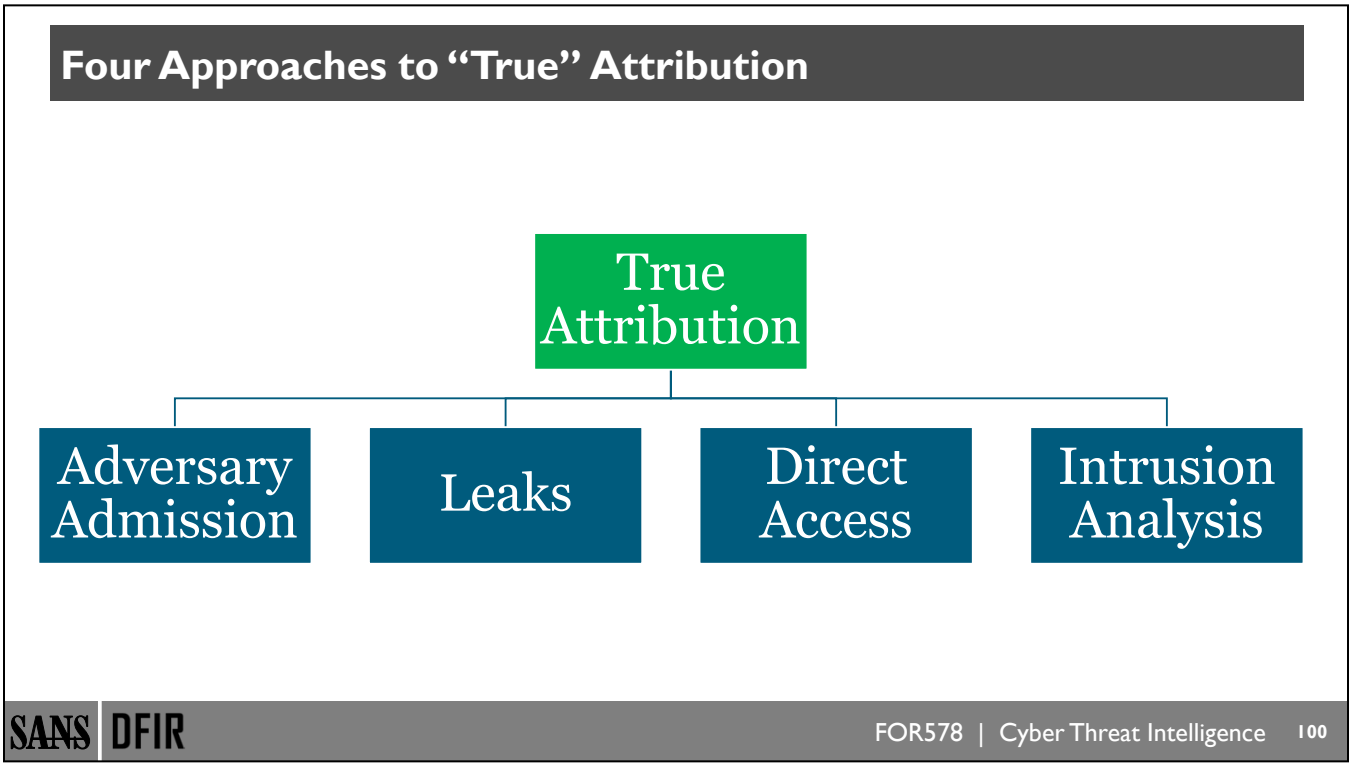

Four Approaches to Attribution

There are four common ways to achieve attribution. Two are required to do "true attribution" (the nation or group responsible) with a high level of confidence. The use of the word "true attribution" here is just meant to distinguish when people actually name states like China or operator names directly; it is not meant to imply that there is such a thing as false attribution. Technically, any type of clustering of intrusions to a group is a type of attribution, but in this context, attribution is being used in terms of revealing the actual identity of the operator and group.

The first type is Adversary Admission. This is when the adversary admits that they did the intrusions or campaign and take credit for their actions. As surprising as it sounds, this happens more often than people realize. As an example, there have been many Israeli-, US-, and UK-based individuals and national leaders that have openly taken credit for successful operations they'd run even if it was years later.

The second type is Leaks, which also cover operational security issues. In these cases, the adversary has substantially released information or had it released about them. The Edward Snowden leaks are a good example of leaks that helped identify various US-based operations. There were also leaks that attributed the French government to various operations they ran known as "Animal Farm."

The third type is Direct Access. This is when someone directly interacts with the adversary or their systems or collects on them, such as classical SIGINT and HUMINT. This is classic spy-type work where someone overhears someone talking about the operations, where they collect their telephone calls about them doing it, or even break into their systems to see the intrusions happen firsthand.

The fourth type is Intrusion Analysis. This activity covers intrusion analysis, as well as the campaigns and groups that orchestrate the intrusions. This is the role of most organizations in identifying intrusions into their environment and patterning out the activity.

Any one of these methods can be highly misleading. Intrusion analysis can be subject to our field of view and biases on how we interpret the data, direct access might overhear something and misunderstand it or collect information that is planted by the adversary, leaks might include information that was wrong or fake information, and adversaries have, in the past, taken credit for operations they did not run for various political or accidental reasons. Having to facilitate better attribution and in the field of CTI, so far, we've seen the most convincing attribution combine Intrusion Analysis with one of the other three types.

The Simpsons Did It

- "The Simpsons Did It" is a pop culture reference that jokes that everything that someone wants to do, especially a plot for a TV show, the Simpsons already did
- Apply "The Simpsons Rule of Attribution" when determining if your use case justifies doing real attribution

- If the attribution determined that the adversary was the Simpsons, would that change the actions you take?
  - If not, then you do not need true attribution
  - Or said another way; if you can't action attribution, don't perform the costly effort

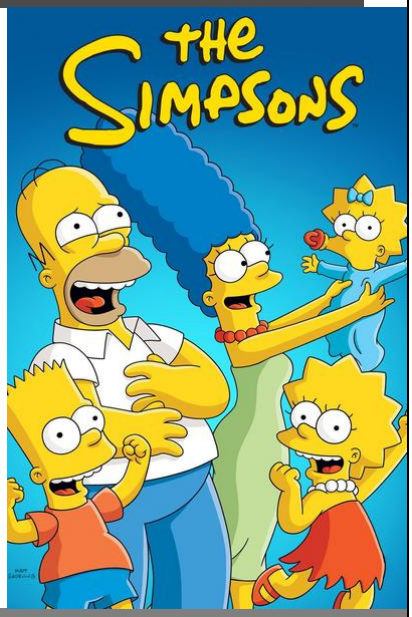

The Simpsons Did It

There are some use cases for attribution but there are even more use cases that people claim need attribution that do not. Attribution is a particularly challenging intelligence requirement that can exhaust analyst resources and thus be an expensive endeavor that jeopardizes the ability to satisfy other intelligence requirements. Thus, it is important to be highly critical of whether or not attribution is really needed for the goal of the intelligence requirement.

A common way to think of this is the common pop culture reference "the Simpsons did it." This reference generally states that there's very little on TV shows and sometimes in real life that the Simpsons haven't already done. From predicting Super Bowl playoffs to having the exact same plot concepts that other TV shows tried; a popular example was South Park, which had the episode title "Simpsons Already Did It" https://en.wikipedia.org/wiki/Simpsons_Already_Did_It

For the use in an attribution construct, imagine that the culprit is in fact the Simpsons; after all, they are commonly accused of having done "it." If your actions as a defender do not change based on whether or not the culprit is in fact the Simpsons, then you do not need attribution.

Achieving the Value of Attribution without Attribution

- A common use case of attribution is the prioritization of defenses
  - Geopolitical unrest between states could hypothetically justify a priority shift
  - This can be obtained without actually doing real attribution
- Define an analytical model of what that entity or state represents to you
  - Define motivations (intent), their opportunities, and capabilities
    - If you cannot do this, then even having true attribution would be useless to you
  - Compare intrusions to the model and do not worry about the true attribution
    - You would note that your cluster operates "in the interests of" the state whether or not it's them and prioritize them as needed

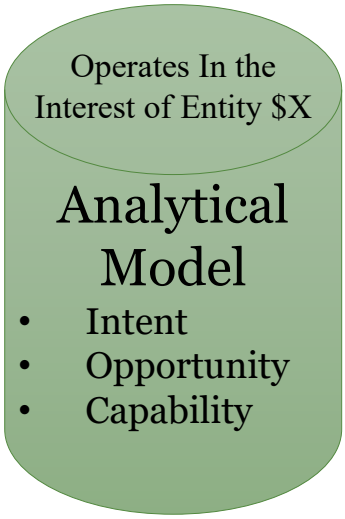

Achieving the Value of Attribution without Attribution

A common argument is that attribution is required for the prioritization of defenses. As an example, if you are a United Kingdom-based electric utility, and the UK and North Korea have strained relationships, then you might try to prioritize activity groups and threat actors that have been attributed to North Korea. The logic seems sound, but there are numerous problems in this. For example, what has been attributed as North Korea before could be highly inaccurate, and you likely have little visibility into the analytical model that went into that attribution from the originating source. Additionally, state entities are diverse and have a variety of intelligence requirements; prioritizing "North Korea" may not lead to actually reducing risk against North Korean-based teams at all, especially if there is a change in targeting and tradecraft by those entities. Further, North Korea may be leveraging allies, and thus the mixing of tradecraft from multiple groups or actors could confuse the prioritization. The chief problem, though, is that there's a more efficient way to get the same value without having to attribute North Korea for sure.

For example, if you define an analytical model on what it means to be operating *in the interests* of the North Korean state, you might include various types of geopolitical motivations, types of opportunities that the state would have to launch operations, and maybe even some aspects of the anticipated tradecraft that those actors might use. To prioritize "North Korean-based actors," you'd have to understand an analytical model anyway. So, define the analytical model and instead define it as "operating in the interests of the North Korean state." It's the same amount of base effort without the rigor needed to do attribution completely. That is, after defining the analytical model, you can use it to compare to any other intrusions and activity and prioritize the ones that fit the model. If the model is done correctly, then it does not matter if the entity responsible is North Korea or Syria, so long as the model is defined and consistent to allow prioritization in different events. As an example, if the group or actor that aligns with the model has an uptick in activity when there is geopolitical unrest between the UK and North Korea, then it doesn't matter who they actually are; for all it matters, it could be the Simpsons.

The alternative is focusing on attribution, which acts as a crutch instead of a structured defined analytical model, meaning that analysts' bias will quickly get in the way of using that information regardless of its accuracy or not.

Example Use Cases of Attribution

- Criminal Prosecution
  - Some entities have a need for criminal prosecution or supporting these actions
  - Financial organizations, as an example, commonly support international law enforcement on this
- State Tools (Diplomatic, Information, Military, or Economic)
  - States have a variety of tools at their disposal, from sanctions to cruise missiles; in these cases, getting true attribution is particularly important
- Strategic Organization Changes
  - Interactions with other entities, including how international business or travel is conducted

Example Use Cases of Attribution

Some entities can leverage attribution. As an example, law enforcement agencies (LEAs) try to prosecute cyber criminals and rogue members of states. There are arguments over the effectiveness of this strategy for deterrence, but this is a completely legitimate use case. Entities can contribute their intrusion analysis to an LEA in concert with the LEA's sources and methods to help provide true attribution that can lead to prosecution.

States also have the need for true attribution, especially as it relates to utilizing state tools such as sanctions and trade agreements. True attribution can help those states leverage the tools at their disposal while also signaling to the aggressors what will and will not be allowed.

Organizations can strategically use attribution as well. For instance, many executives no longer travel to China with their personal devices because they fear being compromised and having intellectual property stolen. This should probably be a rule of thumb for every country, but understanding attribution can help drive cultural change around business dealings internationally or how travel is conducted. Additionally, it may matter for acquisitions and investments as well in foreign companies and the type of business conducted.

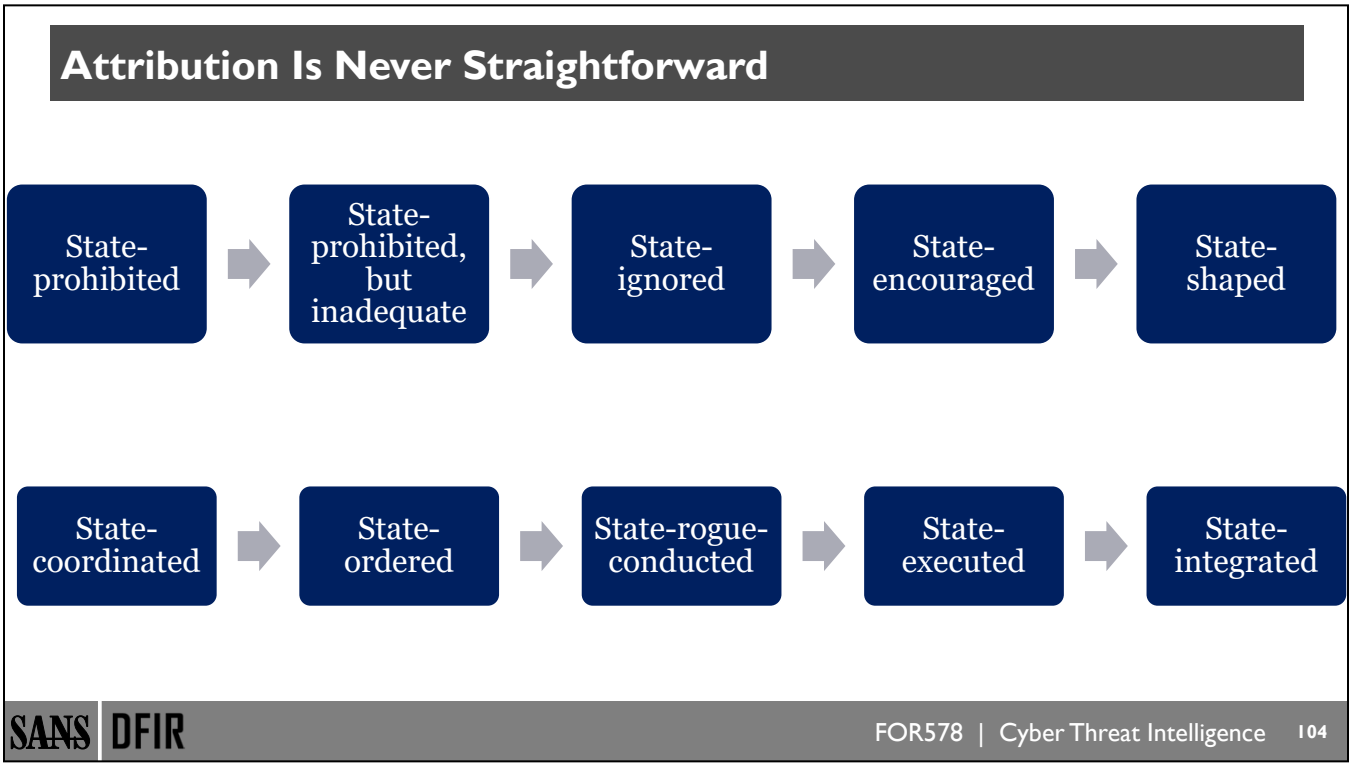

Attribution Is Never Straightforward

Jason Healey put together an exceptional piece at the Atlantic Council on attribution. His effort was to get there to be some responsibility in the digital domain regardless of "attribution," i.e., a state could be culpable for an attack even if they did not push enter on the keyboard.

Jason's scale is useful to understand as we think about attribution so that we do not fall prey to "attribution fixation," where, because we cannot get to definitive attribution, we do not hold people or states responsible for their actions.

The Spectrum of State Responsibility 1. State-prohibited. The national government will help stop the third-party attack. 2. State-prohibited-but-inadequate. The national government is cooperative but unable to stop the third-party attack. 3. State-ignored. The national government knows about the third-party attacks but is unwilling to take any official action. 4. State-encouraged. Third parties control and conduct the attack, but the national government encourages them as a matter of policy. 5. State-shaped. Third parties control and conduct the attack, but the state provides some support. 6. State-coordinated. The national government coordinates third-party attackers such as by "suggesting" operational details. 7. State-ordered. The national government directs third-party proxies to conduct the attack on its behalf. 8. State-rogue-conducted. Out-of-control elements of cyber forces of the national government conduct the attack. 9. State-executed. The national government conducts the attack using cyber forces under their direct control. 10. State-integrated. The national government attacks using integrated third-party proxies and government cyber forces.

Reference:

- https://www.fbiic.gov/public/2012/mar/National_Responsibility_for_CyberAttacks,_2012.pdf

Example: Merged State and Criminal Activity

- CrowdStrike identified "Pioneer Kitten" which is assessed to be an Iranian state-based group
- However, the team has also specifically been seen contracting their operations to others and have compromised organizations to sell access to criminals on underground forums
- Most state teams interact with contractors, development shops (internal or external), infrastructure teams, operators, analysis teams, etc.

Example: Merged State and Criminal Activity

Most analysts understand and accept that there are criminal and state actors and sometimes the state actors leverage resources from the criminals. But, in reality, it also works the other way around as well. In this example provided by CrowdStrike, it was an Iranian state sponsored group that was targeting systems around the world to monetize them by selling them to criminals on underground forums. CrowdStrike assessed that PIONEER KITTEN themselves may just be a contractor. And that may be true in this case as more information is presented; however, this type of activity has been observed by Russian, North Korean, Chinese, and other teams as well where government operators will moonlight (work after hours and use their same capabilities and infrastructure) or even sell access and the byproducts of their operations with or without government approval.

The reality is the threat landscape is an incredibly volatile and mixed landscape with a variety of different requirements and motivations from each team that operates there. Simply stating that "this state likes these types of operations" or "this actor does these things" betrays the complex web of human and geopolitical interactions that can take place.

References:

- https://www.zdnet.com/article/iranian-hackers-are-selling-access-to-compromised-companies-on-an-underground-forum/
- https://www.crowdstrike.com/blog/who-is-pioneer-kitten/

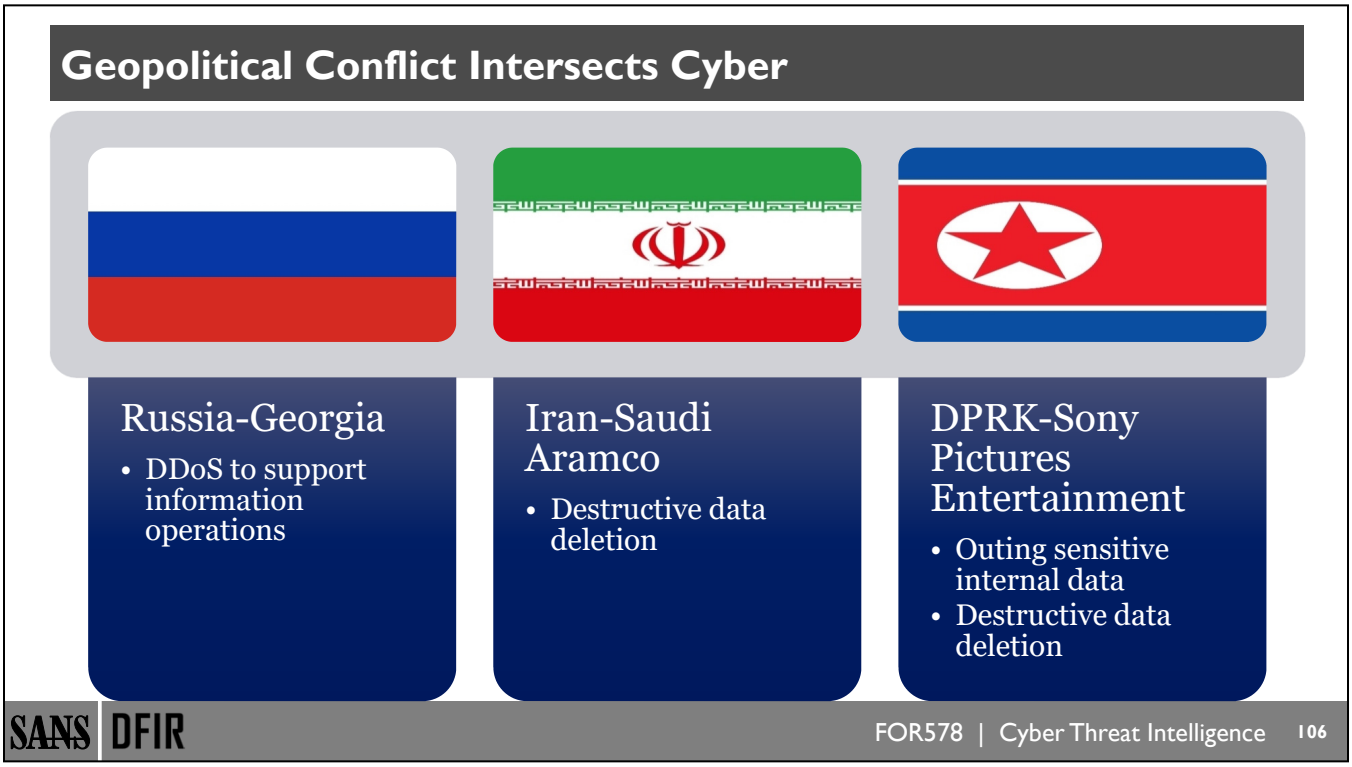

Geopolitical Conflict Intersects Cyber

The geopolitical climate can be a strong motivation for countries to use computer network operations to collect intelligence or influence through information operations. Some examples of this include the Russia-Georgia, Iran-Saudi Aramco, and DPRK-Sony events that have played out on the world stage. There have obviously been others as well, but let's take a few moments to discuss some of the finer points of these incidents as they related to CTI.

In 2008, many suspected the Russian government of conducting cyber attacks via distributed denial of service (DDoS) against Georgian government and media websites, effectively "blacking" them out from normal use. This affected the capability of the government and media to disseminate relevant information to the Georgian citizens about the conflict with Russia (https://www.researchgate.net/publication/230898147_The_2008_Russian_Cyber-Campaign_Against_Georgia).

In a more destructive attack that occurred at Saudi Aramco, a major global oil and gas company, more than 30,000 Windows computers were destroyed by a malware identified by security vendors as Shamoon. Although never officially attributed to the Iranian government, it certainly had the motivation and apparent capability to pull off a somewhat sophisticated attack. (Bronk, Chris and Tikk-Ringas, Eneken, "Hack or Attack? Shamoon and the Evolution of Cyber Conflict" (Feb 01, 2013). Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2270860.

Finally, when the DPRK attacked Sony Pictures Entertainment (SPE) in November 2014, probably one of the most destructive computer network exploitation and attack incidents was highlighted globally. Many didn't suspect that North Korea could pull off such a highly successful compromise and destructive actions. According to various reports, the attackers collected more than 100GB of data from SPE and probably surveyed its entire network. It then created custom destruction tools to target servers and workstations with the SPE network to completely knock SPE off the internet. It was reported that the primary motivation for this attack was to stop the release of a new motion picture that portrayed North Korea in a disrespectful way because it involved killing an actor playing the role of Kim Jong-Un. Prior to the intrusion and destruction, North Korea has publicly released statements condemning the film and threatened retaliation.

As a CTI analyst, it is important to understand how the geopolitical landscape plays into each of these intrusions. Unless you have a background in the culture of a particular country, you probably haven't spent much time understanding each country with a militarized cyber capability and understanding how, when, and to what extent they will use it to both collect intelligence and perform destructive actions.

References:

James Cook (December 16, 2014). "Sony Hackers Have Over 100 Terabytes Of Documents. Only Released 200 Gigabytes So Far" *Business Insider*. Retrieved December 18, 2014.

Ben Child. "Hackers demand Sony cancel release of Kim Jong-un-baiting comedy" *The Guardian*. 9 December 2014.

"Flag of Russia." Licensed under Public Domain via Wikipedia: https://en.wikipedia.org/wiki/File:Flag_of_Russia.svg#/media/File:Flag_of_Russia.svg

"Flag of Iran": English translation / interpretation at https://www.fotw.info/flags/ir'.html. Licensed under Public

Domain via Wikimedia Commons: https://commons.wikimedia.org/wiki/File:Flag_of_Iran.svg#/media/File:Flag_of_Iran.svg

"Flag of North Korea" by Zscout370 - Template:조선 민주주의 인민 공화국. Licensed under Public Domain via Wikimedia Commons: https://commons.wikimedia.org/wiki/File:Flag_of_North_Korea.svg#/media/File:Flag_of_North_Korea.svg

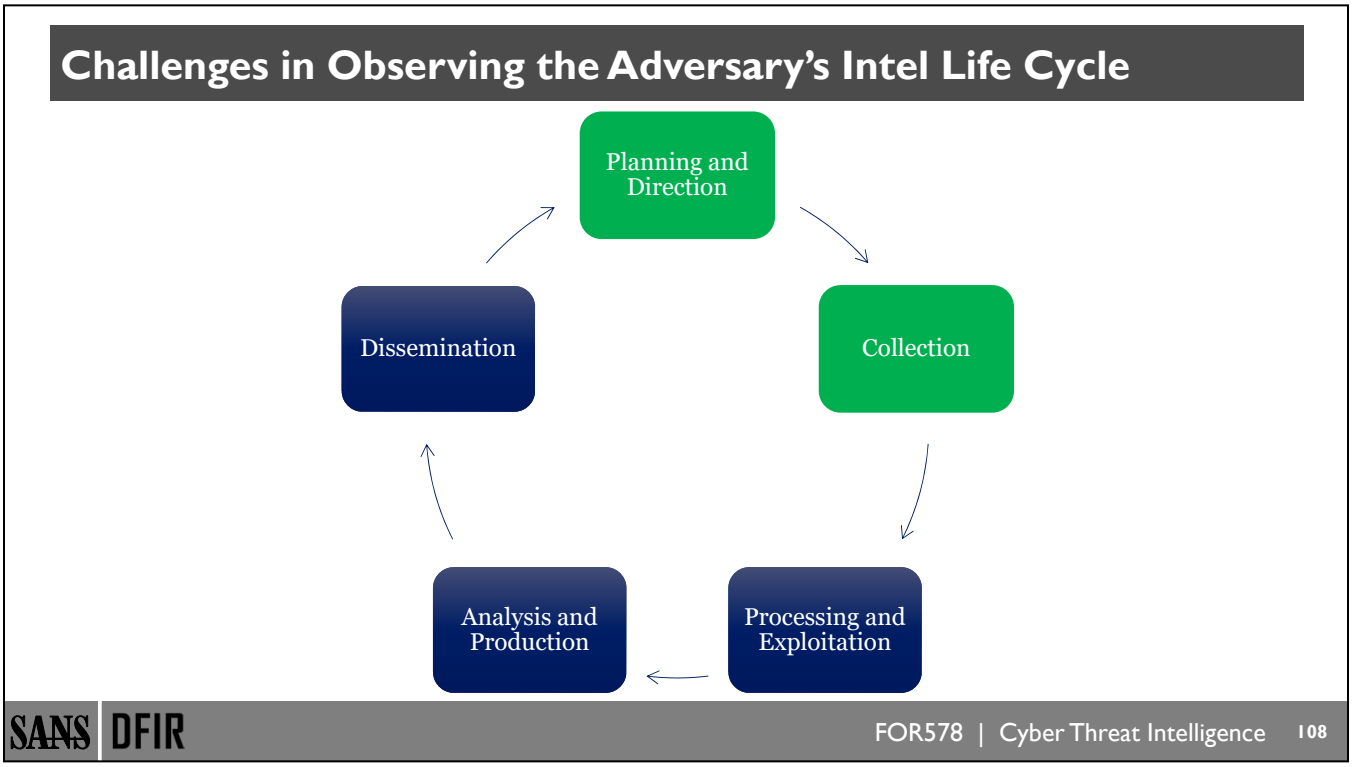

Challenges in Observing the Adversary's Intel Life Cycle

In analyzing the adversary conducting the campaign, let's borrow the intelligence cycle slide from Day 1. When we introduced it, we were focusing on how we as intelligence analysts go about our business. Let's now flip the perspective and look at how the adversary uses a similar intelligence cycle. Computer network exploitation is intelligence collection, but the other phases of the intelligence cycle play their parts within this process.

For the adversary aspect and looking at building out our campaigns or intrusion sets, we need to think about both who is tasking and who is collecting. The respective elements responsible for each phase may be contained within a single element or perhaps a specialized unit or freelance group that is collecting at the request of another organization. The likelihood is that, in many cases, we might gain attributional information about the collectors or operators, but rarely the taskers/consumers. Unfortunately, depending on their organization, trying to determine what an adversary group targets may be difficult without the visibility into the taskers/consumers.

Let's dive deeper into this topic.

Deriving Intent

- Impact != intent
  - We see "perceived" intent, not actual intent or motivation
- Effects suggesting attack may be unintended
- Must understand adversary objectives
  - Respond appropriately
  - Plan appropriate defenses for future incidents
  - LEA/.mil response

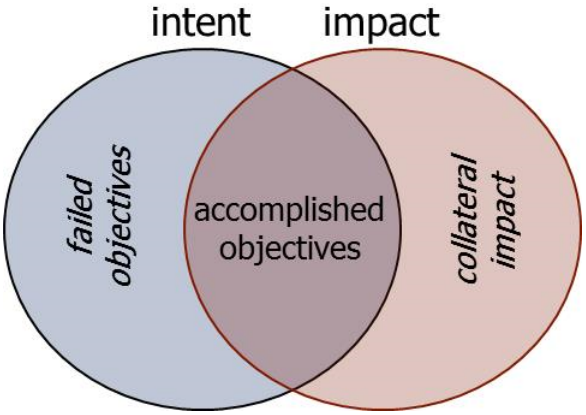

Deriving Intent

Intent and impact are not always the same thing. This is an important point to always consider when trying to infer intent from observed impact. For example, unintended modification of data, or incidental compromises of availability, in the course of exploitation may be said to have had the effect, or impact, of what we normally associate with attack. However, if that were not the intent of the adversary, the incident should be treated as though it were espionage, NOT sabotage. Although this may not make a difference to those responding to the incident, the distinction has significant implications for the organization's strategy to mitigate future actions of this type. It also can have profound implications for law enforcement or national policymakers.

Here is a longer discussion on adversary intent and the challenges with it:

https://www.youtube.com/watch?v=gtLAgvCo8TA&t

The Basics of State Attribution

- Understand the state deeply
  - History, Culture, Language, etc.
- Define Models for Each Entity
- Follow ACH steps
- Classify evidence based on threat definition:
  - Intent
  - Opportunity
  - Capability
- Confidence in assessment informed by support in each group of evidence

The Basics of State Attribution
Just like we discussed with intrusion to campaign attribution, it is important to identify the basics for doing state campaign attribution. In essence, it is the same ACH process, but there is also the element of a true threat—those adversaries possessing the intent, opportunity, and capability to do harm. Think back to the Stuxnet example: Not every country had the intent to do Iran harm, not many had the opportunity, and very few had the technical capability. These helped determine national candidates for attribution.

Linguists in Cyber Threat Intelligence
One significant skill that cannot be understated is the need for a mature and robust CTI program to have capable linguists at its disposal. These can be internal or external outsourced resources for those less established organizations. In looking at the previous APT1 report, for example, a number of the data points that served as evidence to support their assessment that APT1 is conducted by the Chinese military were translated from posts in Mandarin Chinese. In such a technology-focused arena as CTI, the linguist is an often-overlooked skill set for many organizations. However, we will spend some time covering additional uses for linguists that you might not realize that can help shape a CTI program in tangible ways.

Tactical Linguists
Linguists can provide value in other ways to a CTI program as well. Many adversaries reside in foreign countries and in many cases, often their primary language is not the same as spoken by your CTI analysts. Although tools such as Google Translate, Babelfish, and other commercial bulk translation products can aid in trying to triage foreign language content, because of the underlying hacker culture and "leet speak," foreign adversaries have their own unique cyber lexicon. A well-seasoned linguist can bring clarity to conversations that, otherwise, would easily be fouled-up and mistaken for garbage output from a machine translation. A common term among long-time CTI analysts that provides a good example is the Chinese pinyin word "rouji." This word translates literally to "meat chicken" in English and doesn't make much sense. However, in Chinese hacking culture, this word actually refers to zombie computers, bots, or, more generically, "hop points" or operational relay boxes (ORBs).

For another consideration on performing attribution, read Thomas Rid and Ben Buchanan's paper, "Attributing Cyber Attacks" and their Q Model: https://ridt.co/d/rid-buchanan-attributing-cyber-attacks.pdf

Analytical Model for Each Entity

- Define your analytical model for each entity you want to consider for attribution
  - Doing this haphazardly will lead to high levels of bias and uncertainty
    - E.g., you do not "know" what China wants out of operations, but you can go through an assessment to create an analytical model that you then leverage for analysis in attribution
  - Analytical models combined with Leaks, Direct, or Admission can yield true attribution

Analytical Model for Each Entity

Simply trying to guess what an entity wants is neither consistent nor professional. You might as well roll attribution dice at that point. However, if you properly define and justify an analytical model, then you can use that as a basis for what it means to operate in the interests of a specific entity. From there, tracking that data over long periods of time and then looking for OPSEC issues/leaks, adversary admission, or direct intelligence can inform the intrusion analysis and analytical model to help justify whether or not the analytical model is simply "in the interests of" or actually true attribution.

In essence, do intrusion analysis and then complement it with other intelligence and insights that were highlighted in the four types of attribution. Justifying attribution through two different categories of attribution analysis can help reach a moderate to high confidence assessment on attribution.

Or said bluntly: If you do not create the model, you're going to view the intrusion through ahead of time, you will naturally become extremely biased when you try to create a model to assist you on the intrusion you know have familiarity with. I.e., don't find an intrusion and then think "why might this be China." Instead, think "here's what would make me associate something with China" and then find the intrusion and see if it fits. Build the lens before applying it to help minimize your biases.

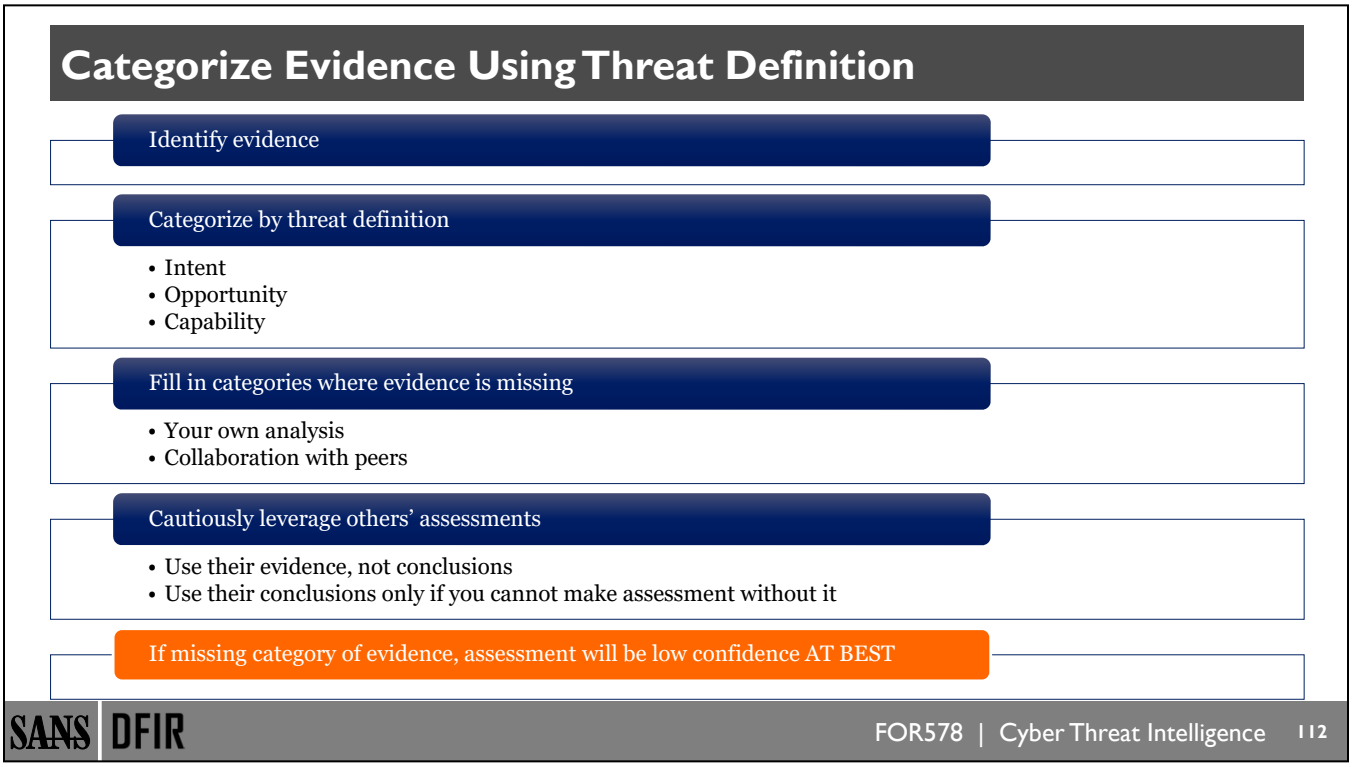

Categorize Evidence Using Threat Definition

When we seek evidence for this exercise, we want to identify what we have. To enumerate our gaps, for which we will seek additional intelligence (through our own analysis or collaboration with peers), we leverage the definition of threat: intent, capability, and opportunity. Categorize your evidence as falling into one of these three factors.

If your evidence doesn't fall into one of these three pieces of evidence, often it's because you're considering as evidence the assessment of a third party. Where possible, read into the detail of **why** that third party made attribution the way it did, and extract that evidence, not its conclusion. If such evidence does not exist, carefully consider the source. It might be beneficial to attempt to make an assessment without the assessments of others first and include only others' assessments if you cannot do so yourself without including it.

If you have one category of evidence that is completely empty, **this is a red light that you probably have insufficient evidence to form a hypothesis**. At best, hypotheses with one category of evidence completely missing will be made with **low confidence**.

State Hypotheses Enumeration Tips

When formulating your hypotheses, you may want to consider multiple states, OR you may simply want to hypothesize that the attribution is "some other entity." Be mindful that, in the end, your analysis may be inconclusive!

Also, and this is important to remember: ACH emphasizes that we try to REJECT hypotheses, not support them. Typically, analysts will be convinced going into an analysis like this that the outcome will be one country or another. Although this is natural, we must work against this tendency, which results in anchoring and confirmation bias!

Understanding Opportunity

- Availability of means to accomplish objectives
- Related but different from vulnerability
- May be:
  - Technical
  - Political
  - Logistical
- Countermeasures mapping to capabilities reduce opportunity to actualize intent

Technical examples
- Email systems
- Zero-day with no patch
- Private registrars
- Access to protected network

Political examples
- Legal authority
- Willful LE inaction
- Failed states

Logistical examples
- Delayed CIRT action
- Organizations merger

Understanding Opportunity

Opportunity is the availability of the means in which an adversary can accomplish her objectives. It's easy to confuse this with vulnerability: Don't! They're two different but related elements that feed into risk. Opportunity can be of a technical, political, or logistical nature. Examples of each follow. As you go through them, you can see how opportunity must align with intent and capability for operations to take place; these are three interdependent elements of threat and why we define it in those terms!

As a corollary, *countermeasures* put in place by defenders, informed by adversarial *capabilities* reduce the *opportunity* for them to convert on their *intent*. This is another example of how CTI facilitates network defense against persistent adversaries.

Following are some examples of technical opportunity:

- Email that sends attachments to users on a target network presents an *opportunity* to deliver malicious code.
- A zero-day with no patch released presents an *opportunity* to exploit client systems.
- Private domain registrars present an *opportunity* to register many domains without discovery by potential targets.

Following are some examples of political opportunity:

- Authority under law to execute operations (that is, constitutional title in the U.S.)
- Willful inattention or application of laws prohibiting operations
- Ineffective governance

Following are some examples of logistical opportunities:

- Accomplishing objectives before security apparatus can react
- Timing of operations exploits work hours of the victim organization
- Merger between two organizations, one of which the adversary has already penetrated

ACH Matrix Template for State Attribution

Category	Evidence	Country 1	Country 2	Other Entity
Intent	E1	+		+
	E2		+	
	E3	-	+	-
Opportunity	E4	+	+	+
	E5			
	E6	-	-	-
Capability	E7	+	+	
	E8	++		
	E9		+	+

When refined, all evidence for opportunity will be removed:

- Seek evidence to fill intel gap
- Determine remaining evidence sufficient for low-confidence assessment
- Determine evidence is inconclusive

ACH Matrix Template for State Attribution

Here is an example of the layout of the ACH matrix after the evidence has been considered against each hypothetical state attribution. Notice how the evidence is grouped by category.

In this generic example, after we refine the matrix, we see that there is no diagnostic evidence remaining in the Opportunity category, as each piece of evidence weighs the same for each hypothesis. At this point, we have a few options:

- Seek additional evidence to fill this intelligence gap and reassess against the hypothesis.
- Continue if the remaining evidence is strong enough to formulate a low-confidence assessment.
- Stop and consider the evidence inconclusive.

Be Prepared for Information to Change

- Understand and document the key evidence that went into the intelligence requirements you satisfied
  - Over time sometimes key evidence or our understanding of it changes
- If you are relying on leaked information, especially for strategic intelligence, be wary of active measures and plan for the worst-case scenario

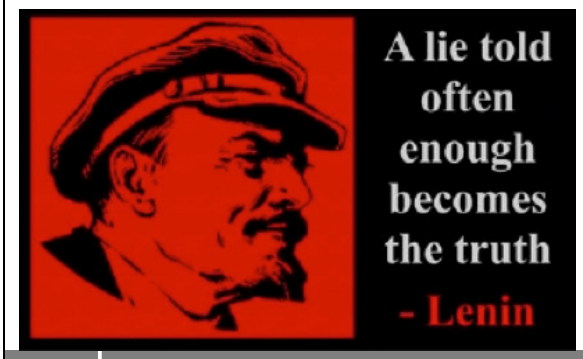

Active measures are semi-covert or covert intelligence operations to shape an adversary's political decisions

- They almost always conceal or falsify the source (anonymity or false flags)
- They also spread forged or partly forged content

Be Prepared for Information to Change

Information can change, as can our understanding of it. Being prepared for information to change is not just an ACH process, but instead an understanding of intelligence. Likewise, a customer may feel the intelligence requirement is satisfied and then determine later on that they did not get all the information they needed. Document all completed intelligence requirements and the key questions, challenges, and information required to satisfy the intelligence requirement. These will not only serve as lessons learned to make your process better but allow you to adapt to changes over time.

If you are relying on leaked information (which tends to be common these days), be aware that it is very common for leaked information to contain fake information as well. The Soviet-era disinformation campaigns routinely would leak information that were lies; however, they also routinely leaked real information and changed only 10–15% of it, knowing that the rest that was factually true could help embolden the lies. Information changes—have a plan to adapt to that fact.

Case-Study: Soviet Disinformation Operations

- In the mid-1960s, Russia's intelligence services pioneered *dezinformatsiya* (active measures), particularly through the KGB and the Stasi's HVA
  - The Cold War saw more than 10,000 individual operations

U.S. invented HIV as a bioweapon	President Carter ordered operations to disrupt black organizations	Threats to African athletes in the 1984 Olympics
KGB started the story in 1985	Soviet-forged document leaks to San Francisco newspaper 18 September 1980	Soviets forged two leaflets from Ku Klux Klan leaders threatening the lives of African athletes
Three "French doctors" published the "research" blaming U.S. military labs	Soviet news agency TASS distributed it in different languages, pushing White House to publicly protest it	TASS released them and KGB planted the stories around the world to gain protest to US-held Olympics in response to US protests of 1980 Olympics

Case-Study: Soviet Disinformation Operations

Disinformation campaigns were pioneered by Russian intelligence services. During the Cold War, more than 10,000 operations were run, ranging from the U.S. inventing HIV to death threats by the KKK on African athletes during the Olympics to encourage protests of the U.S. Olympics. Each time, the Soviets faked the source of the information and then used state-run media to push out the stories into multiple languages, where operatives around the world were able to work it into local newspapers and media coverage. It was common for these campaigns to be focused as a tit-for-tat type effort. As an example, the U.S. protested the Soviet-held 1980 Olympics, which scholars assess is why the 1984 Olympics came under fire with multiple disinformation operations.

References:

- Thomas Rid's testimony to the Senate Intelligence Committee on Russian disinformation operations: https://www.intelligence.senate.gov/sites/default/files/documents/os-trid-033017.pdf
- https://twitter.com/RidT/status/799395555101306880

False Flags

- False flags are a very specific type of operation where the purpose of the operation is redirecting the blame of the attack onto a third party
- Obfuscation, distractions, anti-forensic techniques, etc., are not false flags

- E.g., RT claiming CIA malware tried to impersonate Kaspersky Labs led to claims of false flags; but that is just obfuscation

SANS DFIR FOR578 | Cyber Threat Intelligence 117

False Flags

False flag is a term that is getting used a lot more commonly amongst analysts. However, malware authors leveraging overlaps in code, using certificates to hide their traffic, infrastructure that looks like other adversaries or benign communications, etc., are not false flags. The discussion of false flags is an operation, not a type of data or forensic artifact.

The purpose of a false flag is for a third party to be blamed for an attack. As an example, when TV5 Monde was compromised by "Cyber Caliphate" but was actually Russian government operators. The Russian team specifically wanted France to blame extremists which was then leveraged to gain credibility for the Cyber Caliphate with extremists for the Russian operators to further penetrate their operations while also increasing tension between France and extremists.

References:

- https://www.bbc.com/news/technology-37590375
- https://www.rt.com/news/409376-cia-wrote-code-to-impersonate-kaspersky/

False Flag Example: South Korean Winter Olympics

- A cyber attack on the South Korean Winter Olympics took place during the opening ceremony on Feb 9
- The Olympic Destroyer malware used was built to look like samples from Lazarus; C2 servers and other artifacts overlapped as well
- The intent was to increase pressure between North Korea and South Korea
- Analysis and indictments showed it was Russia's GRU


threatpost Cloud Security / Malware / Vulnerabilities / InfoSec Insiders / Podcasts

Romance Scams Drive Necurs Botnet Activity in Run Up to Valentine's Day

Unicode Technique Used

'Olympic Destroyer' Malware Behind Winter Olympics Cyberattack, Researchers Say


Six Russian GRU Officers Charged in Connection with Worldwide Deployment of Destructive Malware and Other Disruptive Actions in Cyberspace

Defendants' Malware Attacks Caused Nearly One Billion USD in Losses to Three Victims Alone; Also Sought to Disrupt the 2017 French Elections and the 2018 Winter Olympic Games

- **PyeongChang Winter Olympics Hosts, Participants, Partners, and Attendees:** December 2017 through February 2018 spearphishing campaigns and malicious mobile applications targeting South Korean citizens and officials, Olympic athletes, partners, and visitors, and International Olympic Committee (IOC) officials;
- **PyeongChang Winter Olympics IT Systems (Olympic Destroyer):** December 2017 through February 2018 intrusions into computers supporting the 2018 PyeongChang Winter Olympic Games, which culminated in the Feb. 9, 2018, destructive malware attack against the opening ceremony, using malware known as Olympic Destroyer;

False Flag Example: South Korean Winter Olympics

On Feb 9, 2018, a group attacked the PyeongChang Winter Olympics. The adversary went to extraordinary effort to compromise known Lazarus command-and-control servers, use the "human fingerprints" such as Rich Headers and encoding keys, and TTPs of the North Korean linked Lazarus group. Security researchers around the world took the bait and started to blame Lazarus and North Korea. Detailed analysis by Talos, Kaspersky, FireEye, Microsoft, and others showed though that there were oddities in some of the evidence. A few observables, including some slight TTP variations and a backdoor leveraged, helped to identify that something was amiss. It was later apparent that this attack was likely not from Lazarus but more likely APT28, a Russian linked group.

Over time, the analysis would become clear and the US DOJ even indicted six Russian GRU officers for this and other attacks (including the Ukraine 2015 and Ukraine 2016 electric power attacks).

This was a great example of a false flag as the intent was not to obfuscate the malware or attack but instead misdirect the blame to North Korea and increase tension between them and the South Koreans during a key international event.

Reference:

- https://www.justice.gov/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and

Coming to the End…Reassess Intelligence Requirements

- Were you able to satisfy the intelligence requirement?
  - If not, are you going to be able to with additional time or resources?
  - If so, is it an intelligence requirement that will continue on, or is it completed?
  - Are there new knowledge gaps now requiring new intelligence requirements?
- You will have lessons learned at the end of your intelligence process
  - Leverage the lessons learned to adapt the requirements or your process
- You may find that you need to adapt collection as well
  - Collection Management Framework may need to be adapted
  - Collection requirements can generate Requests for Information (RFIs) as well

Reassess Intelligence Requirements

At the end of the intelligence process, you will find that the product you created satisfied the intelligence requirement and you have completed that one; you might find that you satisfied the requirement, but it's an ongoing requirement, which means you should have lessons learned, or you might find that you were not able to satisfy the requirement for a variety of reasons. In all circumstances, you need to attempt to reassess the intelligence requirements.

You will either determine that it's completed, not completed, needs to modification, or isn't something you are going to be able to satisfy. In all cases, you should document lessons learned, key evidence, and data sources, and ensure that you adapt your process through this effort.

(Optional) Case Study: Lazarus Group

A Good Example of a Threat Actor Poorly Defined

Depending on the time in the class this can be moved to an optional case study. It's a great case study and is on the class recording but in live classes is often moved to optional to give you more time on the last lab.

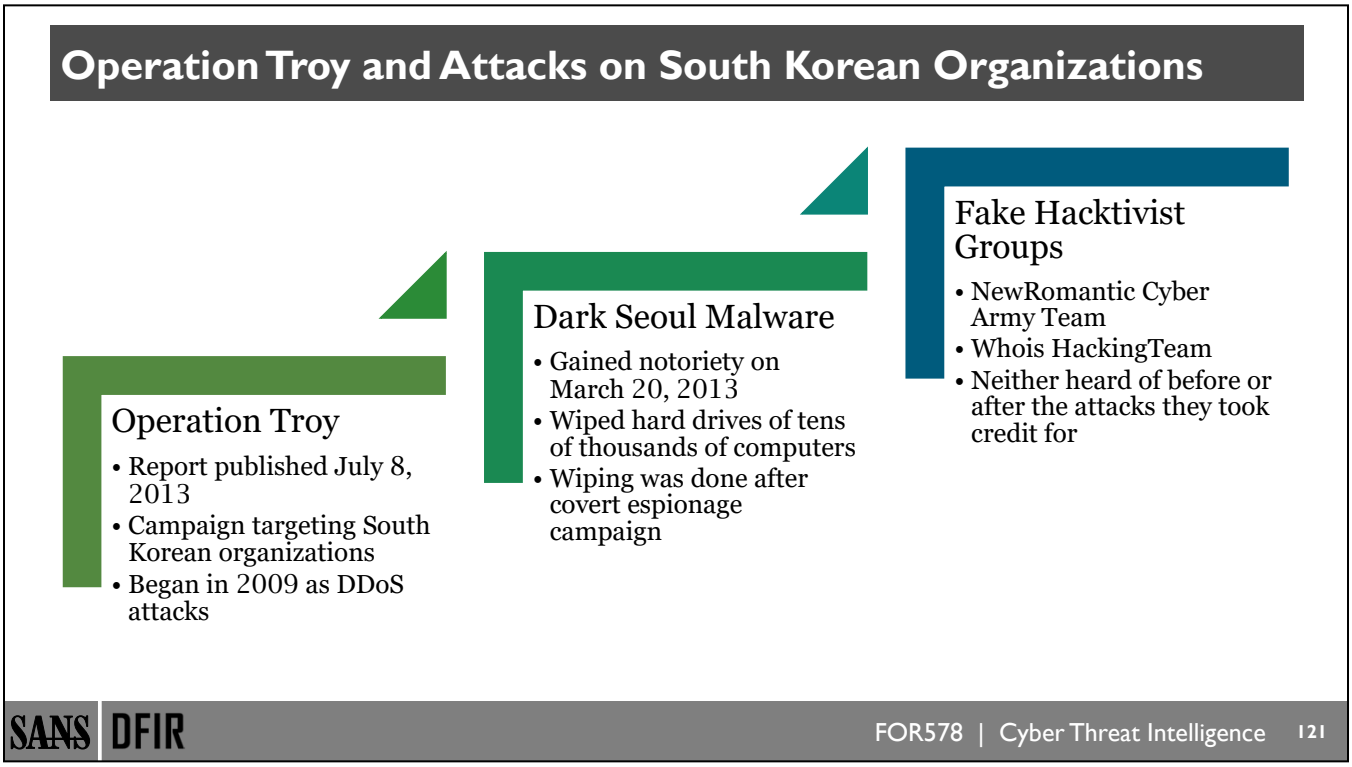

Operation Troy and Attacks on South Korean Organizations
In 2013, McAfee published a whitepaper covering the Dark Seoul malware that was leveraged against South Korean organizations. The paper dubbed the campaign Operation Troy and detailed how the master boot record (MBR) wiping malware was first leveraged as a sophisticated espionage campaign. The Dark Seoul malware contained remote access Trojan (RAT) functionality, which allowed the adversaries to conduct espionage after infecting the organization through spear phishing emails. The targets of the attack were largely banks and media organizations as well as government networks. Interestingly, the attackers took credit for the attacks as "NewRomantic Cyber Army Team" and on other occasions "Whois HackingTeam." Neither of those hacktivist groups were ever heard of before each attack nor were they active following the attacks. Each seemed to be made-up hacktivist groups designed to take credit for the attacks.

Although 2013 was the first major public attention paid to the malware, the analysis of Operation Troy detailed that there were numerous attacks going back to DDoS efforts in 2009 against other South Korean targets.

Reference:
- https://paper.seebug.org/papers/APT/APT_CyberCriminal_Campagin/2013/dissecting-operation-troy.pdf

Dark Seoul
Dark Seoul was a piece of malware that began to be identified around 2009. For the first few years, it was only observed as being espionage focused in nature and targeted South Korean targets such as news organizations and government offices. Whenever there were U.S. and South Korean military drills, there was a spike in activity, and U.S. personnel and facilities would also be targeted.

In 2013, there was a data destruction module added to the malware.

Reference:
- https://www.mcafee.com/enterprise/en-us/resource-library.html

The Sony Attack

- On November 24, 2014, a group identifying itself as "Guardians of Peace" leaked sensitive Sony Pictures Entertainment data
- Warnings continued that if Sony did not cancel the film debut of "The Interview," that it would do more damage
- Data wiping modules deleted numerous systems MBR's, and films were leaked online
- In December, the group made claims referring to September 11, 2001, that ultimately drew the attention of the US Government

The Sony Attack

In 2013, attacks launched against South Korea banks and news organizations wiped systems and overwrote hard drives with the word "hastati," which is a reference to military classes of ancient Rome signifying that hastati were younger and weaker soldiers.

The attackers used infrastructure in China and left a pop-up banner for their victims claiming that "Whois" attacked them. There was also a note stating that the "NewRomantic Cyber Army Team" hacked them (pasted below). Researchers at McAfee noted that the groups were likely the same group and neither of the "hacktivist" groups had been heard of before.

"Hi, Dear Friends, We are very happy to inform you the following news. We, NewRomanic Cyber Army Team, verified our #OPFuckKorea2003. We have now a great deal of personal information in our hands. Those include 2.49M of [redacted by McAfee] member table data, cms_info more than 50M from [redacted]. Much information from [redacted] Bank. We destroyed more than 0.18M of PCs. Many auth Hope you are lucky. 11th, 12th, 13th, 21st, 23rd and 27th HASTATI Detachment. Part of PRINCIPES Elements. p.s For more information, please visit www.dropbox.com login with joseph.r.ulatoski@gmail.com::lqaz@WSX3edc$RFV. Please also visit pastebin.com."

In relation to the Sony Attack, it was this time the "Guardians of Peace," and they leaked movies and contracts from Sony in an apparent retaliation for the publishing of "The Interview."

References:

- https://www.nytimes.com/2013/03/21/world/asia/south-korea-computer-network-crashes.html
- https://en.wikipedia.org/wiki/Sony_Pictures_Entertainment_hack
- https://krebsonsecurity.com/2014/12/the-case-for-n-koreas-role-in-sony-hack/

Government Attribution

- December 17, 2014, U.S. government officials stated that the North Korean government was "centrally involved"
- President Obama stated that the U.S. would wage consequences when it sought best
  - In January 2015, the U.S. levied sanctions on North Korea
- The Federal Bureau of Investigation formally stated on December 19 that the North Korean cyber warfare agency Bureau 121 was involved in the attacks
- The speed to which the attribution was obtained made many, including the *New York Times,* speculate and put forth evidence that the National Security Agency was already in the North Korean networks and observed the attack

Government Attribution

On December 17, 2014, the U.S. disclosed through government officials that the U.S. believed that North Korea was responsible for the attacks. President Obama commented on the attacks, followed by the FBI formally stating that North Korea cyber warfare agency Bureau 121 was involved in the attacks. This move shocked many, as the president responding to an attack on a private company was rare. However, many analysts believe this was due in part to the 9/11-styled threats made by the attackers and also because Sony Pictures Entertainment pulled "The Interview"; this is a stifling of freedom of speech that is central to American ideals and seen as a national security issue.

The speed of the attribution obtained made many speculate that the NSA was already inside the networks of the North Korean actors. This was stated by the *New York Times,* but the U.S. government has never responded to those claims.

References:

- https://en.wikipedia.org/wiki/Sony_Pictures_Entertainment_hack
- https://www.nytimes.com/2015/01/19/world/asia/nsa-tapped-into-north-korean-networks-before-sony-attack-officials-say.html

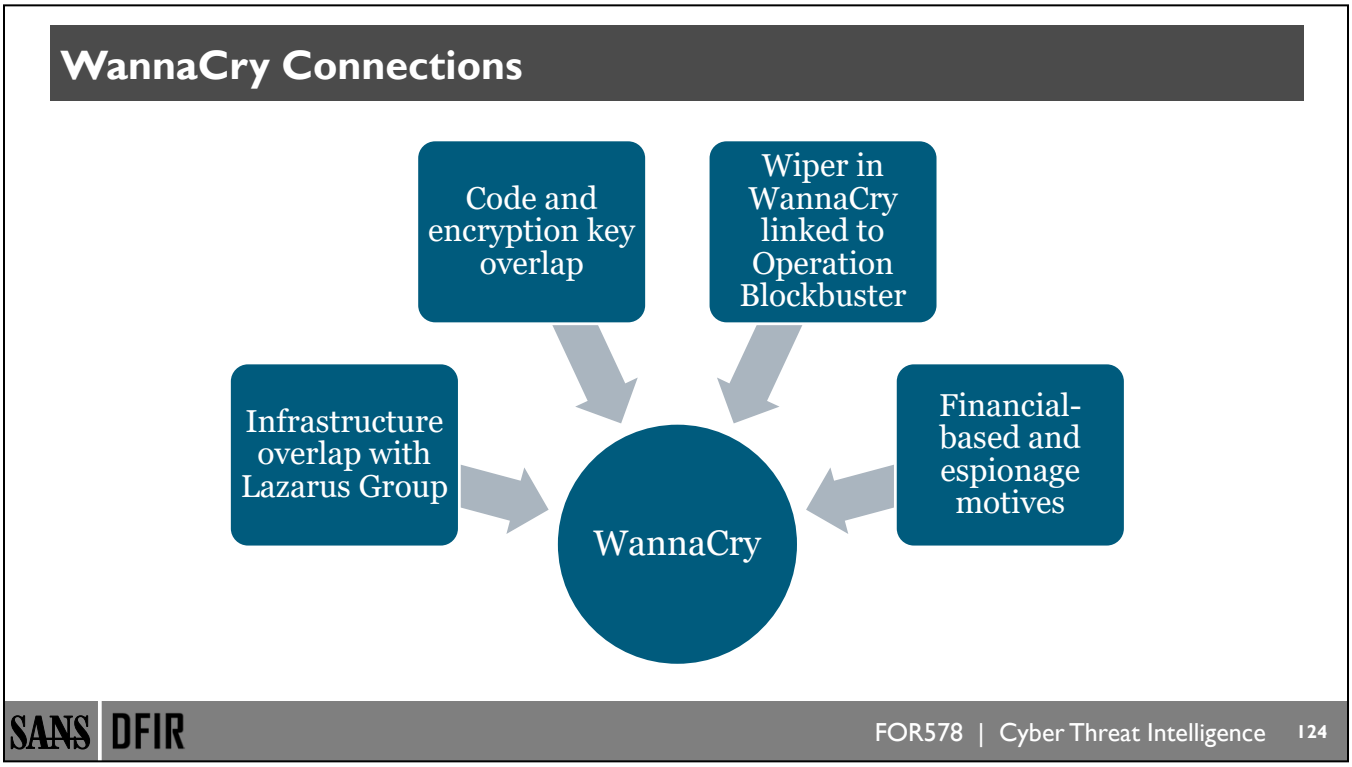

WannaCry Connections

Following the WannaCry infections of 2017, multiple security companies, including BAE and Kaspersky Labs, performed code analysis between the Lazarus Group samples and the WannaCry malware. There was particular overlap, including specific encryption key usage, between the malware samples unique to the adversary group. More important than just code overlap was also command-and-control (C2) overlap between the infrastructure the adversaries were using to distribute WannaCry and infrastructure sites the Lazarus group has used before.

There were many theories around WannaCry and why it worked like it did—including how poorly the ransom component was orchestrated—but as analysts, we have to be careful not to let ourselves be overly biased. The ransomware functionality of the malware could have been a disguise like the wiper component of Dark Seoul after the espionage program. It could have been a mistake as well, though, letting it loose on the internet. Additionally, maybe it was simply handled very poorly, and one team worked on the malware while another team inside the Lazarus group worked, less successfully, on the ransom portion.

As organizations try to piece together the ties to the group, though, it is important to think about whether or not the Lazarus group represents a threat to your organization, what you would do as cyber threat intelligence analysts, and how to use structured analytical techniques to help in your analysis.

References:

- https://securelist.com/wannacry-and-lazarus-group-the-missing-link/78431/
- https://community.broadcom.com/symantecenterprise/communities/community-home/librarydocuments/viewdocument?DocumentKey=b2b00f1b-e553-47df-920d-f79281a80269&CommunityKey=1ecf5f55-9545-44d6-b0f4-4e4a7f5f5e68&tab=librarydocuments

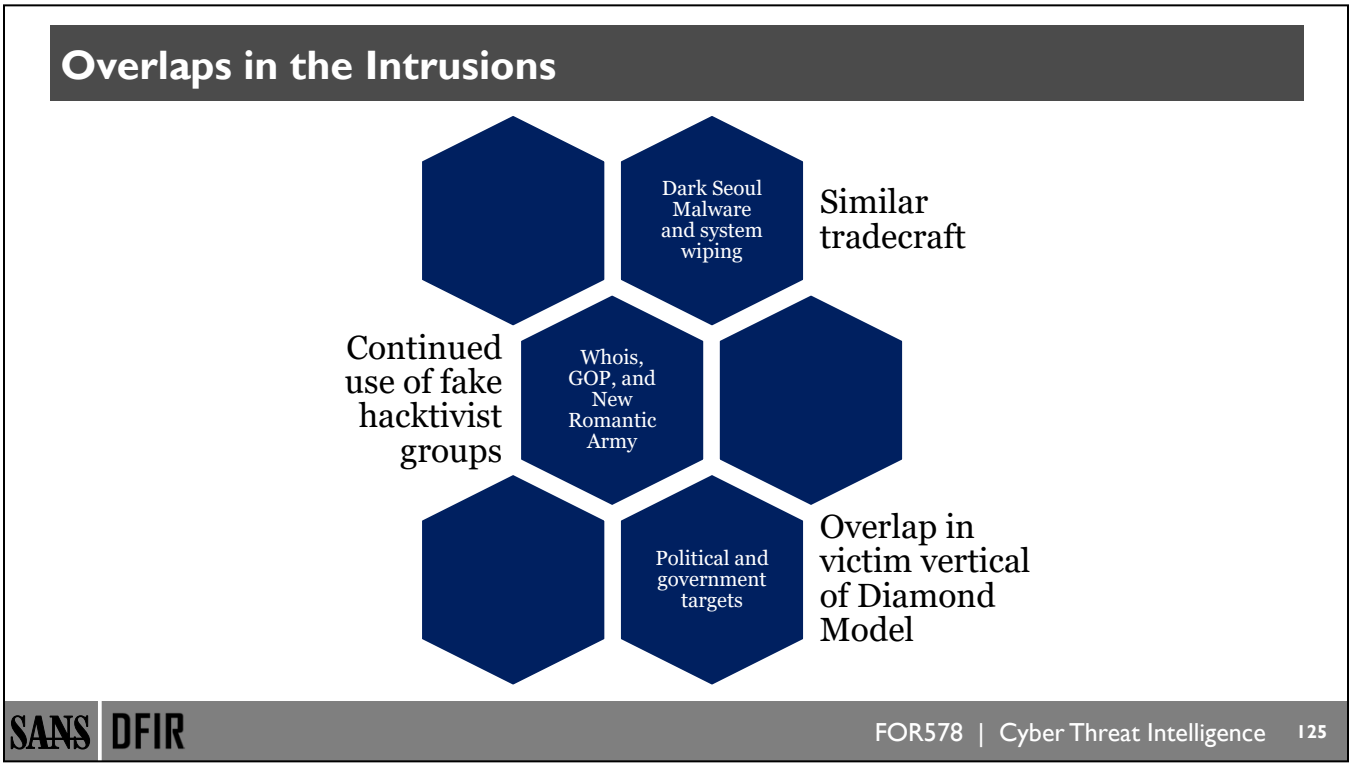

Overlaps in the Intrusions

Previous North Korean-attributed threat actors that utilized malware such as Dark Seoul and targeted locations such as South Korean media-based companies had many similarities with what occurred in Sony.

There were significant tradecraft similarities in the infection vectors then corresponding with data deletion, and then false attribution by the adversaries pretending to be a hacking collective that was previously unheard of; even the style of the graphics used was very similar.

One of the more significant overlaps was in infrastructure reuse, as well as malware and code reuse, including a spelling error in the code.

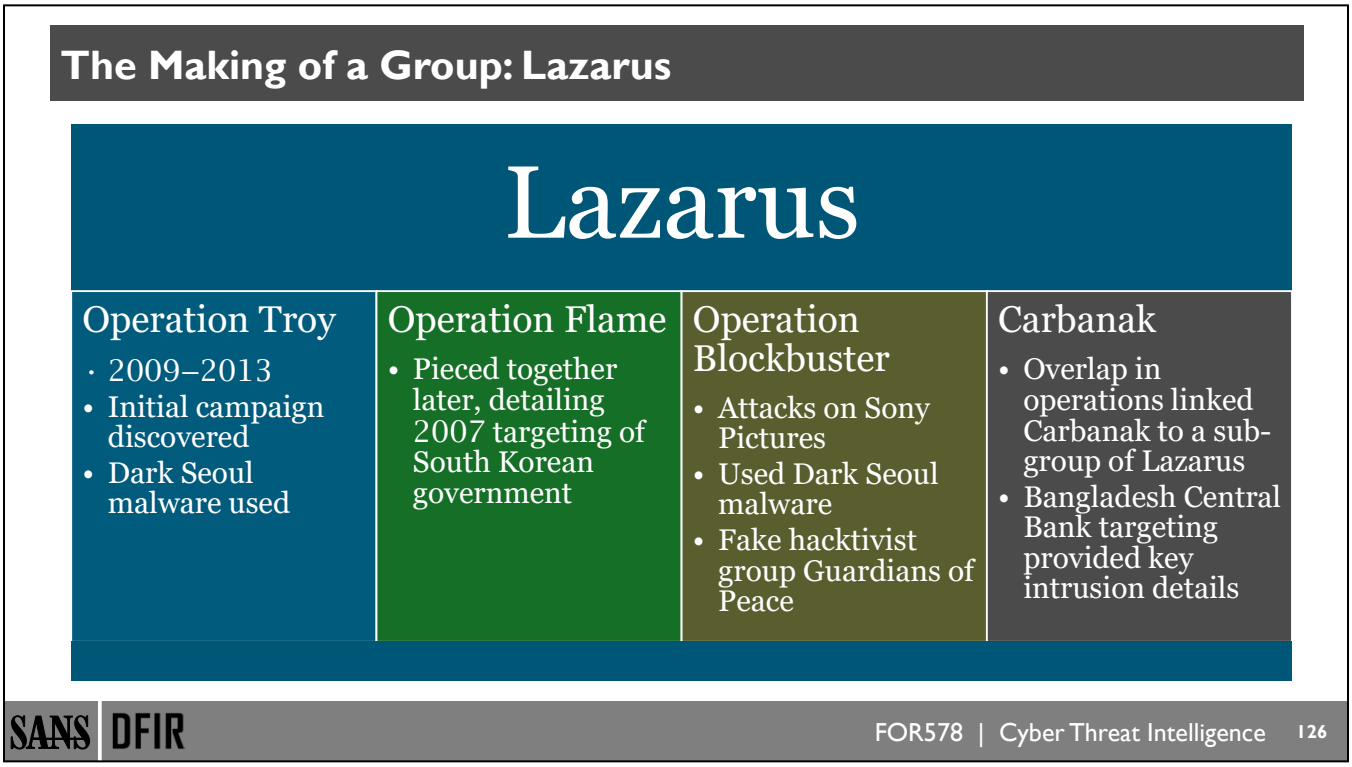

The Making of a Group: Lazarus

As the individual intrusions were analyzed across numerous targets and years' worth of data, there were distinctive patterns that formed. There was a commonality in the malware, in the style of the intrusions, and even in the naming convention of fake hacktivist groups each time an attack was orchestrated. But Operation Troy was simply one look at what would become known as the Lazarus Group. Over the course of a few years, there were various campaigns orchestrated by what appeared to be a single group. This group consistently targeted South Korean media and financial companies and eventually began targeting other organizations, such as the Bangladesh Central Bank and Vietnamese Tien Phong bank.

Interestingly enough, the analysis helped tie together an even earlier campaign than Operation Troy called Operation Flame. This was the first generation of the Dark Seoul malware leveraged against the South Korean government for a long-term espionage program.

References:

- https://community.broadcom.com/symantecenterprise/communities/community-home/librarydocuments/viewdocument?DocumentKey=257dd693-5986-41bf-bc33-f9dc76d9c6a8&CommunityKey=1ecf5f55-9545-44d6-b0f4-4e4a7f5f5e68&tab=librarydocuments
- https://www.wired.com/2016/02/sony-hackers-causing-mayhem-years-hit-company/
- https://media.kasperskycontenthub.com/wp-content/uploads/sites/43/2018/03/07180244/Lazarus_Under_The_Hood_PDF_final.pdf

Problem with Extending Too Far

- Each individual "Operation" and clustering was interesting
- Combining various clusters together over the years abstracted away the specifics
  - The focus became *who* was responsible, not *how*
    - Defenders told to defend against "Lazarus" must ask numerous follow-on questions such as:
      - What period of time of Lazarus?
      - What victimology?
      - What tradecraft?
- Lazarus is effectively a blob of activity that does not mean much more than "activity associated with the North Korean government"
  - Not all the activity is North Korean
  - Tradecraft overlaps clustered too much together
  - Still, it can be a useful abstraction to some

Problem with Extending Too Far

Lazarus contains activity dating back a decade. It is a collection of clusters of intrusions from various teams against various intelligence requirements against various collections. The analysis done along the way has morphed and been largely useless to any defenders. Trying to defend against "Lazarus" isn't specific enough, so individual clusters need to be broken out. Additionally, stating that the clusters are all "Lazarus" is an attribution-based schema, but it itself is not well defined publicly. The clusters effectively morph into a blog of activity that all is supposed to represent "North Korean activity." However, it's not feasible that every intrusion or cluster pulled into Lazarus is 100% correct attribution. That can create links that are spurious and guide inaccurate analysis.

However, the blob still might be useful as an abstraction, depending on your intelligence requirements.

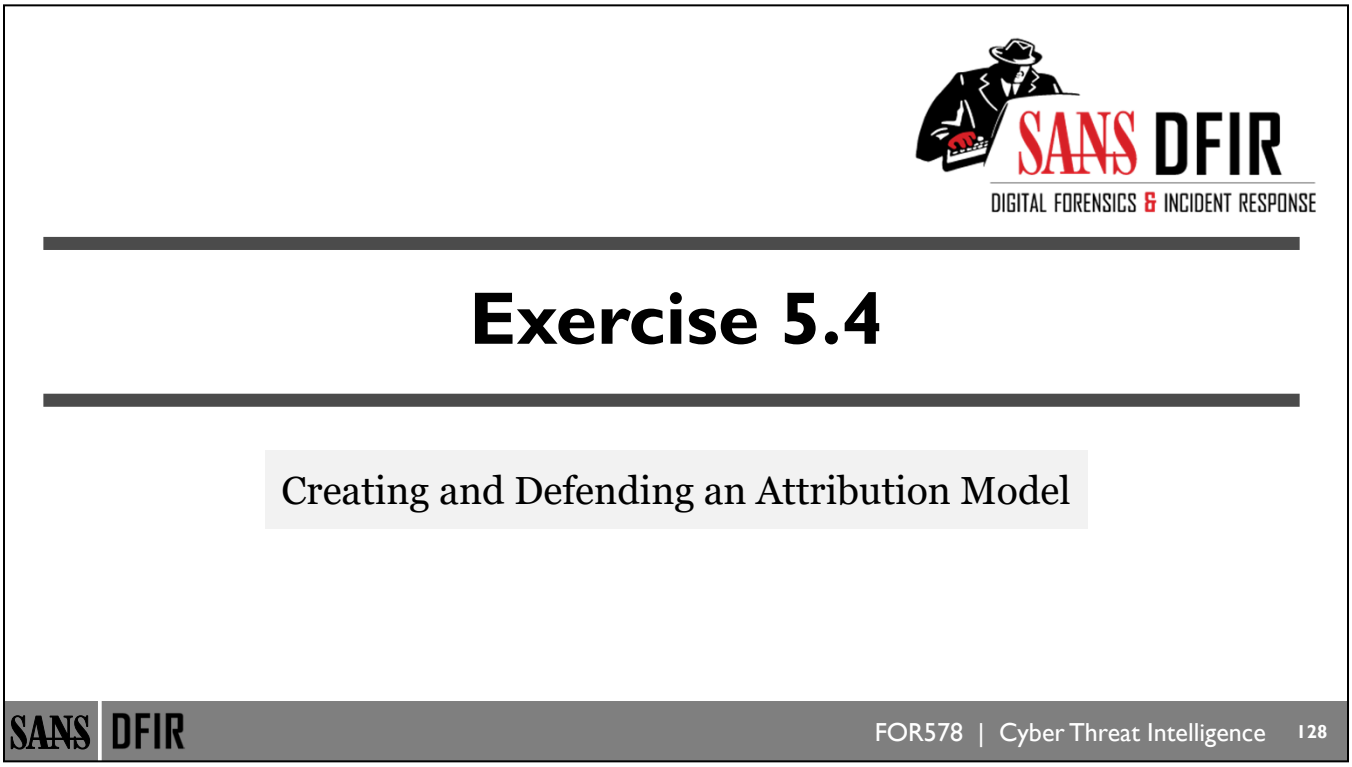

This page intentionally left blank.

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