THREAT INTELLIGENCE AND THE LIMITS OF MALWARE ANALYSIS

Joe Slowik, Dragos Inc.

JANUARY 2020

Table of Contents

Executive Summary

Malware analysis forms a vital part of cyber threat intelligence operations. Yet the proliferation of binary-focused analysis enabled by tools such as VirusTotal that make samples available widely absent victim and use context yield analysis lacking significant amplifying information. This is not to say that any specific malware analysis performed is wrong, but rather that malware-exclusive analysis may miss contextuality, significance, and use-cases that are vital aspects of understanding a security incident.

Looking at some high-profile instances from the past four years of security reporting, there are many examples of excellent technical reports in isolation which nonetheless miss some critical aspects of certain security incidents or appear to indicate connections that greater context reveals as unsupported. By incorporating other aspects of security event analysis – host artifacts, network infrastructure, network traffic, and where possible adversary motivations and objectives – cyber threat intelligence analysts can gain greater, more accurate insight into activities.

Ultimately, threat intelligence producers will rarely have the “full picture” of an incident, but whether limited to malware or some other single aspect of an event, analysts must ensure that resulting products are properly dispositioned and resulting conclusions supported by the available evidence. Threat intelligence consumers must realize the limitations faced by producers and formulate their own analysis and integration of third-party threat intelligence to incorporate other sources to fill in gaps where possible. Recognizing the limitations of analysis based on small sample sizes using only a single analytical method means network defenders overall, and threat intelligence practitioners specifically, can accurately categorize observations and apply controls and defenses in a supportable fashion.

Introduction

The commercial threat intelligence field largely came into existence with the publication of the Mandiant APT1 report in 2013.[1] While previous reports on cyber threats existed, many took the form of one-off blogs, or media stories lacking the analytic rigor currently associated with cyber threat intelligence reporting.[2] Initially, non-government threat intelligence was limited to entities with direct access to data as a result of incident response work, security tool telemetry, or other direct exposure to events. However, the launch of VirusTotal by Spanish security company Hispasec significantly changed matters.[3] Initially designed as a mechanism for users to perform security scans of files using multiple anti-virus (AV) products on its debut in 2004, the service gradually expanded (especially following acquisition by Google in 2012) to provide a paid service enabling researchers (and companies) beyond the participating AV vendors to retrieve files for analysis.

With this move, an entirely new field of security research opened with the availability of a large corpus of malware samples beyond any researcher's or company's individual experience and actions. While this resulted in an explosion of research and analysis of malicious software, this shift also generated an ecosystem where threat intelligence and analysis became increasingly dominated by pure malware analysis.

The question we as practitioners now face is, while the benefits of greater exposure, sharing, and availability have produced a dramatically larger corpus of analysis and information, has this effort come at a cost? This paper will argue that while malware analysis offers significant benefits for generating threat intelligence, it is not perfect. Particularly, malware analysis on its own imposes certain limitations on contextuality and purpose, important items that are typically unavailable in pure binary or malware sample examination. Understanding these limitations and adjusting for them nonetheless allows defenders to incorporate this information, provided the results can be placed in proper context.

Goal of Threat Intelligence

Threat intelligence is typically defined as knowledge that enables defensive action,[4] or knowledge that allows for prevention or mitigation of attacks.[5] There are many examinations producing more specific definitions, as well as multiple frameworks for organizing or systematizing threat intelligence information and knowledge, but these general definitions are sufficient for our current purposes. More importantly, the fundamental goal of threat intelligence is to provide some mechanism for an organization to prepare for, or defend against, an event or attack which it has not already been the victim of; or, to provide mechanisms to identify an intrusion which may have otherwise gone unnoticed.

Thus, threat intelligence's value proposition to an organization comes explicitly from its ability to enable and enhance operations. This can range from something as simple as distributing raw observables or more refined indicators of compromise (IOC) to detailing attacker techniques and methodologies around which more complex (and robust) defense can be built.[6] In either case, network defenders, or those charged with evaluating (and accepting) network security risk, should be able to extract actionable, practical value for threat intelligence to be of use. While information and reporting that does not meet these criteria of usefulness and actionability may be quite interesting, for practical defensive purposes they become so much "frictionless spinning in a void".[7]

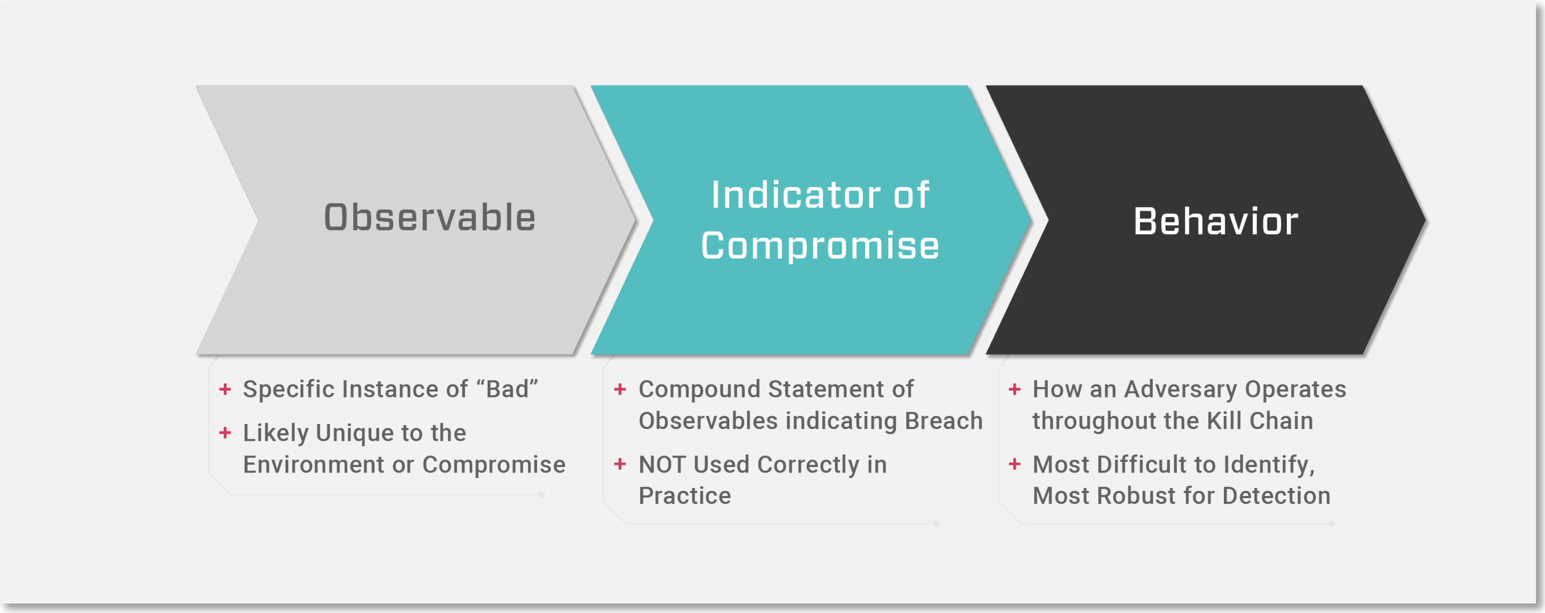
Observable
Indicator of Compromise
Behavior
+ Specific Instance of "Bad"
+ Likely Unique to the Environment or Compromise
+ Compound Statement of Observables indicating Breach
+ NOT Used Correctly in Practice
+ How an Adversary Operates throughout the Kill Chain
+ Most Difficult to Identify, Most Robust for Detection

Figure 1: Cyber Threat Intelligence Evolution

Ultimately, we want cyber threat intelligence to follow a sequence of actions yielding greater understanding of the threat environment. Raw data analysis and IOC tracking generates the foundation upon which we can then begin to piece together different observables to yield more general adversary behaviors. Based on these behaviors and adversary tendencies, defenders can begin articulating specific defensive measures and mitigations to counteract adversary behavior (or queries and searches to identify if that actor is already present). Additionally, organizational leadership can take portions of this information to judge the risk faced by their operation, either from specific behaviors or techniques to which they may be susceptible or from certain types of actor based on assessed objectives and goals.

Investigation and Malware Analysis

Malicious software, or malware, forms part of many (if not most) network security incidents, from highly complex, custom-developed binaries to scripting objects built from public repositories. Malware analysis is the practice of "dissecting malware to understand how it works, how to identify it, and how to defeat or eliminate it."[8]

The progression of deliverables from understanding functionality, developing mechanisms to recognize, then building countermeasures against malware roughly fits the criteria established above for threat intelligence. Analysis of a given sample should yield more than an investigation of an esoteric encoding/decoding routine or evasion technique, but also seek to identify practical mechanisms to identify and defeat these techniques, or some other aspect of the malware's functionality so as to inhibit its effectiveness and usefulness.

Historically, malware analysis took place either as part of a broader security investigation or incident response, or among small groups of researchers who could share samples – and context around where the samples came from and their assessed purpose. In these cases, malware analysis came with contextuality: why it was significant, where it was located, how it was used, and what its immediate defensive implications were. In the last decade or so, the proliferation of sample sharing and distribution portals, whether commercial (VirusTotal) or free (Any.Run, Malshare)[9] have enabled wider distribution and greater availability of malware samples – but at the cost of stripping context from them.

While samples now lack valuable contextual information surrounding their use and origin, much of the process of malware analysis can nonetheless take place. As a result, purely technical analysis can flourish, removed from any grounding in network or security operations.

Issues, Concerns, and Limitations

Malware analysis is an important, and often invaluable, portion of overall threat intelligence operations. Yet malicious file analysis also represents but one technique among many for understanding and analyzing adversary behavior. Understanding malware functionality (and how to defeat it) is undeniably important, but in isolation from other factors or absent execution context it may be significantly less useful than it otherwise would be with additional, enriching information.

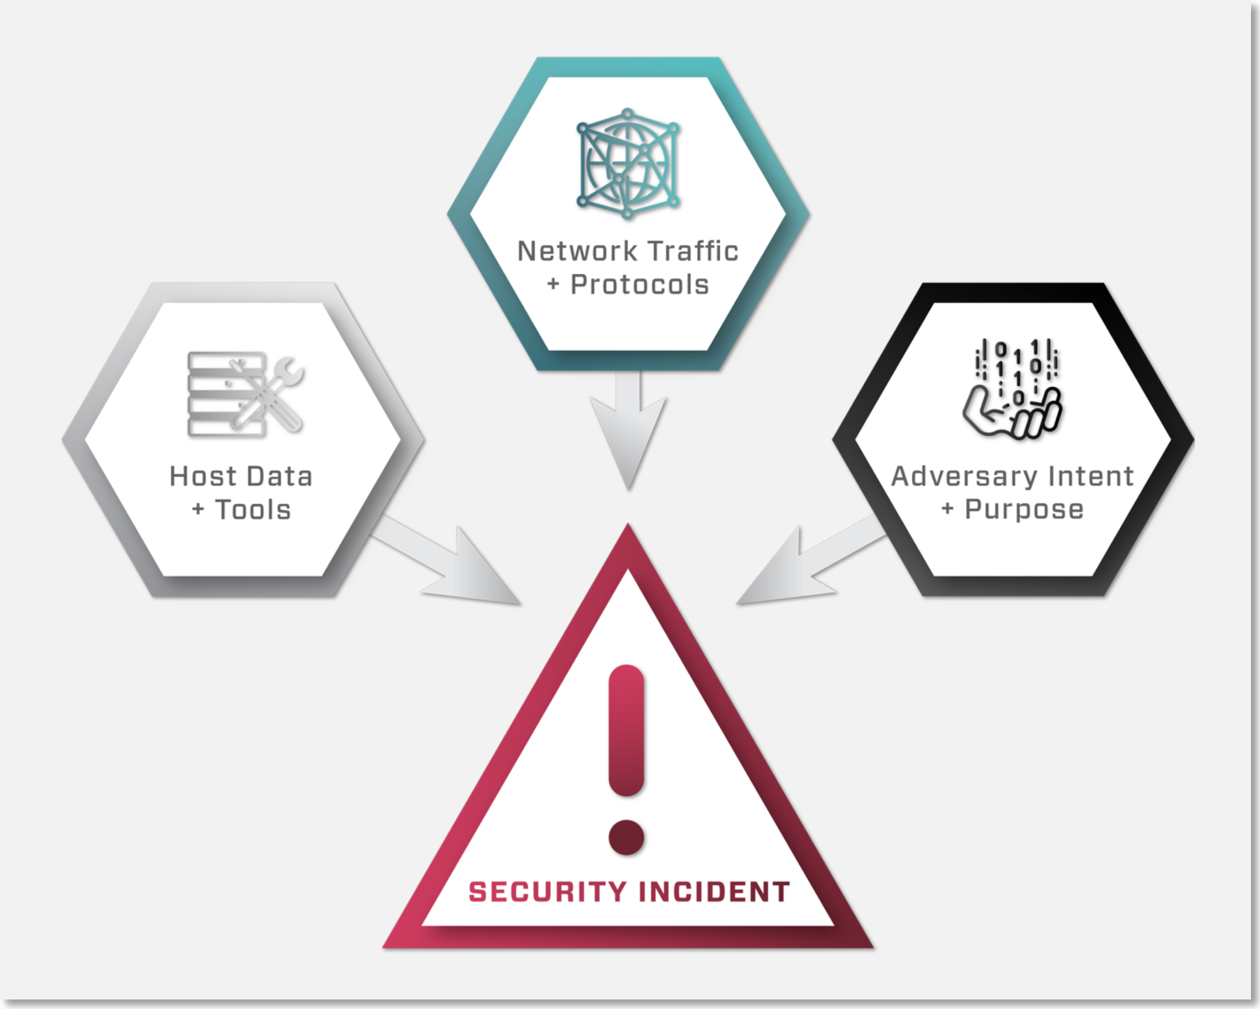

Figure 2: Security Incident Components

A security incident consists of three general components: host artifacts and tools; network data and traffic analysis; and adversary objectives or intentions. A ransomware attack would consist of the specific ransomware malware involved, the means through which it propagated across the network, or entered into the environment in the first place and the adversary's purpose (ostensibly to make money, although items such as NotPetya can muddy these waters)[10]. Yet a malware-specific approach to such an investigation, especially if the only data source is a sample sharing service, provides only the ransomware as an object for investigation and analysis. There is much that can be learned from the sample, but depending on its functionality, there may be multiple other steps or necessary observables required to identify critical defensive needs – such as propagation mechanisms, host artifacts specific to the victim's network (as opposed to the analysis environment), network traffic patterns, and related information. Malware analysis on its own only gets us ***part*** of the way to answering these questions in many circumstances.

Examining specific circumstances, where an overemphasis on malware analysis was used to formulate conclusions, we can identify cases not necessarily representing mistakes, but where matters were less useful than they could be or failed to capture certain degrees of nuance. Thus, analysts generated and disseminated threat intelligence, but the resulting product may not have been as complete, accurate, or useful lacking broader event context from other sources and analytical mechanisms.

2016 UKRAINE POWER EVENT

The 2016 Ukraine power event represented the first known electric power incident induced through malware.[11] While some preliminary conference talks addressed the event along with media reporting, the first technically detailed, published report was ESET's analysis of Industroyer, followed by publication under the name CRASHOVERRIDE by Dragos. ESET's report appeared to cover multiple phases of the intrusion, which initially seemed to originate from a remote access tool (RAT) facilitating network access which enabled deployment of the actual electric transmission manipulation software. The tools were grouped as distinct stages in a single intruder event, building up along typical kill chain steps to final actions on objectives through the ICS payload. From a pure malware analysis perspective, this approach makes sense as the identification of the set of tools associated with the incident including a backdoor or RAT combined with industrial control-specific malware indicates that the former would be used to position and then execute the latter.

While initial analysis of the malware is accurate in terms of each sample's capability, the absence of contextual incident information left some items (such as additional adversary actions to enable malware installation) unexplained. There were some hints in public talks and media reporting as to the greater scope of the event, but no in-depth, written technical coverage.[12] Reports emerging years after the event, in 2018 and 2019, eventually addressed in detail how multiple tools were in play in the victim network, and execution of the industrial control system (ICS) specific portion of the attack does not appear necessarily linked to the backdoor.[13] Furthermore, while some such as Oleksii Yasynskyi had publicly hinted that multiple entities were likely involved in the 2016 power event,[14] more recent review of events shows apparently different adversary intrusion techniques in play. Enterprise IT intrusion steps displayed a significantly different set of behaviors than the ICS-focused portion of the attack. This observed bifurcation appeared to confirm a multiple actor intrusion hypothesis.

A few years after the incident, ESET identified new malware, called EXARAMEL, with significant functional overlap with the 2016 Ukraine event backdoor. Based on the technical overlap and a malware-centric investigative approach, previous analysis looking at the power event as a single-actor incident supported transitive attribution that the users (or at least authors) of EXARAMEL were also responsible for not merely the

Industroyer backdoor, but the overall 2016 Ukraine power event. In the case of EXARAMEL, ESET analysts were able to link this malware to TeleBots, known elsewhere as Sandworm.[15] Given past analysis and the backdoor code overlap, by extension 2016 Ukraine appears to become a Sandworm event.

Yet analysis incorporating information from the victim environment and the overall intrusion – not available if conducting or limited to pure binary analysis – indicates a more subtle operation as outlined above. Instead of a single, monolithic intrusion, the 2016 Ukraine event appears to be a case where at least two distinct teams were operating at different stages of the incident: Sandworm serving as an initial access and penetration team, handing over access to ELECTRUM as an ICS-specialist group to perform final control system network penetration and the ICS attack.

Key points establishing this differentiation center around the Sandworm-linked backdoor similar to EXARAMEL. Analyzing available event data shows no evidence that the backdoor is linked to the introduction of the ICS attack payloads, or in their execution in the victim environment. Instead, these stages of the incident use separate sets of techniques, including simple scripting, built-in operating system tools, and credential re-use. Additionally, network infrastructure associated with the backdoor maps to known Sandworm tendencies (e.g., using servers also functioning as TOR infrastructure at the time of the incident), while other network infrastructure more closely aligned with the ICS-specific portion of events lack these characteristics. A summary of these differences is provided in Figure 3 below.

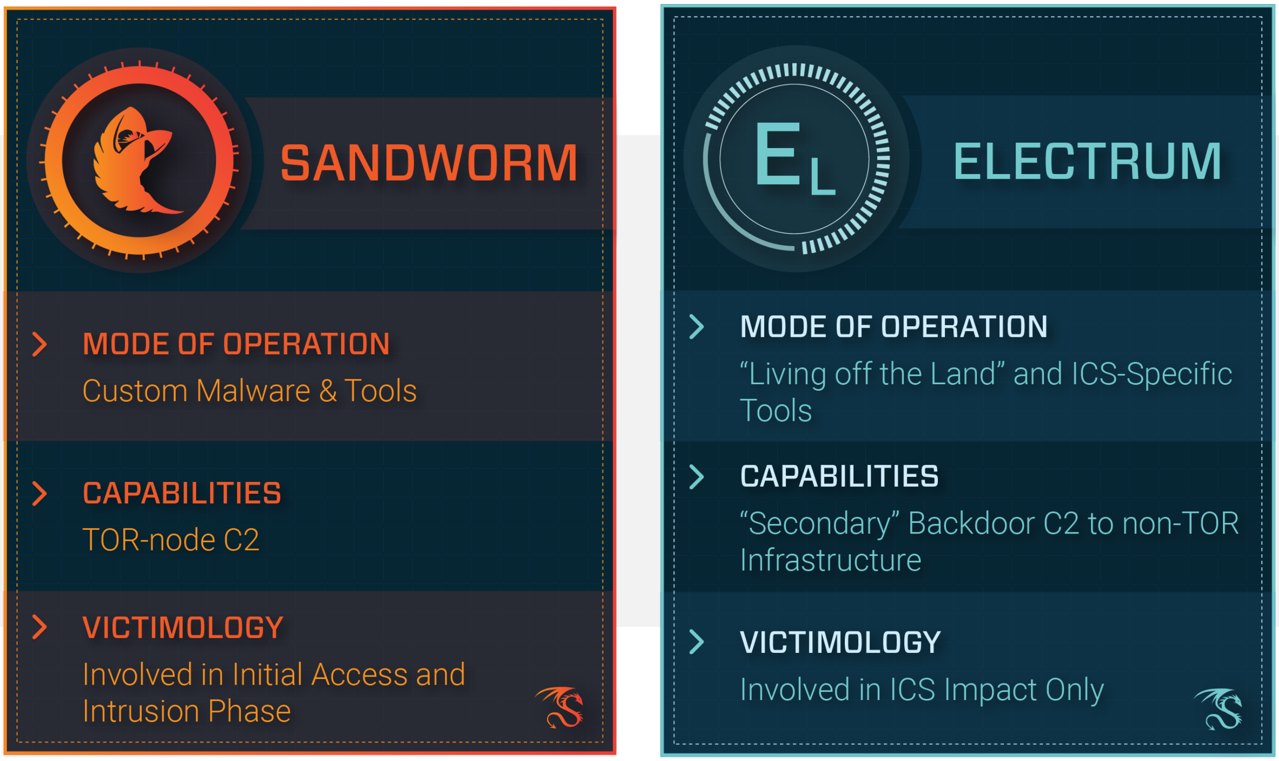
SANDWORM
MODE OF OPERATION
Custom Malware & Tools
CAPABILITIES
TOR-node C2
VICTIMOLOGY
Involved in Initial Access and Intrusion Phase
EL
ELECTRUM
MODE OF OPERATION
“Living off the Land” and ICS-Specific Tools
CAPABILITIES
“Secondary” Backdoor C2 to non-TOR Infrastructure
VICTIMOLOGY
Involved in ICS Impact Only

Figure 3: SANDWORM and ELECTRUM Comparison in 2016 Ukraine

While initial reporting on Industroyer/CRASHOVERRIDE within the context of the 2016 Ukraine event provided significant information for defensive response, it did so largely in context of observed malware alone. Thus, the picture provided was incomplete and lacked insight into, and coverage of, critical aspects of the event: penetration of the control system network and execution of potentially destructive malware in the environment.[16] Once additional, detailed information became available, a more thorough analysis became possible, yielding not only distinct phases of adversary operation (one relying mostly on custom malware, the other primarily on "living off the land" techniques), but possibly distinct entities involved in the event as well. From a threat intelligence application, this more complete picture empowers more effective, broader reaching defense.

ACTOR LINKING ON TOOLS

One common methodology in malware-focused threat intelligence, observed to some degree in the discussion above given connections between the Industroyer backdoor and EXARAMEL, is grouping activity based on tool use and design. However, while ESET's analysis of EXARAMEL strongly suggests additional telemetry was used to tie EXARAMEL to Sandworm, many other similar assessments are made purely on malware analysis alone. In these cases, analysts can quickly reach problematic conclusions as they effectively focus more on developers or capability sources rather than actors or executors.

The most obvious cases where researchers can get into difficulty concerns publicly- or commercially available tools. As attackers have broadened arsenals away from primary reliance on custom-developed malware for operations to embrace a combination of open source (e.g., Mimikatz)[17], criminal source (e.g., PlugX)[18], and commercial (e.g., CobaltStrike)[19] tools, linking activity to an entity or campaign based purely on tool use becomes very problematic. The proliferation of tools such as PlugX, ChinaChopper, njRAT, and the host of open-source libraries now used by attackers make tool-centric attribution increasingly weak in many cases. This can result in attribution errors, which in turn can lead to false assumptions on other aspects of the given attack lifecycle – such as accompanying, expected behaviors and attacker objectives. The result is inaccurate or faulty overall threat intelligence that can lead to problematic or erroneous responses.

A more interesting variant of this phenomenon of tool tracking occurs when considering development and division of labor for threat actors. Malicious cyber entities need not be monolithic bodies with dedicated development resources working exclusively to support a certain threat actor. Instead, for the most complex and well-resourced entities, separate operational teams are likely supported by development shops, contractors, or even commercial acquisition. These entities in turn may contribute tools and capabilities to multiple distinct entities. Furthermore, in shallow pools of elite talent, individual developers and related entities may shift jobs or roles resulting in similar coding patterns and technical "tells" migrating with individuals or even contracted teams when they begin supporting a different entity. Thus, tracking even custom tools may reveal more about developers than it yields on operational teams and their objectives.

Some headline discoveries on what appear to be clusters of activity – such as Animal Farm or Equation Group-related malware[20] – may therefore track development resources and not actual operational teams.

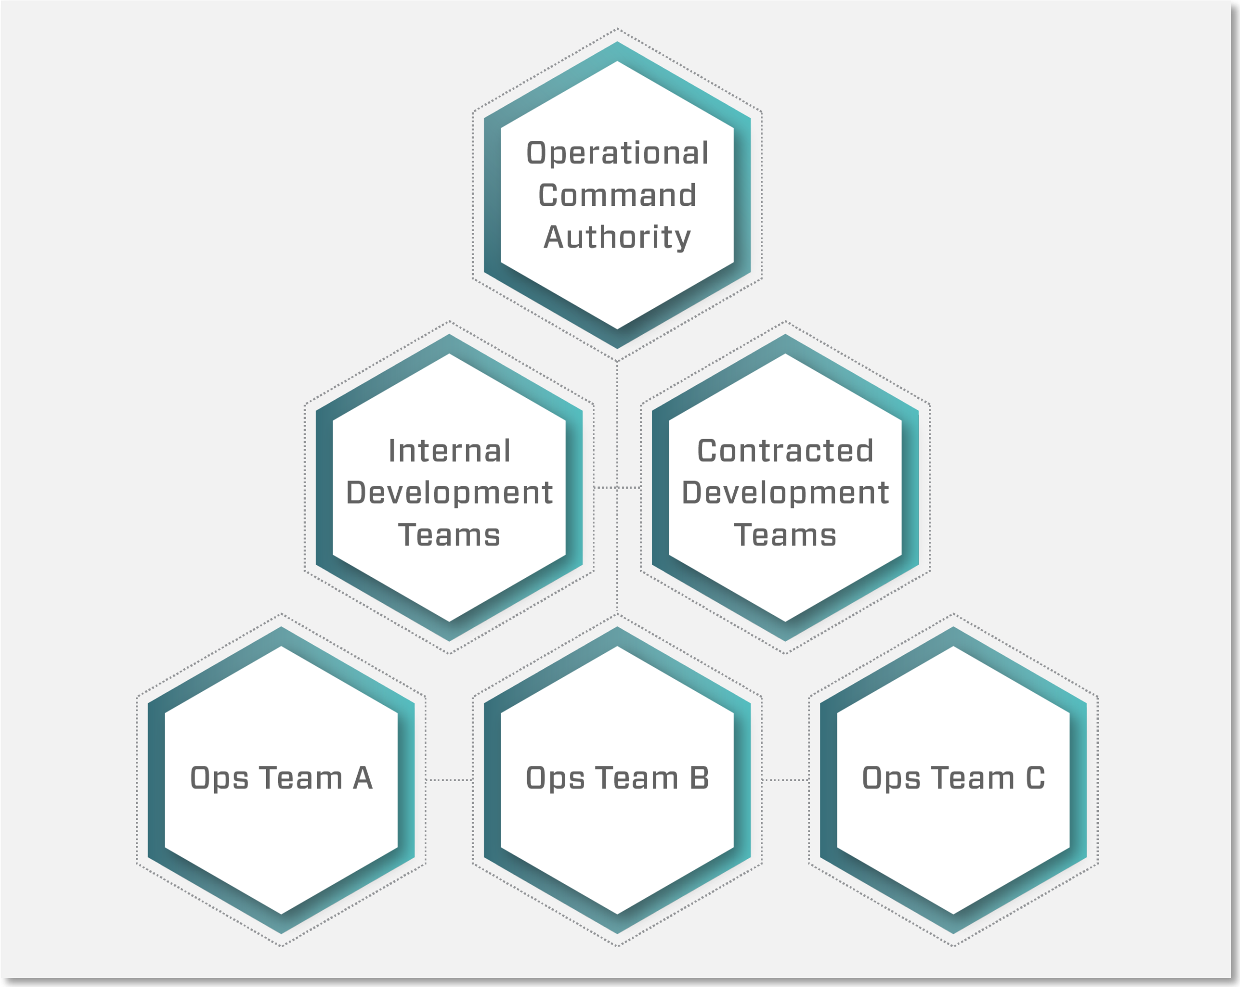
Operational Command Authority
Internal Development Teams
Contracted Development Teams
Ops Team A
Ops Team B
Ops Team C

Figure 4: Operational and Tasking Relationships

As outlined in Figure 4, complex, well-resourced entities (such as state-sponsored activities) feature multiple tiers and centers of activity from operational or national command authority through various tool and capability development shops until finally reaching operational teams. The various entities operating may overlap, follow similar operational and technical guidance, or in some cases (such as distinct "splits" in state power centers, such as NSA vs. CIA in the United States, FSB/SVR vs. GRU in Russia, or MOIS vs. IRGC in Iran) entities serving similar interests or goals may operate in completely different ways.

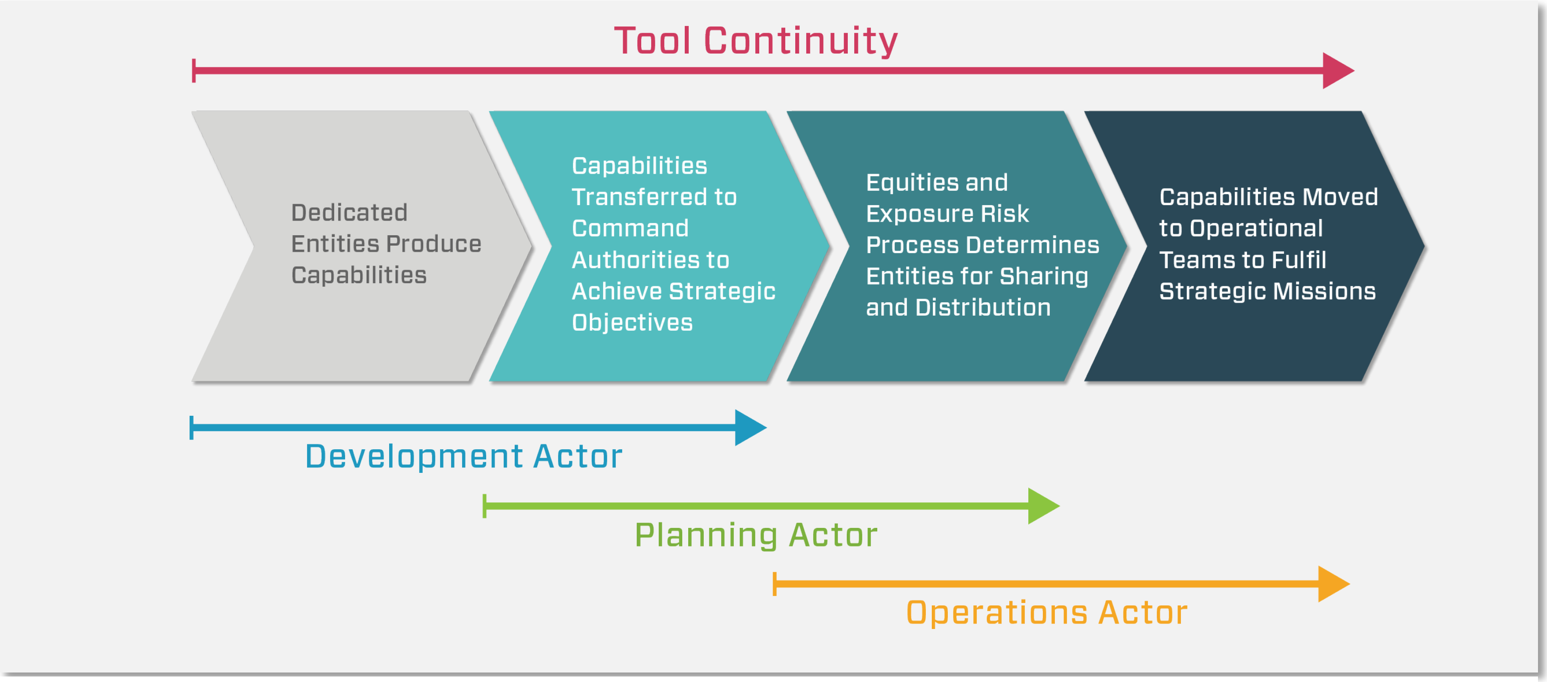
Tool Continuity
Dedicated Entities Produce Capabilities
Capabilities Transferred to Command Authorities to Achieve Strategic Objectives
Equities and Exposure Risk Process Determines Entities for Sharing and Distribution
Capabilities Moved to Operational Teams to Fulfil Strategic Missions
Development Actor
Planning Actor
Operations Actor

Figure 5: Tool Usage Compared to Operational Authority and Execution

While the technical analysis linking groups in efforts such as Animal Farm, Equation Group, the totality of legacy Sandworm reporting, and similar items reflects excellent malware analysis capability, lack of additional context or exposure to precise lines of development, communication, and operation mean specific relationships can be obscured. Looking at something like Equation Group-related tooling, researchers may have identified a common contractor working for multiple distinct entities instead of identifying some overarching, coordinated development program for advanced malware tools shared among partnering organizations. The latter possibility certainly exists, but its provability remains relatively low in likelihood absent additional information beyond malware samples to tie identified tools to specific actors and objectives.

Visualizing a division of labor emblematic of complex, bureaucratic systems means intentions, capabilities, and actions are handed off between multiple discrete entities. Some of these entities may solely support given actors at different levels, while others shift depending on tasking, priorities, and resource requirements. Thus, as shown in Figure 5, specific tool continuity may persist among different actors for a specific operation or goal, but need not hold for future operations that require different coalitions of actors and entities to achieve actions on objectives.

None of this should be construed as saying the binary analysis in work cited above is somehow wrong. Rather, this serves as a caution that the perceived relationships yielded through a single type of data and a single type of analysis can (and almost certainly will) possess holes or permit misconceptions to proliferate. Drawing overarching conclusions from such activity and bundling this into guidance and threat intelligence may therefore broadcast relationships that are far different to what they may seem through a circumscribed field of view.

LOOKBACK AND APT10

Finally, malware analysis is subject to adversaries who may deliberately manipulate tooling or take advantage of potential researcher biases to induce inconclusive or misleading results. While examples of interesting "false flags" exist such as Olympic Destroyer,[21] a more recent and potentially interesting example can be found in phishing activity targeting North American electric utilities in 2019. From August to September 2019, researchers at Proofpoint identified a phishing campaigned they call LookBack, utilizing spoofed network infrastructure and emails to deliver malware to electric utilities.[22]

Specific aspects of this phishing campaign were exceptionally strange as they seemed to completely replicate techniques employed by a different entity, thought to be APT10, that were publicly disclosed in great detail nearly a year prior by FireEye.[23] Examples of near identical overlap include campaign-specific items from 2018 events such as mimicking the Generic Updater (GUP) service used by Notepad++ and using a version of Windows LIBCURL with a malicious exported function to obscure execution. Overall, LookBack appears to be either APT10 completely replaying known tradecraft in a new incident, or a very deliberate attempt to mimic well-known behaviors associated with APT10.

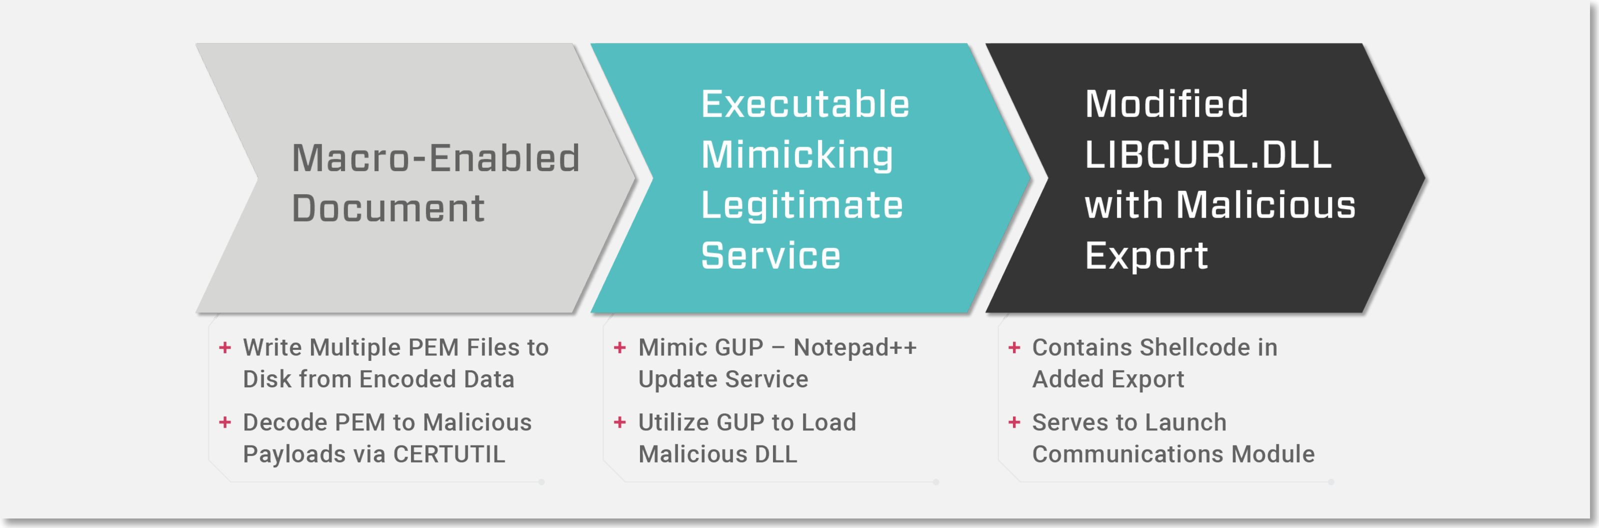
Macro-Enabled Document
Executable Mimicking Legitimate Service
Modified LIBCURL.DLL with Malicious Export
+ Write Multiple PEM Files to Disk from Encoded Data
+ Decode PEM to Malicious Payloads via CERTUTIL
+ Mimic GUP – Notepad++ Update Service
+ Utilize GUP to Load Malicious DLL
+ Contains Shellcode in Added Export
+ Serves to Launch Communications Module

Figure 6: LockBack Exploitation and Execution Chain

This precise adherence to disclosed, legacy tradecraft raises many questions. Presumably, given detailed previous public reporting, an adversary demonstrating past sophistication and dedicated support or development resources to enable campaigns such as Cloud Hopper would not resort to such an obvious replay of known techniques.[24] If we accept this premise (and it must be noted, this is an assumption, and adversaries can certainly be lazy), then some other entity would need to be involved. The resulting possible conclusion would be an entity working to ***look like*** APT10.

Although incomplete given the relatively small number of available samples from this event, there are some subtle indications that the document dropper may share some design and functionality overlap with document droppers previously used by entities linked to North Korean interests from 2016 through 2017.[25] In addition to this potential technical overlap, there is also precedent for North Korean-related entities targeting the US electric sector via phishing with spoofed documents to drop and execute encoded malware.[26] Thus two datapoints emerge which may suggest an alternative explanation for the observed activity.

While threat attribution may seem a lower priority than identifying the attack mechanisms and defenses, from a threat intelligence informing operations perspective it is significant in this case given the substantial differences in likely adversary objectives. For example, APT10 is associated with long-running access and information gathering operations with no publicly known instances leveraging intrusions for deliberate disruption or destruction. Meanwhile, the North Korean linked activity, while also encompassing some traditional intelligence collection, extends to potentially disruptive activity such as cryptocurrency mining and financial theft, and deliberate disruptive behavior such as ransomware or wipers. Thus, attribution based on this information has implications for just what risks a potential victim would face from a potential LookBack-related intrusion, and the appropriate responses for mitigating them.

Overall, this campaign represents a case where it is simply too early (given available information) to make either assessment, at least with a high degree of confidence. To the original reporter’s credit, this uncertainty is documented and reflected in public releases around LookBack. While Proofpoint articulated their reporting with proper caveats and estimative language, in many cases temptations exist to take limited information born solely of analysis of a handful of samples to make claims which may not withstand further scrutiny. Worse yet, in some cases analysis based on small sample sets may become revealed as outright false with additional evidence. Given limited information, threat intelligence based on a single analytical technique and a handful of samples must realize the limitations imposed on any conclusions. Finally, these limitations must be clearly and explicitly communicated to recipients so they can disposition the resulting report and intelligence appropriately.

Defensive Options and Implications

None of the above represent existential threat intelligence issues, but these examples illustrate how analytic tendencies and dependencies can yield less than ideal results. This covers both threat intelligence producers, who must be cognizant of the limitations that narrow, incomplete source analysis place on conclusions, and threat intelligence consumers, who need to assess received reporting considering how it was produced to determine actual relevance and efficacy.

Malware analysis will remain a very important aspect of threat intelligence production for the foreseeable future. But understanding precisely how it fits in to the overall intelligence analysis and production process is necessary to ensure practitioners and consumers do not assign greater confidence to matters than necessary. Thus, defenders can still extract value from even a single source of malicious software – but must appropriately couch analysis, recognizing what inherent limitations are thus placed on conclusions given lack of amplifying information or contextual clues. Use of proper estimative language, confidence levels, and related techniques can allow threat intelligence producers to adequately and accurately disposition results.

From a threat intelligence consumer perspective, the nature of threat intelligence production means that in many cases reports will be based on limited or incomplete data, especially for new or emerging items. Knowing this limitation, especially in the case of single-source analysis such as a malware-only report, can allow defenders to therefore properly assess the strength and applicability of such a report and its findings. Furthermore, knowing this limitation means threat intelligence consumers cannot be merely passive recipients, but must engage in active research, analysis, and correlation of sources to identify potential holes and limitations in received information. This may seem undesirable as essentially reproducing the analysis and production stage of the intelligence cycle. Yet defenders and threat intelligence consumers must apply all-source analysis and cross-source correlation to identify and, where possible, fill in important contextual gaps.

Threat intelligence practitioners must understand the limitations placed on certain types of analysis, accept that limited sourcing or single-analysis examination will leave holes, and report such items with the proper use of estimative language and justifiable conclusions. Organizations can still produce valuable, actional intelligence, but by adding these extra steps and recognizing limitations in techniques, practitioners can generate better reporting at appropriate levels of confidence. Furthermore, while this specific report focused only on malware analysis, this is just an example as similar issues can emerge in pure network infrastructure tracking (especially when pivoting to identify related items), or analysis based on a single incident where many observations may be specific to the victim environment.

Conclusion

The increased availability of malware samples to companies and researchers has energized and broadened the field of cyber threat intelligence, allowing for new discoveries and enhancing defense. However, while beneficial in multiple ways, many aspects of the industry may rely too much on this single form of analysis and place greater confidence in discoveries from isolated malware samples than is justified. Threat intelligence practitioners and consumers must therefore work to understand the limitations of any single-source, small sample size examination, and measure expectations and evaluations appropriately. Furthermore, this concern is neither unique nor limited to malware analysis as a discipline, but applies to all facets of threat intelligence.

When analyzing events or campaigns, threat intelligence professionals must work toward integrating as many data sources and samples as possible to produce high-confidence analysis. Recognizing that this is not always possible, organizations must then ensure appropriate estimative language and confidence assessments are applied to ensure conclusions are sound and that analytical leaps and pivots are clearly identified, underlying logic is justified, and analytical leaps supportable through some evidence. Working in this fashion, the field of cyber threat intelligence can continue moving forward, yielding appropriately beneficial and actionable results to consumers, enabling improved defenses.

Notes

[1] APT1 – FireEye (https://www.fireeye.com/content/dam/fireeye-www/services/pdfs/mandiant-apt1-report.pdf)

[2] One possible exception predating the APT1 report is McAfee's Operation Shady Rat (http://www.csri.info/wp-content/uploads/2012/08/wp-operation-shady-rat1.pdf). Although released prior to the Mandiant report and receiving significant media attention at the time, this report – though significant – has since been eclipsed in the popular imagination by the Mandiant paper.

[3] VirusTotal Tips, Tricks, and Myths – Randy Abrams, VirusBulletin (https://www.virusbulletin.com/uploads/pdf/magazine/2017/VB2017-Abrams.pdf)

[4] Threat Intelligence Defined – CrowdStrike (https://www.crowdstrike.com/epp-101/threat-intelligence/)

[5] What is Threat Intelligence? – RecordedFuture (https://www.recordedfuture.com/threat-intelligence/)

[6] Indicators and Network Defense – Joe Slowik (https://pylos.co/2018/05/16/indicators-and-network-defense/)

[7] ***Mind and World,*** Malcom McDowell

[8] ***Practical Malware Analysis***, Michael Sikorski and Andrew Honig

[9] Any.Run (https://any.run/), Malshare (https://malshare.com/)

[10] The Untold Story of NotPetya – Andy Greenberg, Wired (https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/)

[11] CRASHOVERRIDE: Analysis of the Threat to Electric Grid Operations – Dragos (https://dragos.com/wp-content/uploads/CrashOverride-01.pdf); Win32/Industroyer: A New Threat for Industrial Control Systems – Anton Cherepanov, ESET (https://www.welivesecurity.com/wp-content/uploads/2017/06/Win32_Industroyer.pdf)

[12] Examples include: Cyber Attacks on Ukraine Power and Critical Infrastructure – Marina Krotofil and Oleksii Yasynskyi, S4 Events (https://www.youtube.com/watch?v=lTwsDLO3C44); The Ukrainian Power Grid was Hacked Again – Kim Zetter, Motherboard (https://www.vice.com/en_us/article/bmvkn4/ukrainian-power-station-hacking-december-2016-report); Hackers' Methods Feel Familiar in Ukraine Power Grid Cyberattack – Fifth Domain (https://www.fifthdomain.com/home/2017/01/29/how-a-power-grid-got-hacked/)

[13] Anatomy of an Attack: Detecting and Defeating CRASHOVERRIDE – Joe Slowik, VirusBulletin (https://www.virusbulletin.com/uploads/pdf/magazine/2018/VB2018-Slowik.pdf); CRASHOVERRIDE: Reassessing the 2016 Ukraine Electric Power Event as a Protection-Focused

Attack – Joe Slowik, Dragos (https://dragos.com/wp-content/uploads/CRASHOVERRIDE.pdf); Dragos WorldView customers are also encouraged to review TR-2019-34, CRASHOVERRIDE, SANDWORM, and ELECTRUM and TR-2018-19 CRASHOVERRIDE Attack in Review for additional analysis.

[14] See quotes in: Hackers' Methods Feel Familiar in Ukraine Power Grid Cyberattack – Fifth Domain (https://www.fifthdomain.com/home/2017/01/29/how-a-power-grid-got-hacked/)

[15] New TeleBots Backdoor: First Evidence Linking Industroyer to NotPetya – Aanton Cherepanov and Robert Lipovsky, ESET (https://www.welivesecurity.com/2018/10/11/new-telebots-backdoor-linking-industroyer-notpetya/); The Rise of TeleBots: Analyzing Disruptive KillDisk Attacks – Anton Cherepanov, ESET (https://www.welivesecurity.com/2016/12/13/rise-telebots-analyzing-disruptive-killdisk-attacks/)

[16] The potentially destructive aspects of the 2016 Ukraine event are described in greater detail in the following: Reassessing the 2016 Ukraine Electric Power Event as a Protection-Focused Attack – Joe Slowik, Dragos (https://dragos.com/wp-content/uploads/CRASHOVERRIDE.pdf)

[17] Alert (AA18-284A) Publicly Available Tools Seen in Cyber Incidents Worldwide – US Department of Homeland Security (https://www.us-cert.gov/ncas/alerts/AA18-284A)

[18] Tracking Down the Author of the PlugX RAT – Jaime Blasco, AT&T Cybersecurity (https://cybersecurity.att.com/blogs/labs-research/tracking-down-the-author-of-the-plugx-rat)

[19] Cobalt Group – MITRE ATT&CK (https://attack.mitre.org/groups/G0080/)

[20] Who is GOSSIPGIRL? – Chronicle (https://medium.com/chronicle-blog/who-is-gossipgirl-3b4170f846c0); Equation Group: The Crown Creator of Cyber-Espionage – Kaspersky GReAT (https://www.kaspersky.com/about/press-releases/2015_equation-group-the-crown-creator-of-cyber-espionage); Animals in the APT Farm – Kaspersky GReAT (https://securelist.com/animals-in-the-apt-farm/69114/); Babar: Espionage Software Finally Found and Put Under the Microscope – Gdata (https://www.gdatasoftware.com/blog/2015/02/24270-babar-espionage-software-finally-found-and-put-under-the-microscope)

[21] Who Wasn't Responsible for Olympic Destroyer? – Paul Rascagneres and Martin Lee, Cisco Talos (https://blog.talosintelligence.com/2018/02/who-wasnt-responsible-for-olympic.html); Who Wasn't Responsible for Olympic Destroyer – Paul Rascagneres and Warren Mercer, VirusBulletin (https://www.virusbulletin.com/uploads/pdf/magazine/2018/VB2018-Rascagneres-Mercer.pdf)

[22] LookBack Malware Targets the United States Utilities Sector with Phishing Attacks Impersonating Engineering Licensing Boards – Michael Raggi and Dennis Schwarz, Proofpoint (https://www.proofpoint.com/us/threat-insight/post/lookback-malware-targets-united-states-utilities-sector-phishing-attacks); LookBack Forges Ahead: Continued Targeting of the United States' Utilities Sector Reveals Additional Adversary TTPs – Michael Raggi, Proofpoint (https://www.proofpoint.com/us/threat-insight/post/lookback-forges-ahead-continued-targeting-united-states-utilities-sector-reveals)

[23] APT10 Targeting Japanese Corporations Using Updated TTPs – Ayako Matsuda and Irshad Muhammad, FireEye (https://www.fireeye.com/blog/threat-research/2018/09/apt10-targeting-japanese-corporations-using-updated-ttps.html)

[24] Operation Cloud Hopper – PWC and BAE Systems (https://www.pwc.co.uk/cyber-security/pdf/cloud-hopper-report-final-v4.pdf)

[25] Dragos WorldView customers can review TR-2017-34 COVELLITE Dropper Documents and TR-2018-1 New COVELLITE-Related Dropper Document for additional information on precise technical items.

[26] North Korea Likely Targeted US Power Companies in Spear Phishing Campaign, Firm Says – Krysti Shallenberger, Utility Dive (https://www.utilitydive.com/news/north-korea-likely-targeted-us-power-companies-in-spear-phishing-campaign/507150/); North Korean Actors Spear Phish U.S. Electric Companies – FireEye (https://www.fireeye.com/blog/threat-research/2017/10/north-korean-actors-spear-phish-us-electric-companies.html)