White Paper

Conceptualizing a Continuum of Cyber Threat Attribution

Joe Slowik

Table of Contents

Introduction

Few topics in the field of Cyber Threat Intelligence (CTI) prompt as much passion and debate as the concept of threat attribution. From numerous conference talks, to blogs and papers, to various applications in CTI analysis, the question of threat attribution repeatedly emerges. While CTI attribution discussions can take many forms and aim at specific audiences—for example, policy-makers and state strategy[1]—this discussion will focus on the technical analyst's perspective. In adopting this viewpoint, the question of attribution typically manifests in a very binary fashion.[2] Whereas attribution, as described below, represents various gradations, most discussion limits itself to "yes or no" discussions as to the value and need for CTI attribution, when the actual answer (as with most things in CTI) is, "it depends."

In this paper, a concept of attribution that moves the CTI community away from binary conceptions of CTI attribution value and instead approaches a continuum of attribution types will be introduced. In doing so, multiple possibilities emerge for CTI attributive statements, of different values and significance for different parties—as

[1] "Publicly attributing cyber attacks: a framework" - Florian J. Egloff & Max Smeets, Journal of Strategic Studies (https://www.tandfonline.com/doi/pdf/10.1080/01402390.2021.1895117)
[2] "Achieving Effective Attribution: Case Study on ICS Threats" - Robert M. Lee, SANS ICS Summit (https://www.youtube.com/watch?v=ntBTVUMTFok); "A Brief History of Attribution Mistakes" - Sarah Jones, SANS CTI Summit (https://www.youtube.com/watch?v=Y3EPkDUoGyc)

well as different degrees of relevance for those who wish to make such statements. Through this discussion, the relative value of different types of statements will be examined. Additionally, critical consideration will be applied to why some positions along the emerging continuum of attribution types may be less than desirable for all parties, and ultimately best avoided.

Defining an Attribution Continuum

At present, no shortage exists of discussions on CTI attribution. From brief, general media overviews to multiple in-depth academic works to vendor-specific methodology discussions,[3] various perspectives exist on the subject of how and why to perform CTI attribution.[4] Yet common to nearly all these approaches is a focus on a binary perspective with respect to the value of attribution: whether such activity is desirable, or should be ignored. This view results in a false dichotomy of either an "all in" approach of threat actor attribution or ignoring the subject as either not useful for defenders, or not possible for most CTI analysts.

Resulting debates differentiate between attribution as a way of identifying perpetrators and actors responsible for given cyber operations, and more behavior-centric clustering of identified activity to delineate sets of common tactics, techniques, and procedures (TTPs). This separation of behavioral clustering from "attribution" is reflected in academic literature, where attribution is frequently defined in the following, "who-centric" manner:

> *"[A]ttribution" [is] "determining the identity or location of an attacker or an attacker's intermediary." A resulting identity may be a person's name, an account, an alias, or similar information associated with a person. A location may include physical (geographic) location, or a virtual location such as an IP address or Ethernet address.*[5]

This view of attribution—focusing on the definition of a specific organization or even person behind operations as well as characteristics of that organization—seems far too narrow to reflect operational realities for the overall concept. If we instead adopt a perspective where "attribution" represents the identification or classification of an incident to a given *entity* where *entity* is a flexible identifier ranging from "how-centric" attribution, linking observations to clusters of similar behaviors or TTPs, to very specific "who-focused" attribution, referencing a specific person or organization, multiple possibilities emerge. In adopting this view, analysts approach a range of attribution options which can be mapped to a continuum, such as the following representation:

---

[3] Examples include: "Threat Attribution: Misunderstood & Abused" - MalwareBytes & DarkReading (https://www.darkreading.com/partner-perspectives/malwarebytes/threat-attribution-misunderstood-and-abused/a/d-id/1327915); "Strategies for Resolving the Cyber Attribution Challenge" - Panayotis A. Yannakogeorgos, Air Force University (https://media.defense.gov/2017/May/11/2001745613/-1/-1/0/CPP_0001_YANNAKOGEORGOS_CYBER_TTRIBUTION_CHALLENGE.PDF); "Techniques for Cyber Attack Attribution" - David A. Wheeler & Gregory N. Larsen, Institute for Defense Analysis (https://apps.dtic.mil/dtic/tr/fulltext/u2/a468859.pdf); "Attributing Cyber Attacks" - Thomas Rid & Ben Buchanan, The Journal of Strategic Studies (https://ridt.co/d/rid-buchanan-attributing-cyber-attacks.pdf); "DebUNCing Attribution: How Mandiant Tracks Uncategorized Threat Actors" - Kelli Vanderlee, FireEye (https://www.fireeye.com/blog/products-and-services/2020/12/how-mandiant-tracks-uncategorized-threat-actors.html); "Threat Analytics and Activity Groups" - Joe Slowik, Dragos (https://www.dragos.com/blog/industry-news/threat-analytics-and-activity-groups/)

[4] "Attribution of Advanced Persistent Threats" - Timo Steffens (https://www.springer.com/gp/book/9783662613122)

[5] Wheeler & Larsen

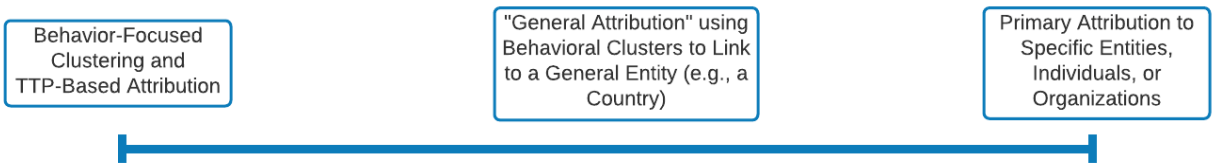

Continuum of Cyber Threat Intelligence Attribution

In this view, CTI attribution does not become a question of "yes or no" but rather "how and to what degree." All organizations, analysts, and CTI vendors therefore perform attribution, but do so to different degrees as outlined in the above continuum of options. By exploring the "poles" of this continuum—Behavioral Attribution on the far left, and Primary Attribution on the far right—we can gain greater insights into the attribution process and its implications.

Behavioral Attribution

The left pole of the attribution continuum focuses primarily on technical observables and evidence. Related items include malware samples, network infrastructure, logs and system artifacts, and similar observations. Residing within technical observations and observed activities, this type of attribution centers on *how* a given event took place while abstracting away from *who* might be responsible. Network defenders as well as CTI vendors and analysts will typically reside primarily in this realm, as the majority of (if not all) available evidence relating to cyber operations typically resides in this region.

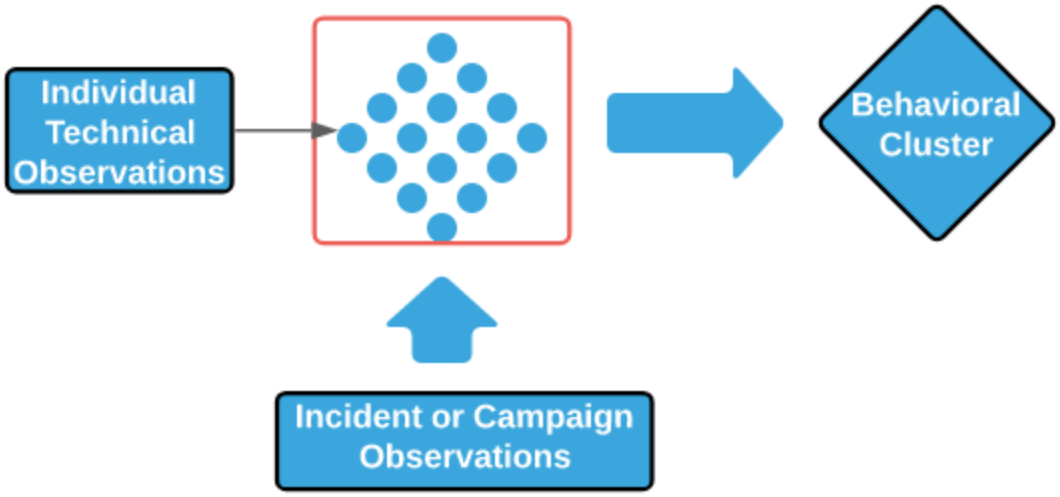

Clustering related activity based on technical observations, behaviors, or TTPs represents a fundamental element of cybersecurity, and is applied by multiple vendors and analysts when generating intelligence products. Central to the concept is remaining faithful to the evidence on hand, and not projecting further than available information allows. In this fashion, Behavioral Attribution—though still quite complex and often difficult to execute in practice—remains both grounded in existing observations while remaining accessible to a variety of parties.

Typical names for Behavioral Attribution include activity groups, behavioral clusters, and similar nomenclature.[6] Importantly, the resulting names and "designators" assigned to the resulting clusters are representative of an analytical methodology—how the observations were clustered—as opposed to any inherent attribute of the responsible party. Thus the names and cryptonyms used by vendors and analysts in Behavioral Attribution are meaningful primarily in the sense of acting as placeholders.[7] An APT number or adjective describing a behavioral cluster simply reflects the collection of evidence and the methodology used to link observations, rather than adhering to the specific entity or party responsible for creating the observations in the first place.

Overall Behavioral Attribution seeks to bucket and group observations as they are linked through technical analysis or observation within technical artifacts. As such, this type of attribution can be quite powerful in grouping similar activities while respecting the boundaries of available evidence. However, such an approach also will, by definition, fail to generate any understanding as to the identity or motivations of the party responsible for a given operation. Certainly there is some scope to infer intentionality—for example, the presence of code designed to delete data, as in Shamoon, or a "ransomware" routine that does not allow for decryption, as seen in NotPetya, gives insight into adversary intent[8]—but further links to purpose, desire, and motivation require inferences that are difficult to defend absent significant non-technical evidence.

Primary Attribution

The right pole of the attribution continuum incorporates significant non-technical data to flesh out information about the *perpetrators* of a given cyber incident as opposed to details and aspects of the incident itself. Unlike Behavioral Attribution, this "who-centric" approach to attribution extends beyond technical observables to incorporate evidence relating to or describing the entity responsible.

---

[6] "The Diamond Model: An Analyst's Best Friend" - Sergio Caltagirone, Dragos (https://www.dragos.com/resource/the-diamond-model-an-analysts-best-friend/); Vanderlee

[7] "Naming, Necessity, and Activity Group Attribution" - Joe Slowik, Stranded on Pylos (https://pylos.co/2018/06/04/naming-necessity-and-activity-group-attribution/)

[8] "Shamoon The Wiper: Further Details (Part II)" - Dmitry Tarakanov, Kaspersky (https://securelist.com/shamoon-the-wiper-further-details-part-ii/57784/); "ExPetr/Petya/NotPetya is a Wiper, Not Ransomware" - Anton Ivanov & Orkhan Mamedov, Kaspersky (https://securelist.com/expetrpetyanotpetya-is-a-wiper-not-ransomware/78902/)

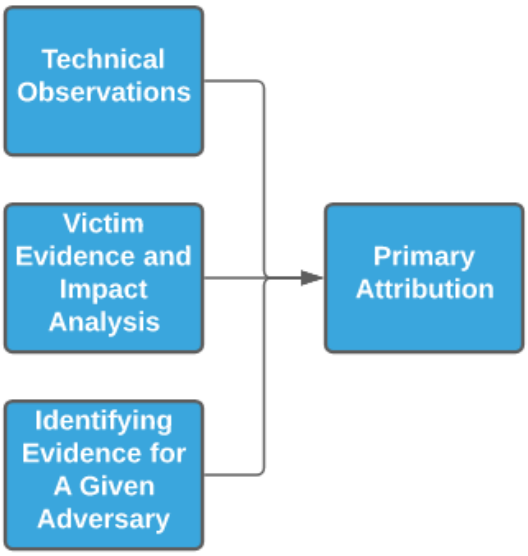

Primary Attribution delves to the source of activity, rather than just its externally observable features. As such, Primary Attribution seeks to determine precisely who is responsible for a given event, potentially down to the identity of specific persons. The information required to make such assessments is difficult to gather, and in many cases not available to private parties. Occasionally CTI vendors are able to gather information allowing for Primary Attribution statements, such as Mandiant's seminal APT1 report linking the identified activity to the People's Liberation Army (PLA) Unit 61398.[9] Other entities, such as Citizen Lab based at the University of Toronto or the Bellingcat organization, may also occasionally reach supportable Primary Attribution conclusions with respect to cyber activity. More often, the required information to link activity to specific organizations or persons resides in capabilities reserved for government use: law enforcement authorities such as warrants or subpoenas; lawful interception of communications; or Signals Intelligence (SIGINT) and Human Intelligence (HUMINT) collection and analysis conducted by intelligence agencies.

Central to these capabilities is the ability to move from technical observation to identification of a persona responsible for that activity. To satisfy the evidentiary needs of Primary Attribution, investigations must move into "who-focused" avenues where private entities not only are ill equipped to operate, but often are prohibited from doing so by law. As a result, independent researchers and analysts often get the best glimpses of Primary Attribution statements through public releases by government entities in the form of indictments, sanctions, or similar actions specifying parties related to a given cyber operation. By correlating technical observations to the publicly documented incident, CTI analysts can perform transitive attribution through this available evidence to get close to Primary Attribution statements.

While Behavioral Attribution remains foremost in importance for frontline defenders, in that such an approach focuses on the actual *how* behind or shaping operations, Primary Attribution, when done well and properly, has significant benefits as well. When correctly performed with proper evidence, Primary Attribution can link events to specific groups or personas which allows for reasonable estimation of intention or purpose. When combined with technical observations, such visibility can be very powerful in enabling defender response. Yet given the evidentiary requirements of Primary Attribution, typically such information (and more importantly, finalized analysis) is not available until well after events have transpired, making this approach valuable but typically not applicable to operational network defense.

[9] "APT1: Exposing one of China's Cyber Espionage Units" - Mandiant (https://www.fireeye.com/content/dam/fireeye-www/services/pdfs/mandiant-apt1-report.pdf)

General Attribution

Finally, while significant "left-" and "right-tail" space exists adjacent to the poles of Behavioral and Primary Attribution—where either some evidence exists to migrate slightly beyond behaviors, or not enough allows for definitive assignment to a specific group—a middle ground is also present. In this case, analysts move beyond pure technical observations, but lack the evidence necessary to link activity to a specific group, person, or related entity.

This middle ground typically represents the most accessible degree of "who-centric" attribution for most analysts and vendors. Yet its value is questionable, along with the possibility of allowing for assumptions and bias to leak into analysis. The easiest way to conceptualize this type of attribution, General Attribution, is through statements such as "this incident is linked to China" or "this event is tied to ransomware operators working out of Romania" or similar generic statements.

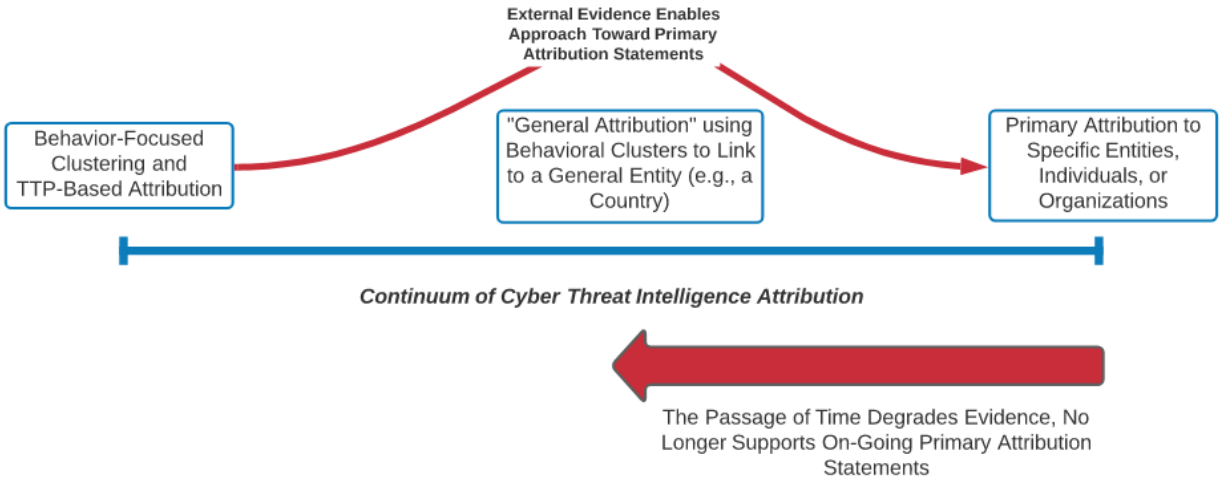

In addition to approaches where limited evidence or projections of intent allow migration from Behavioral Attribution toward General Attribution, it is also possible for attributive statements to degrade over time from Primary Attribution. In this scenario, an analyst encounters external evidence (an indictment, sanctions, or similar primary source evidence) enabling Primary Attribution statements related to specific activity. But, as illustrated in the above diagram, such statements are valid only for the activity covered by the given evidence and its circumstances. Subsequently, such evidence decays toward General Attribution statements where we can link subsequent activity to historical incidents tied to a specific party, but such links are no longer as firm as they were initially.

General Attribution claims identify a geographic region and potentially general interest, but fail to move beyond such statements to precise identification of responsible parties or motivations. While these may seem as on their face useful in providing rough guidance to what entity may be responsible for a given operation, such an approach carries significant and worrisome assumptions. Most importantly, this approach assumes that entities such as "China" or "Romanian criminal elements" are unitary actors with similar (if not identical) motivations and goals,

when all available evidence indicates significant diversity among such higher-level groupings of cyber operators. Just as a continuum exists reflecting degrees of attribution, we can also view state-directed cyber operations as residing along a continuum of motivations and degrees of responsibility.[10] General Attribution to a diverse entity, such as a country, may generate headlines or a sense of satisfaction, but does little to further either tactical or even strategic planning against a given adversary.

Essentially, General Attribution represents a minimally-enriched "best guess" half-way mark toward more robust "who-based" Primary Attribution. Such statements may be supportable and defensible given evidence (e.g., activity identified is ultimately linked to a given state), but the utility of doing so is debatable while the scope for misinterpretation or assumption is concerningly high. While General Attribution statements are accessible to most analysts and organizations, their value, as will be demonstrated in greater detail below, is questionable for all parties with an interest in CTI attributive statements.

Attribution Action Constraints

Attributive actions, although representing some degree of *choice* on the analyst's part, also reflect degrees of *constraint*. Particularly, an analyst may wish to perform Primary Attribution, yet only possess technical evidence of the activity in question. In this case, irrespective of analyst desire, constraints should compel the analyst to focus on Behavioral Attribution. Yet the danger inherent to this situation is a migration to General Attribution as a "good enough" substitute for the rigors of Primary Attribution statements.

CTI analysts and vendors frequently demonstrate a mismatch between choice and constraints along the limitations described above. Most often, we observe this disconnect where analysts desire greater "who-centric" attribution than evidence permits, searching for a way to make statements such as "This organization is responsible for the campaign." Less frequently, analysts possessing significant additional knowledge beyond mere technical observables may omit these items (or the conclusions that can be derived from them), typically to preserve sources and methods of collection for such evidence which are often sensitive in nature.

In the latter case, the broader community is left wanting for more information and detail where such items are actually available. While not necessarily desirable, such an action does not impose any significant cost or lead to potentially incorrect analysis or misconceptions, at least from a pure network defense perspective. Such concerns are, however, inherent to the former, when analysts wish to move beyond constraints to make pseudo-Primary Attribution claims when necessary evidence is lacking. In these cases, a desire to move beyond the evidence available, either through hypothesis or outright guesswork, can lead to outcomes that are indefensible and outright incorrect. Examples would include misattributing identified behaviors to a definite entity, or somewhat less concerningly making such claims (which may nonetheless appear likely) when the evidence simply doesn't exist to support.

That constraints exist on reporting ability and potential accuracy is undeniable, yet overcoming such constraints raises several questions. Items include:

---

[10] "Beyond Attribution: Seeking National Responsibility for Cyber Attacks" - Jason Healey, The Atlantic Council (https://www.atlanticcouncil.org/wp-content/uploads/2012/02/022212_ACUS_NatlResponsibilityCyber.PDF)

- How can we move beyond evidence constraints on attribution?
- How can analysts properly describe activity when constraints limit us to technical details?
- Is it even desirable for analysts to push against evidentiary constraints toward more robust, "who-centric" forms of attribution (the right side of the continuum)?

Moving beyond constraints is easily satisfied: analysts simply require more (and better) evidence. Yet the *type* of evidence necessary is often beyond the reach of private sector CTI operations and usually cannot be distributed or referenced in public reporting from government or law enforcement operations except in anomalous situations. The type of evidence in question relates to "who-specific" findings typically associated with government-run intelligence collection or law enforcement investigations. At times, private sector entities and independent researchers can utilize this information when it is made public, and even more exceptionally uncover such evidence independently. Especially valuable over the past few years are detailed indictments produced by the United States Department of Justice (DoJ). DoJ indictments are typically treasure troves of information that simply cannot be found elsewhere or through other (legal) means, including details such as organizational hierarchies and even the identification of specific individuals involved in operations. In these cases, third-party analysts can leverage this work and, through connection with known technical observables, make links between information in indictments and previously-gathered technical evidence to make statements closer (although not necessarily identical) to Primary Attribution.

When the evidence is simply not available to move beyond constraints, such limitations must be recognized and embraced. In these situations, analysts face the frustrating—but necessary—task of grouping observations purely based on technical observations. Multiple frameworks exist to either assist such activity, such as MITRE ATT&CK,[11] or outright define the process, like the Diamond Model of Intrusion Analysis.[12] In any event, analysts must remain faithful to the information available (technical observations) and not work to move beyond what such evidence can support. This essentially means dwelling on the left side of the continuum, focusing on Behavioral Attribution, linking observations to named clusters of activity as opposed to discrete entities or personas.

Yet in cases where evidence falls short, but desire (or necessity) requires pushing further to the right on the attribution continuum toward Primary Attribution, what avenues are available? Some may argue that going "just far enough" towards "who-centric" attribution to identify a plausible sponsor or directing state authority, even if this falls well short of a specific identity, may suffice to add value. In this scenario, where available information is insufficient to make a Primary Attribution statement such as "Sandworm, identified as Russian GRU Unit 74455,"[13] simply saying "Russian-related" or "likely of Russian origin" would appear to be acceptable. This type of General Attribution, linking to a very broad entity (which consists of multiple subgroups and interests which may not all work in concert with one another), seems enticing at first as a "halfway" point between Behavioral Attribution

[11] MITRE ATT&CK (https://attack.mitre.org/)
[12] "The Diamond Model of Intrusion Analysis" - Sergio Caltagirone, Andrew Pendergast & Christopher Betz (https://www.activeresponse.org/wp-content/uploads/2013/07/diamond.pdf)
[13] Sandworm - MITRE ATT&CK (https://attack.mitre.org/groups/G0034/); "Six Russian GRU Officers Charged in Connection with Worldwide Deployment of Destructive Malware and Other Disruptive Actions in Cyberspace" - United States Department of Justice (DoJ) (https://www.justice.gov/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and)

and incredibly difficult Primary Attribution. But on further analysis, this concept seems to create far more issues than it solves.

The "Mushy Middle" of Attribution

Extremes of attribution activity—behavior-only clustering and externally-enriched Primary Attribution of threat actors— have been established and explored thus far. While there remains significant left- and right-tail space for items that move slightly beyond Behavioral Attribution with additional evidence and items which may fall short of complete Primary Attribution, a significant space exists between the two. In a previous section, we described this as General Attribution—although here we will refer to this area primarily as the "mushy middle."

The "mushy middle" consists of attribution statements which go far enough to name a general party—"China", "Romanian criminal gangs", "Iranian interests"—but that do not go further to concrete entities within these broader groupings. On the one hand, narrowing down activity from the universe of "all possible actors" to a subset such as "Chinese-related threat actors" seems to be somewhat beneficial. Yet any such benefit derives from a (mis)conception of country-level unitary actions with respect to cyber operations, their objectives, and desired impacts.

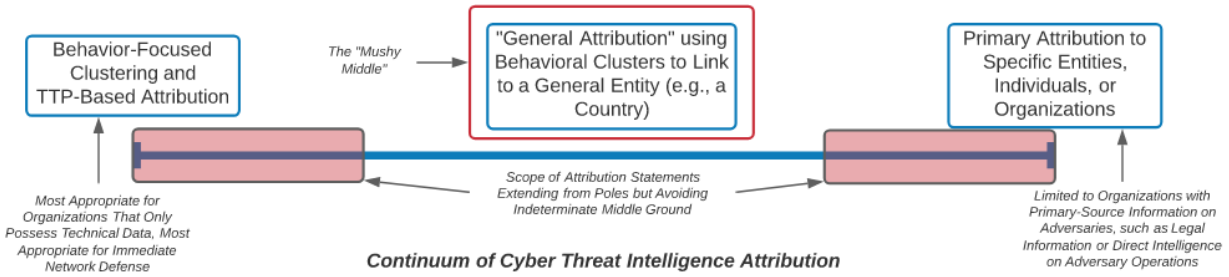

Continuum of Cyber Threat Intelligence Attribution

Further examination and reflection reveals that general groupings—whether "Chinese interests" or "Bulgarian criminal elements"—mask a diverse set of actors with varying motivations, TTPs, and intentions. Lazily lumping activity under such a general label may seem productive at first, but under additional scrutiny such actions are likely to facilitate bias or assumption rather than lead to greater analytical rigor and improved defensive response. To further explore the "mushy middle" and its implications, we will consider several examples.

Reviewing Russian-Related Intrusion Operations and SUNBURST

Several high-profile incidents and resulting statements and analysis from multiple parties highlight a number of worrying cyber operations linked to various Russian entities. As part of this, statements ranging from specific identification of Russian government entities (and even the personnel staffing such agencies) exist, along with more general claims of links to Russian-related "interests." While saying "Russian-related" for a given incident—such as the recently disclosed SUNBURST activity[14]—may appear on its face somewhat beneficial, this

[14] "Continuous Eruption: Further Analysis of the SolarWinds Supply Chain Incident" - Joe Slowik, DomainTools (https://www.domaintools.com/resources/blog/continuous-eruption-further-analysis-of-the-solarwinds-supply-incident); "Highly Evasive Attacker Leverages SolarWinds Supply Chain to Compromise Multiple Global Victims With SUNBURST Backdoor" - FireEye

classification papers over significant differences in Russian-nexus cyber operations. As documented elsewhere, Russian intelligence operations are not only diverse, but at times functioning at cross-purposes or in competition. [15] Such rivalries make unitary attribution at just the level of "Russia" problematic, and this is before considering operational and behavioral distinctions.

A significant history exists for Russian-sponsored disruptive cyber operations, ranging from the 2015 and 2016 Ukraine electric sector events to the 2017 NotPetya incident.[16] Yet the common thread uniting these incidents (and subsequent reporting and analysis) is Primary Attribution. As made clear in public statements, indictments, and sanctions from the United States, United Kingdom, and other countries, these operations were all perpetrated by Russian military intelligence (GRU), specifically Unit 74455.[17] While this may appear so much trivia at first, such information must be put in context of other Russian-related operations that did *not* result in disruptive or destructive impacts, which have aligned to other elements of Russia's intelligence services: either the civilian intelligence agencies, the Foreign Intelligence Service (SVR) or Federal Security Service (FSB), or even other elements of the GRU such as Unit 26165 (commonly referred to as APT28 or Fancy Bear).[18]

(https://www.fireeye.com/blog/threat-research/2020/12/evasive-attacker-leverages-solarwinds-supply-chain-compromises-with-sunburst-backdoor.html); "Customer Guidance on Recent Nation-State Cyber Attacks" - Microsoft Security Response Center (https://msrc-blog.microsoft.com/2020/12/13/customer-guidance-on-recent-nation-state-cyber-attacks/)

[15] "Putin's Hydra: Inside Russia's Intelligence Services" - Mark Galeotti (https://ecfr.eu/archive/page/-/ECFR_169_-_INSIDE_RUSSIAS_INTELLIGENCE_SERVICES_(WEB_AND_PRINT)_2.pdf)

[16] "Analysis of the Cyber Attack on the Ukrainian Power Grid" - Robert M. Lee, Mike Assante & Tim Conway, SANS Institute (https://ics.sans.org/media/E-ISAC_SANS_Ukraine_DUC_5.pdf); "CRASHOVERRIDE: Reassessing the 2016 Ukraine Electric Power Event as a Protection-Focused Attack" - Joe Slowik, Dragos (https://www.dragos.com/wp-content/uploads/CRASHOVERRIDE.pdf); "Six Russian GRU Officers Charged in Connection with Worldwide Deployment of Destructive Malware and Other Disruptive Actions in Cyberspace" - United States DoJ (https://www.justice.gov/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and)

[17] "The United States Condemns Russian Cyber Attack Against the Country of Georgia (February 20)" - United States Department of State (https://ge.usembassy.gov/the-united-states-condemns-russian-cyber-attack-against-the-country-of-georgia-february-20/); "UK exposes series of Russian cyber attacks against Olympic and Paralympic Games" - United Kingdom Foreign, Commonwealth & Development Office (https://www.gov.uk/government/news/uk-exposes-series-of-russian-cyber-attacks-against-olympic-and-paralympic-games); "Sandworm Actors Exploiting Vulnerability in Exim Mail Transfer Agent" - United States National Security Agency (https://media.defense.gov/2020/May/28/2002306626/-1/-1/0/CSA-Sandworm-Actors-Exploiting-Vulnerability-in-Exim-Transfer-Agent-20200528.pdf)

[18] "Six Russian GRU Officers Charged in Connection with Worldwide Deployment of Destructive Malware and Other Disruptive Actions in Cyberspace" - US DoJ (https://www.justice.gov/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and); Official Journal of the European Union, L 351 I (https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L:2020:351I:FULL&from=EN); "Russian GRU 85th GTsSS Deploys Previously Undisclosed Drovorub Malware" - US National Security Agency & Federal Bureau of Investigation (https://image.communications.cyber.nj.gov/lib/fe3e15707564047c7c1270/m/2/FBI-NSA+Cybersecurity+Advisory.pdf); APT28 - MITRE ATT&CK (https://attack.mitre.org/groups/G0007/); "Who is Fancy Bear (APT28)?" - CrowdStrike (https://www.crowdstrike.com/blog/who-is-fancy-bear/)

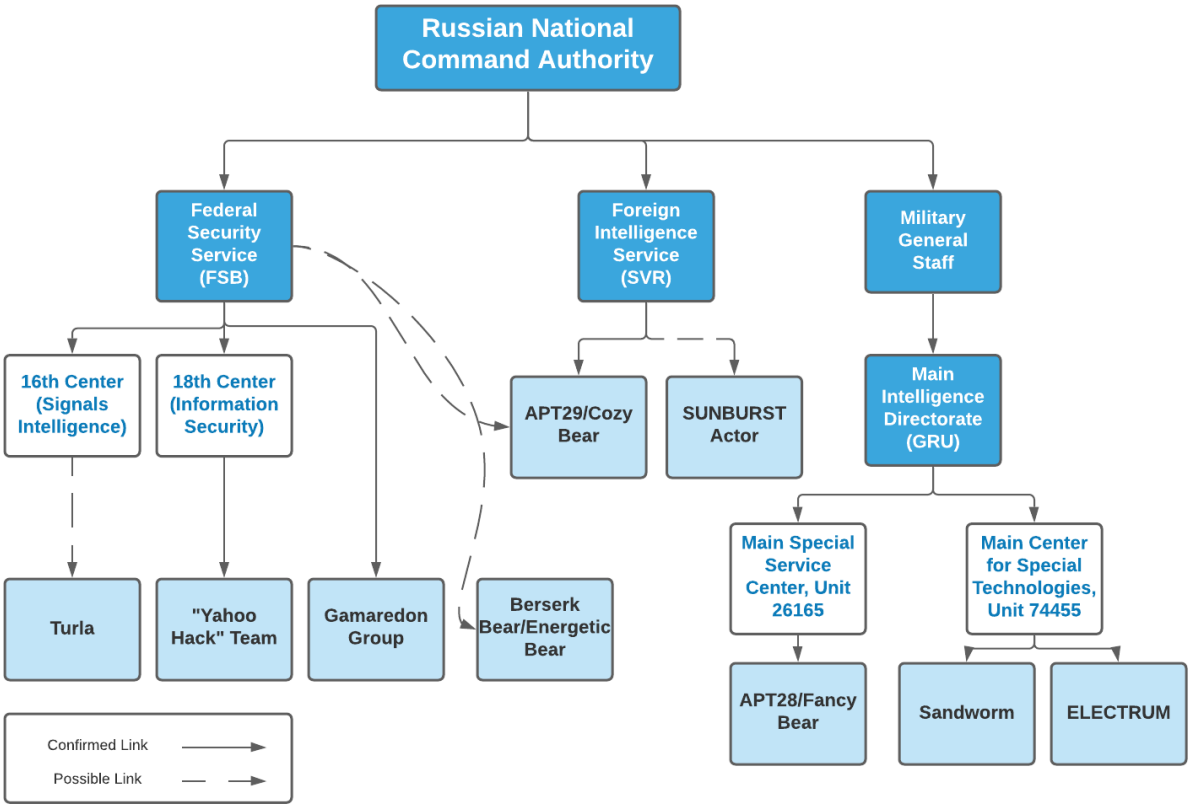

Although at present additional evidence and details are sadly unavailable, multiple leaks from US government officials linking SUNBURST to SVR and *not* GRU activity is significant,[19] even though both groups ultimately take orders from Russian national command authority.[20] The reason is the operational history of these groups. While GRU operations—especially those linked to Sandworm, or Unit 74455—are frequently linked to disruptive or outright destructive operations, no such parallel exists with historical SVR-linked activity. Instead, SVR-linked operations (commonly tracked as APT29 or Cozy Bear, although additional subgroups likely exist) overwhelmingly focus on traditional espionage activities with no evidence of disruptive activity.

In this example, an attribution statement going *just far enough* to say "Russia" is not only unhelpful, but potentially misleading given the differing operational profiles, intentions, and histories of specific Russian-linked threat actors. A General Attribution statement related to SUNBURST operations thus seems less than ideal, and

[19] "Russian Government Hackers are Behind a Broad Espionage Campaign that has Compromised U.S. Agencies including Treasury and Commerce" - Ellen Nakashima &Craig Timberg, The Washington Post (https://www.washingtonpost.com/national-security/russian-government-spies-are-behind-a-broad-hacking-campaign-that-has-breached-us-agencies-and-a-top-cyber-firm/2020/12/13/d5a53b88-3d7d-11eb-9453-fc36ba051781_story.html); "How the Russian hacking group Cozy Bear, suspected in the SolarWinds breach, plays the long game" - Shannon Vavra & Tim Starks, Cyberscoop (https://www.cyberscoop.com/cozy-bear-apt29-solarwinds-russia-persistent/); "The Devil's in the Details: SUNBURST Attribution" - Joe Slowik, DomainTools (https://www.domaintools.com/resources/blog/the-devils-in-the-details-sunburst-attribution)

[20] "Russian Cyber Units" - Congressional Research Service (https://crsreports.congress.gov/product/pdf/IF/IF11718)

perhaps creates more confusion absent additional evidence allowing for a Primary Attribution statement linking the activity to a specific entity under Russian command authority.

Differentiation among actors, taking into consideration their different goals, interests, and likely intentions, is therefore critical in properly dispositioning the SUNBURST and related activity. If such activity could be linked to an entity such as GRU Unit 74455, the commensurate defensive Differentiation among actors, taking into consideration their different goals, interests, and likely intentions, is therefore critical in properly dispositioning the SUNBURST and related activity. If such activity could be linked to an entity such as GRU Unit 74455, the commensurate defensive response and fallout would be dramatically different than if (as seems likely given current information) such actions are instead linked to more traditional espionage operations instead.

HAFNIUM Operations and Widespread Vulnerability Exploitation

Another example of General Attribution statements adding little of value—and potentially leading to misconceptions in understanding and responding to an event—lies in widespread exploitation of CVE-2021-26855 ("ProxyLogon") and related vulnerabilities in March 2021.[21] When initially disclosed by Microsoft on 02 March 2021, the activity in question was limited to an activity group identified as HAFNIUM that paired Exchange exploitation with several privilege escalation vulnerabilities, to both capture email on victimized servers as well as to install webshells for further access. While derived via Behavioral Attribution, Microsoft's original public blog also noted that HAFNIUM is "assessed to be state-sponsored and operating out of China," a General Attribution statement.[22]

[21] "HAFNIUM targeting Exchange Servers with 0-day exploits" - Microsoft (https://www.microsoft.com/security/blog/2021/03/02/hafnium-targeting-exchange-servers/); "Microsoft Exchange Server Remote Code Execution Vulnerability" - Microsoft Security Response Center (https://msrc.microsoft.com/update-guide/vulnerability/CVE-2021-26855); "ProxyLogon" - OrangeTsai/Cheng-Da Tsai (https://proxylogon.com/); "Operation Exchange Marauder: Active Exploitation of Multiple Zero-Day Microsoft Exchange Vulnerabilities" - Josh Grunzweig, Matthew Meltzer, Sean Koessel, Steven Adair & Thomas Lancaster, Volexity (https://www.volexity.com/blog/2021/03/02/active-exploitation-of-microsoft-exchange-zero-day-vulnerabilities/)

[22] "Examining Exchange Exploitation and its Lessons for Defenders" - Joe Slowik, DomainTools (https://www.domaintools.com/resources/blog/examining-exchange-exploitation-and-its-lessons-for-defenders)

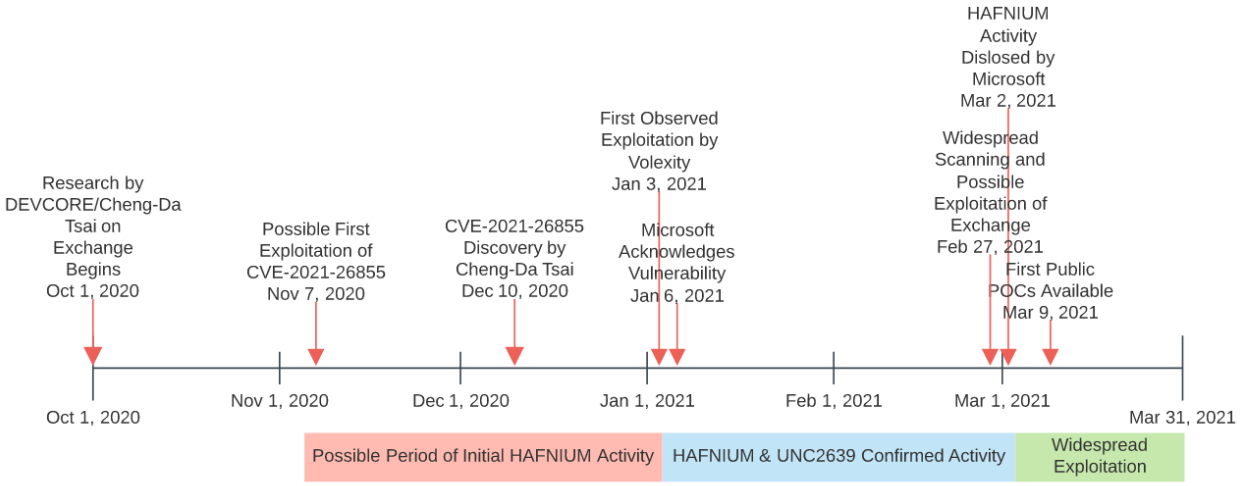

Based on reporting from Microsoft, Volexity, and FireEye, threat actors were exploiting CVE-2021-26855 since at least January 2021, and given other timelines potentially as far back as November 2020.[23] After disclosure, exploitation activity exploded with various media outlets reporting 20-30,000 entities victimized in the U.S. alone within days of public vulnerability disclosure.[24] Most significantly, several additional, distinct entities beyond HAFNIUM started abusing this vulnerability—shortly after public disclosure and in several cases just prior.[25] While non-Microsoft reporting emphasized multiple distinct groups exploiting the Exchange vulnerability from the start media reporting and popular conceptions latched onto the initial General Attribution statement from Microsoft: that CVE-2021-26855 and related Microsoft Exchange exploitation is correlated with HAFNIUM, HAFNIUM is linked to China, and therefore all such activity is linked to Chinese operations.

There are several problems with the logic in the previous statement. First, and as demonstrated in reporting from FireEye and Volexity, it was already debatable as to whether exploitation from January 2021 onward is linked to a single entity. Following public disclosure (and the ability to reverse engineer the Exchange patches), exploitation possibilities rapidly expand as the vulnerability shifts from a "Zero Day" to a "N-Day."[26] Although evidence is uncertain as of this writing, available information suggests that one to *six* distinct groups were exploiting the ProxyLogon vulnerability prior to public disclosure, with multiple potential additional entities opportunistically

---

[23] "Detection and Response to Exploitation of Microsoft Exchange Zero-Day Vulnerabilities" - Matt Bromiley, Chris DiGiamo, Andrew Thompson & Robert Wallace, FireEye (https://www.fireeye.com/blog/threat-research/2021/03/detection-response-to-exploitation-of-microsoft-exchange-zero-day-vulnerabilities.html); "Examining Exchange Exploitation and its Lessons for Defenders" - Joe Slowik, DomainTools (https://www.domaintools.com/resources/blog/examining-exchange-exploitation-and-its-lessons-for-defenders)

[24] "More than 20,000 U.S. organizations compromised through Microsoft flaw" - Joseph Menn, Raphael Satter & Trevor Hunnicutt, Reuters (https://www.reuters.com/article/uk-usa-cyber-microsoft/as-microsoft-email-software-hack-spreads-experts-brace-for-more-impact-idUKKCN2AX23L);

[25] "Four new hacking groups have joined an ongoing offensive against Microsoft's email servers" - Patrick Howell O'Neill, MIT Technology Review (https://www.technologyreview.com/2021/03/06/1020442/four-new-hacking-groups-microsoft-email-servers/)

[26] "The Overlooked Problem of 'N-Day' Vulnerabilities" - Ang Cui, DarkReading (https://www.darkreading.com/vulnerabilities---threats/the-overlooked-problem-of-n-day-vulnerabilities/a/d-id/1331348)

doing so after the patch release and vulnerability identification.[27] Even in cases where use of the chained exploitation of the set of Exchange vulnerabilities first linked to HAFNIUM results in the deployment of the same toolset (such as the widely available China Chopper webshell), diversity in precise tool configuration and other details strongly indicates multiple, distinct entities are engaged in the underlying activity.[28]

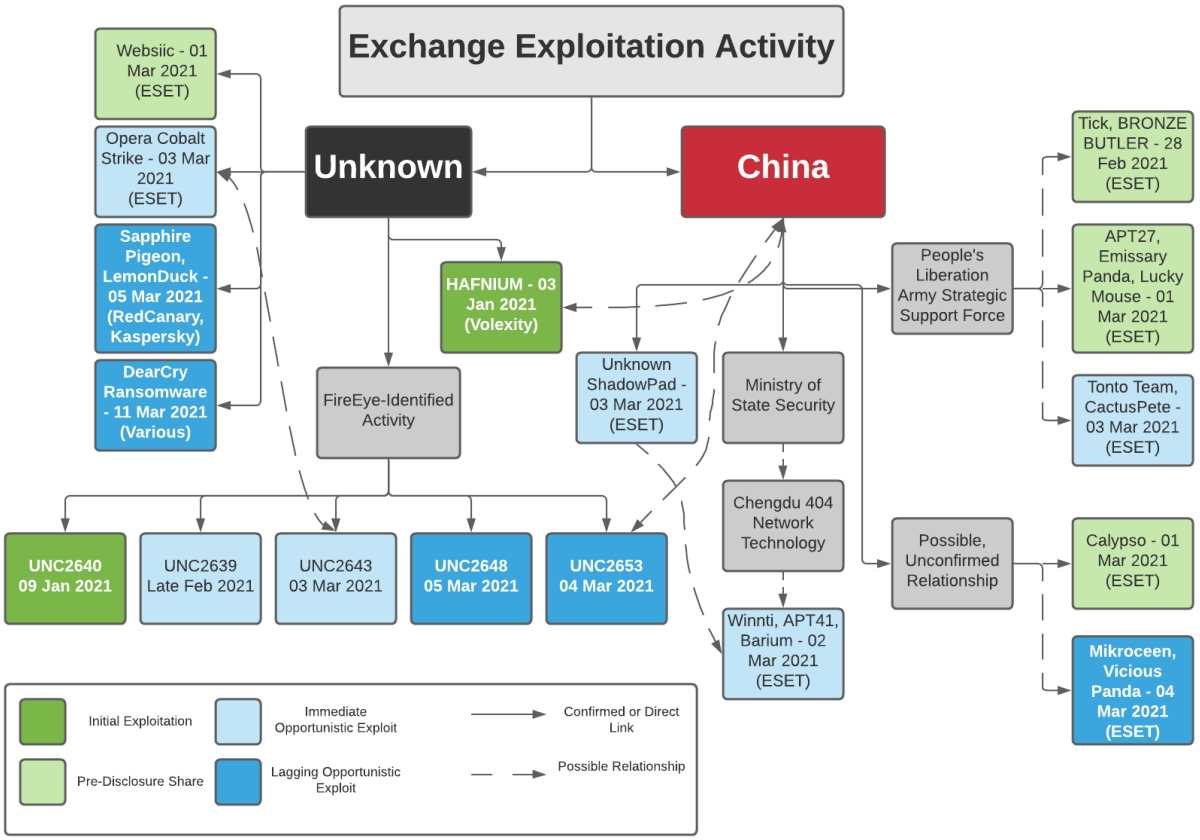

Second, the association of HAFNIUM with "China" raises concerns. HAFNIUM may very well be associated with Chinese authorities or interests, but such a high-level association (to "China" as opposed to anything specific about the People's Republic and its various military, intelligence, contracting, and criminal entities) does little to assist defense, response, or even higher-level strategy. As stated previously, postulating "China" as a unitary object—with uniform behaviors and intentions in cyber operations—simply ignores the reality of diverse entities executing cyber intrusions with ties to Chinese authorities.

---

[27] O'Neill

[28] "Analyzing Attacks Against Microsoft Exchange Server With China Chopper Webshells" - Jeff White, Palo Alto Networks (https://unit42.paloaltonetworks.com/china-chopper-webshell/); "China Chopper Still Active 9 Years Later" - Paul Rascagneres & Vanja Svajcer, Cisco Talos (https://blog.talosintelligence.com/2019/08/china-chopper-still-active-9-years-later.html)

Since Chinese authorities established the People's Liberation Army (PLA) Strategic Support Force (SSF) in December 2015,[29] Chinese-associated cyber activities increasingly bifurcate between traditional espionage operations and military cyber operations focused technical intelligence along with offensive and defensive actions. Typical views of China-linked cyber operations focus on extensive intellectual property theft and espionage activity. The formation of the SSF, and its Network Systems Department (NSD) component, combined former military espionage operations (PLA Third Department) and offensive cyber missions (PLA Fourth Department) into a single entity focused on military specific missions and support. Meanwhile, increased diversity of operations by the Chinese civilian Ministry of State Security (MSS) indicates a possible separation of duties in overall espionage operations between SSF/NSD and MSS from the mid-2010s onward. While the SSF/NSD appears to largely concentrate on military-supporting capabilities in cyber offense and defense, MSS (and its various contractors and regional units) remain focused on espionage and technology theft operations.[30] Although exceptions certainly exist,[31] the overall trajectory of Chinese state computer network operations activity since the SSF's formation reflects an increasingly firm separation of duties between military and civilian intelligence authorities.

Furthermore, in addition to the SSF-MSS bifurcation, a thriving contractor-based, "hack-for-hire" ecosystem also exists within China. Most clearly demonstrated through the confusing cluster of activity referred to as "Winnti,"[32] a complex system of relationships exists between actual operating entities and sponsors or directing authorities. Such tangled links have confounded multiple researchers and analysts, leading to a confusing mess of activity clusters and incomplete attribution. HAFNIUM may fall into this category, making assessment of this group's intentions, motivations, and operational direction more confusing still.

---

[29] "The Strategic Support Force and the Future of Chinese Information Operations" - Elsa B. Kania & John K. Costello, US Army Cyberdefense Review (https://cyberdefensereview.army.mil/Portals/6/Documents/CDR%20Journal%20Articles/The%20Strategic%20Support%20Force_Kania_Costello.pdf); "China's Strategic Support Force: A Force for a New Era" - John Costello & Joe McReynolds, Institute for National Strategic Studies (https://ndupress.ndu.edu/Portals/68/Documents/stratperspective/china/china-perspectives_13.pdf)

[30] "Seven International Cyber Defendants, Including "Apt41" Actors, Charged In Connection With Computer Intrusion Campaigns Against More Than 100 Victims Globally" - US DOJ (https://www.justice.gov/opa/pr/seven-international-cyber-defendants-including-apt41-actors-charged-connection-computer); "Two Chinese Hackers Working with the Ministry of State Security Charged with Global Computer Intrusion Campaign Targeting Intellectual Property and Confidential Business Information, Including COVID-19 Research" - US DOJ (https://www.justice.gov/opa/pr/two-chinese-hackers-working-ministry-state-security-charged-global-computer-intrusion); "Chinese Intelligence Officers and Their Recruited Hackers and Insiders Conspired to Steal Sensitive Commercial Aviation and Technological Data for Years" - US DOJ (https://www.justice.gov/opa/pr/chinese-intelligence-officers-and-their-recruited-hackers-and-insiders-conspired-steal); Chinese Espionage: Operations and Tactics - Nicholas Eftimiades

[31] "Chinese Military Personnel Charged with Computer Fraud, Economic Espionage and Wire Fraud for Hacking into Credit Reporting Agency Equifax" - US DOJ (https://www.justice.gov/opa/pr/chinese-military-personnel-charged-computer-fraud-economic-espionage-and-wire-fraud-hacking)

[32] "What Even is Winnti?" - Daniel Gordon, RiskyBiz (https://risky.biz/whatiswinnti/); "Seven International Cyber Defendants, Including "Apt41" Actors, Charged In Connection With Computer Intrusion Campaigns Against More Than 100 Victims Globally" - US DOJ (https://www.justice.gov/opa/pr/seven-international-cyber-defendants-including-apt41-actors-charged-connection-computer)

From all of these observations, a faulty reading of "China"-associated cyber operations could lead to an assumption that compromised organizations are only of value to the extent they possess valuable intellectual property or similar information. Yet the diversity of Chinese strategic interests and possible operational goals—including the possibility of building up deniable proxy infrastructure through which to conduct future operations, including offensive actions—means that all assumptions of potential Exchange exploitation purpose simply based on a link to "China" are either insupportable or meaningless given numerous plausible options. Furthermore, that post-disclosure opportunistic exploitation of the vulnerabilities in Exchange are now lumped into unitary activity by a single agent engenders additional problems. Given information available at the time of this writing, the Behavioral Attribution statements made by Microsoft, FireEye, and others seem defensible and desirable, while subsequent General Attribution statements by Microsoft and popular media generate little more than confusion with respect to this activity.

The Multitude of Possible Midpoint Errors

The two examples provided—linked to Russian and Chinese cyber operations, respectively—may appear to limit this discussion to large, complex, and well-resourced cyber powers. And yet with minimal investigation, similar concerns arise in other areas as well. The Democratic People's Republic of Korea (North Korea), Iran, and even the United States (and other members of the Five Eyes (FVEY) SIGINT-sharing partnership)[33] do not represent monolithic entities with complete unification over cyber capabilities, missions, and objectives. Rather, these entities represent complex arrangements with multiple units performing missions ranging from traditional espionage collection to internal security operations to pre-positioning for future conflict (commonly referred to as "operational preparation of the environment," or OPE).[34]

Simply saying that a given activity is the act or aligns with the interest of a given country or authority leaves so much space for inferring actual intentionality as to be meaningless at best, or misleading at worst. Although seemingly not as complex, the interaction and overlap between criminal organizations and entities operating in the e-crime ecosystem is not far off from the complexity of state services.[35] In these spaces, as seen in the explosion of actors and service providers facilitating ransomware since at least 2018, an entire ecosystem spanning initial access, subsequent penetration, and ultimate network ransom operations exists.[36] Thus, making general claims about regions or similar does little to shed light on the specifics of adversary operations given the diversity of actors across multiple geographies operating within this space.

---

[33] "This Week in DIA History: Formation of the FVEY Partnership" - US Defense Intelligence Agency (https://www.dia.mil/News/Articles/Article-View/Article/1861392/this-week-in-dia-history-formation-of-the-fvey-partnership/)

[34] "Tactical Cyber: Building a Strategy for Cyber Support to Corps and Below" - Isaac R. Porche III, Christopher Paul, Chad C. Serena, Colin P. Clarke, Erin-Elizabeth Johnson & Drew Herrick, RAND Corporation (https://www.rand.org/pubs/research_reports/RR1600.html)

[35] "2021 Global Threat Report" - CrowdStrike (https://go.crowdstrike.com/rs/281-OBQ-266/images/Report2021GTR.pdf)

[36] "Ransomware-as-a-service: The pandemic within a pandemic" - Intel 471 (https://intel471.com/blog/ransomware-as-a-service-2020-ryuk-maze-revil-egregor-doppelpaymer/); "Analysis of a Cybercrime Infrastructure" - Proofpoint (https://cdn2.vox-cdn.com/uploads/chorus_asset/file/2340876/proofpoint-analysis-cybercrime-infrastructure-20141007.0.pdf)

Overall, General Attribution statements may appear, on their face, to add greater clarity to matters as they seemingly narrow activity down to a generalized actor, but in practice such statements merely give in to inherent biases and preconceived notions on given general adversary tendencies. At the same time, General Attribution does little to actually elucidate specific motivations and intentions for the matter at hand. As a result, General Attribution statements represent so much "junk food" for analysts and external observers, providing some abstracted glimpse of "who" might be responsible for a given activity while lacking the rigor or specificity required for either network defense or strategic planning and response.

Orienting Cyber Threat Intelligence to Defensible Attribution Actions

If CTI attribution remains locked between separate "poles" of Behavioral Attribution and Primary Attribution, with limited ability to migrate from behaviors to specific entities and a likely misleading middle lying in between, what should be the goal of CTI attributive statements? Put exceptionally simply, the goal of attribution within the concepts of operational cyber security is quite succinct: attributive statements should seek to assist network defenders. In this fashion, the "mushy middle" described above fails to reach the necessary evidentiary rigor of Primary Attribution while muddying Behavioral Attribution by tapping into preconceived notions or artificial expectations of state- or group-level conduct as though these were always monolithic, unitary objects.

The above is not to say that Primary Attribution statements are useless for network defenders. They can in fact be quite valuable to the degree that such statements enable direct links to real entities from which we can then reasonably infer intention and purpose. Identification to this level of specificity can cement understanding of adversary behaviors and TTPs, while also facilitating understanding of intruder or attacker goals. When performed properly and accurately, such identification can be invaluable for defensive response and triage. For example, being able to differentiate between a likely espionage campaign and a potentially disruptive operation would be invaluable for response actions to a given breach, as discussed in the SUNBURST example above.

Yet while desirable, the likelihood that such information will be available (given evidence requirements), actionable (due to possible controls over the use and dissemination of information), or timely (since gathering Primary Attribution evidence takes time) will often be incredibly low. Certainly attribution to specific entities (and motivations) is valuable when it can take place, but as evidenced by the months or years-long separation between public statements and actual events in U.S. indictments or the sanctions and public reports of other countries, the likelihood of such information being timely and actionable is quite small.

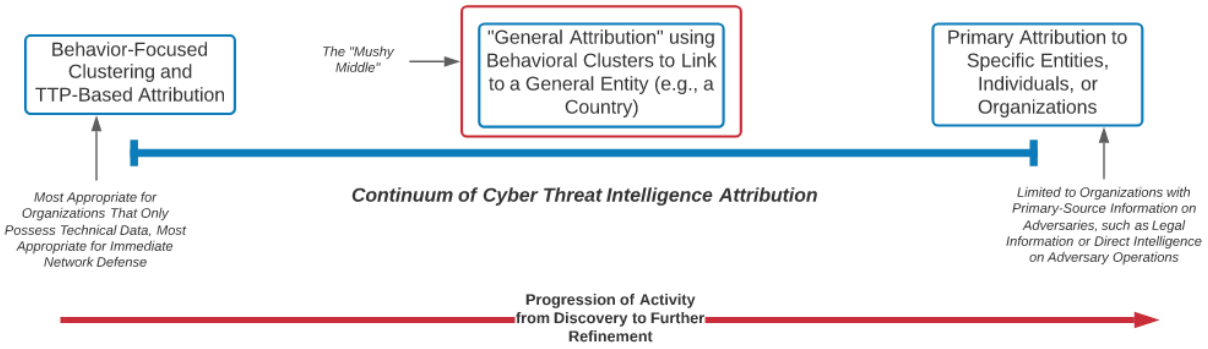

Instead, by remaining within existing and supportable evidentiary boundaries, CTI analysts can ensure attribution claims remain grounded in available data while still providing critical support to network defenders and related parties. Such discipline allows CTI reporting to detail *how* intrusions take place, even if *why* and *by whom* remain mysteries. Furthermore, CTI analysts will also avoid generic, largely unhelpful statements such as "this appears to align with Russian intentions" that provide no appreciable value for defensive operations and may ultimately lead to misconceptions and ill-thought assumptions instead given the diversity of such interests.

While the "mushy middle" may succeed in claiming soundbites, conference presentations, and media mentions, this arena fails to provide much in the way of actionable, useful information for network defenders. Relation to a general "interest" or abstract "intention" simply does not provide enough fidelity into adversary motivation and objectives to engender improved defensive responses. Instead, such an approach generates a superficial feeling of "adversary identification" when such identification resembles little more than the toss of a weighted die.

Furthermore, General Attribution provides little for audiences typically more interested in Primary Attribution statements, such as policymakers and strategic decision makers. For their actions to truly matter, just pointing a finger at "China" or similar does little more than provide a performative statement—actual, concrete action demands narrowing down responsible entities to specific parties. In this way, responses ranging from sanctions to indictments to possible cyber "counterstrikes" are enabled.[37] Failure to provide sufficient detail on such matters to this type of audience therefore does little to enable any impactful or effective action in the realm of cyber operations, deterrence, or response. While "who" can and, depending on the audience, certainly does matter in cyber defensive operations, the evidentiary standard of reaching "who-based" attribution is well beyond that which is expressed by many entities operating in the public sphere.

Given the above arguments. General Attribution represents a lazy midpoint between grounded Behavioral Attribution and more robust (but rarely achieved, and if so even more rarely made public) Primary Attribution activities. General Attribution statements are certainly possible, and potentially (if vacuously) supportable to some degree. Yet given the evidence available to many CTI analysts, the combination of assumptions made in performing General Attribution combined with lack of any noticeable, material value make this an undesirable, pointless endeavor. Where analysts can proceed toward more robust attribution statements defining specific

[37] "U.S. Cyber Command operation disrupted Internet access of Russian troll factory on day of 2018 midterms" - Ellen Nakashima, The Washington Post (https://www.washingtonpost.com/world/national-security/us-cyber-command-operation-disrupted-internet-access-of-russian-troll-factory-on-day-of-2018-midterms/2019/02/26/1827fc9e-36d6-11e9-af5b-b51b7ff322e9_story.html)

groups, organizations, or persons, we should use all tools and evidence in our power to do so. But given the rare situations where this can be achieved, CTI analysts and commercial vendors are advised to remain firmly grounded in making Behavioral Attribution claims which are supported by evidence, which also happen to be the statements of most immediate utility to on-the-ground network defenders.

A Note on Names

One final consideration relates to naming linked to CTI attributive actions.[38] Analysts and external observers are familiar with a diverse series of names from various vendors, organizations, and other parties describing what seems, on its face, to be identical (or at least overlapping) activity. One typical and convincing argument concerning the profusion of threat actor names, including when referring to the same activity, is that analysts should avoid using the terminology of others due to different sources, clustering methodologies, and visibility.[39] This approach makes sense and ties in to the epistemic realities of typical analyst visibility and available evidence. Yet as noted in prior sections, this activity aligns with a certain sector of our attribution continuum, namely Behavioral Attribution based on clustering of observables.

As illustrated below, unless an analyst finds themselves in the improbable position of omniscience with respect to the set of all possible observations of a given threat, campaign, or other action, at best analysts will work with a subset of observations. Furthermore, different analysts, utilizing different sources or possessing different visibility into matters, will have perspectives on matters that, although frequently overlapping, are not necessarily exact—and this is before we delve in to matters of methodology which may result in different clustering actions. In this case, unless third parties have visibility into both the methodology and corpus of evidence used to derive a cluster name, Behavioral Attribution terms *should* be distinct among different analysts.

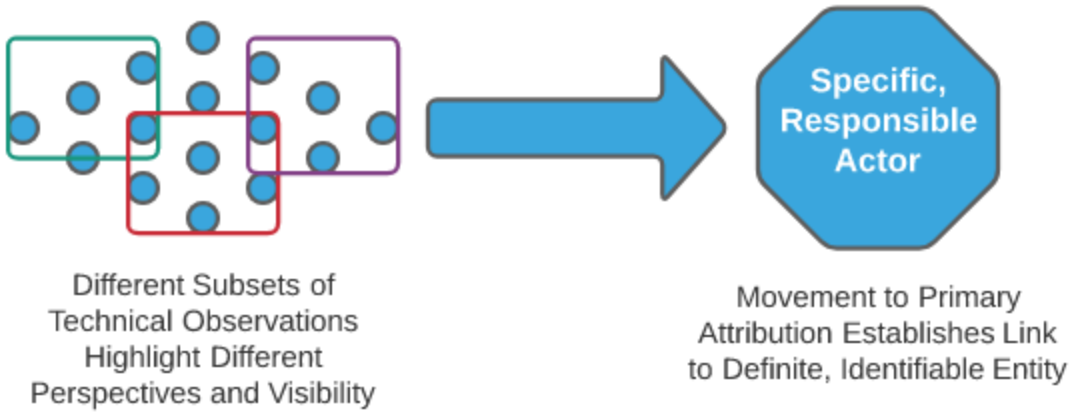

Matters change when we proceed along the attribution continuum towards Primary Attribution. In this situation, names are not linked to varying clusters of related observations, but instead designate a specific entity, organization, or person in the world. In these cases, a profusion of names for *precisely the same thing* is both

[38] "The Newcomer's Guide to Cyber Threat Actor Naming" - Florian Roth (https://cyb3rops.medium.com/the-newcomers-guide-to-cyber-threat-actor-naming-7428e18ee263)
[39] "Making Order out of Chaos: How to Deal with Threat Group Names" - Katie Nickels, SANS (https://www.youtube.com/watch?v=ff1yhdIx0yY)

unhelpful and absurd. As previously documented, CTI analysts will rarely be able to migrate toward Primary Attribution statements, and halfway General Attribution claims are similarly unhelpful as well as imprecise. Therefore the issue of name profusion for specific, existing entities is not common—but it does occur.

To explore further, we can look at the 2018 indictment of two Chinese nationals by the U.S. Department of Justice for intellectual property theft campaigns.[40] Prior to the indictment, CTI analysts and vendors tracked the activity in question under various names such as APT10, Cloud Hopper, POTASSIUM, Stone Panda, and other terms.[41] While such tracking made sense—given different visibility and analytical methodology—prior to the indictment (and its treasure trove of Primary Attribution evidence), after its release the activity was linked to a specific entity: Huaying Haitai Science and Technology Development Company, operating on behalf of or in conjunction with the Chinese MSS's Tianjin State Security Bureau. At this point, the activities tracked as APT10, Stone Panda, Cloud Hopper, etc., coalesce around a specific entity in the world—and the cryptonyms no longer make sense, at least in the context of activity specifically described in the indictment.

For describing the campaigns in the 2018 DoJ document, use of the vendor-specific cryptonyms becomes useful only insofar as it allows us to link the Primary Attribution statements within the document to prior analysis when such evidence was unavailable. A proper way of referencing the activity would therefore resemble something along the lines of, "cyber intrusions executed by Huaying Haitai, previously tracked as APT10," or similar. But continuing to use the cryptonyms exclusively without reference to the named, identified responsible party muddies analysis and enables continuous confusion.

Where matters become less clear are in future research and analysis of activity, similar to previous discussions concerning the "decay" of linked Primary Attribution statements over time. Using the same example as above, Primary Attribution evidence allows for the identification of a specific entity responsible for items previously tracked as Cloud Hopper, etc.—but future activity, absent continued collection and availability of Primary Attribution evidence, will lack equivalent visibility. In these cases, visibility- and methodology-specific names may reemerge as problems of evidence availability and perspective return. However, future statements can at least anchor themselves in prior Primary Attribution statements by noting the historical link to a definite entity, and the *possibility* that future observations reflect a continuation of this link. At this stage, we reach the rightward section of the attribution continuum, but fall short of pure or complete Primary Attribution possibilities. Instead, we can only make *assessments* of a continued link to the named entity in past Prior Attribution work until sufficient evidence becomes available to re-establish the "who"-specific link.

---

[40] "Two Chinese Hackers Associated With the Ministry of State Security Charged with Global Computer Intrusion Campaigns Targeting Intellectual Property and Confidential Business Information" - U.S. Department of Justice (https://www.justice.gov/opa/pr/two-chinese-hackers-associated-ministry-state-security-charged-global-computer-intrusion)

[41] "Operation Cloud Hopper" - PWC & BAE Systems (https://www.pwc.co.uk/cyber-security/pdf/cloud-hopper-report-final-v4.pdf); "APT10 (MenuPass Group): New Tools, Global Campaign Latest Manifestation of Longstanding Threat" - FireEye (https://www.fireeye.com/blog/threat-research/2017/04/apt10_menupass_grou.html); "Two Birds, One STONE PANDA" - Adam Kozy, CrowdStrike (https://www.crowdstrike.com/blog/two-birds-one-stone-panda/)

Conclusion

Attribution is a thorny subject in CTI discussion, but through further examination and analysis we can reveal the concept not as a binary choice but rather a selection along a continuum of options. These range from purely behavior-based analysis and linking on technical information to robust, all-source identification of specific organizations or even individual persons responsible for a given incident.

Although attribution maps to a continuum with a near infinite variety of possibilities, thorough examination of attribution outcomes reveals a mushy, undesirable middle ground between the two poles described above. Due to a variety of factors, many (if not most) CTI analysts and their respective employers will be unable to push very far from the leftward pole of Behavioral Attribution. Yet in desiring to achieve something more akin to Primary Attribution, CTI practitioners must beware a middle ground of General Attribution that provides little real value to either network defenders or higher-level policymakers and strategic audiences, while also introducing assumptions and overly broad classifications that confuse decision-making and follow-on analysis.

Even though the temptation to CTI analysts is strong to make general associative statements linking a given activity or incident to a country or high-level group, further examination indicates these are at best misleading and at worst potentially harmful. Understanding that groups, from state-sponsored cyber operations to complex criminal networks, are not unitary objects but rather contain various interest centers, points of focus, and operational teams highlights that simply saying “this incident is related to Country X” does little to advance the field of network defense on any practical level. By understanding the continuum of attribution choices and the constraints forced on these choices through available data and evidence, CTI practitioners can remain grounded in what sources are accessible while also improving support to consumers of CTI.