THE ETHICS AND PERILS OF APT RESEARCH: AN UNEXPECTED TRANSITION INTO INTELLIGENCE BROKERAGE

Juan Andrés Guerrero-Saade
Kaspersky Lab, USA

Email juan.guerrero@kaspersky.com

ABSTRACT

The top tier of the information security industry has undergone a tectonic shift. Information security researchers are increasingly involved in investigating state-sponsored or geopolitically significant threats. As a result, the affable and community-friendly information security researcher has become the misunderstood and often imperilled intelligence broker. In many ways, researchers have not come to accept this reality, nor have they prepared to act out their new role. Similarly, our industry has yet to gain insights into the complicated playing field of geopolitical intrigue it has set foot into, and as such has fallen into an identity crisis.

Both individual researchers and top-tier infosec firms face drastic changes in embodying their new role as intelligence brokers. Necessary areas of improvement beyond dispute include the enhancement of geopolitical analysis skills and analytical frameworks, coordinated operational security, and strategic decision-making based on a political calculus befitting heightened stakes and disproportionately powerful players. As this new playing field comes into clear view, so will the perils and ethical conundrums that are its permanent features. In the face of investigations with geopolitical weight and consequences, whose final attributions entail unmasking nation-state operations, even the most capable security researcher among us will need drastic preparations, not only to excel but to survive.

INTRODUCTION – RE-SITUATING OUR CONCEPTUAL COORDINATES

In recent years, the information security industry has undergone a tectonic shift as it has embraced 'cyberespionage' research. Research reports on advanced persistent threats (APTs) and targeted attacks (TAs) have become a commonplace offering for high-end outfits and security startups alike, both fighting for the headlines that accompany 'nation-state attacks'. Though the flashy newcomers often lack the visibility or expertise to properly analyse an APT campaign, the top-tier infosec companies have come to pin their legitimacy on intelligence reports. Despite the analytical strength and scrutiny poured into these, the object of study is largely misunderstood.

The terms 'APT', 'targeted attack', 'nation-state sponsored', and even 'cyberespionage' are inaccurate and misrepresent the object of study, which is to say an espionage operation partially carried out with the use of malware. The execution of this complex task is largely determined by a cross-section between the requirements and resources of the attacker, the particular features of the victim systems, and the dynamic and

opportunities between the two. The breadth of interactions that arise therein lend themselves to persistent attacks that are not advanced [1], advanced attacks that are not persistent[1], widely distributed attacks intended for a specific target[2], and targeted attacks with the intention of reaching a wider audience[3]. APT is the generic moniker applied to all of these cases for the sake of convenience and easy marketability, at the expense of accuracy and greater understanding.

Similarly, inexperienced political analysts tilted towards the myopic bias of an intended market audience are wont to botch an investigation already ripe with potential for both misinterpretation and intentional deception on the part of the attacker[4]. Malware is classified as 'nation-state sponsored' regardless of its sophistication based entirely on the stature of its targets and the nature of the data pursued within infected systems. That is to say: 'who, if not the government, is interested in political documents, military dossiers, or industrial control systems?'. One outlier is a subset of attack groups whose interest in political and military secrets is prospectively monetary: with the intention to sell pilfered data to interested parties who may include government institutions but also political opposition, private consulting and political analysts, government contractors, adversarial nation-states, as well as corporations, utilities providers, financial speculators, and a diverse array of institutions whose interests overlap due to governmental regulation and intervention. Further focus on the malware's deployment onto specific systems or its design for stealing specific document types or specific strings may strengthen our conviction that the attackers are after a specific type of data but it gets us no closer to understanding who the final recipient of the pilfered data may be. This goes to say that there exists enough ambiguity to disqualify the term 'nation-space sponsored malware' as vague, if not purposefully inaccurate.

The purpose of this extensive disambiguation is to get at the root of the inadequacy of the broader category of 'cyberespionage' at the level of genuine nation-state attacks. When discussing the operations of a professional institution whose primary activity is intelligence gathering and production, cyberespionage as colloquially defined by the infosec community is largely meaningless. 'Cyber'-capabilities for data gathering are an evolution of the adoption of technological sleight-of-hand that intelligence agencies have spear-headed in order to run uncompromised agent networks as well as more efficiently generate and exfiltrate data. To define espionage by its material means may be momentarily useful insofar as it speaks to attacker capabilities and specific features of the data gathered but to subsequently allow this narrow lens to define the operation as a whole is a mistake. Analogously, it is partially useful to discuss the benefits of concealed microphones in espionage operations but to understand espionage operations through the practice of

---

[1] With malware residing in memory and leaving minimal footprint on disk, as in the case of Duqu 2.0 [2].
[2] Examples include Darkhotel components [3] and Animal Farm's Nbot [4] intended for DDoS.
[3] Regin's use of legitimate institutions as domestic network proxies [5] or Duqu's utilitarian targeting for digital certificates [6].
[4] Common examples of attributory data manipulation centre around language strings and timestamp manipulation. An example of the former is MiniDuke's attempt to hide Russian-speaking developers with consistent use of English [7] or extensive use of 'red herrings' by CloudAtlas/Inception Framework [8].

microphone concealment is bound to lead us astray. What we have come to refer to as cyberespionage continually falls into this trap as analyses are seldom situated in the context of espionage proper.

In order to further break apart these unconducive trends of thought and explore the role of the security researcher investigating espionage operations that employ malware, I will situate the production of threat intelligence reports in terms of an emergent parallel appropriation of the intelligence production cycle already at the heart of intelligence agencies, I will then explore the perilous outcomes that accompany the shortcomings of this appropriation with particular interest in the need for institutional strategic thinking, and finally describe the features of full-hearted threat intelligence brokerage along with the new weak points that organically arise in a mature cyber-threat intelligence market.

A PATCHWORK INTELLIGENCE PRODUCTION CYCLE

Espionage understands the fundamental difference between information and intelligence. Intelligence is the product of coalescing diverse sources of both privileged and readily available information with the intention of empowering a decision-making agent in an information-asymmetrical space. Malware deployed for espionage purposes gathers all sorts of data but, unless its target is an intelligence agency (IA) or its customers, what the malware exfiltrates is largely raw data. In pseudo-algebraic terms, information represents the input for a function whose product is intelligence. That function itself is constituted by a procedural methodology that situates that information in terms of the intended intelligence consumer's requirements. Thus the need for analysts of great breadth cannot be circumvented nor can a product be prepared without insight or concern for the requirements of the ultimate customer.

A complex strategic calculus

Intelligence agencies are required not only to have ready access to privileged information but also a significant number of analysts[5] on hand and the strategic foresight to determine what the customer should ultimately have access to in fulfilling the request at hand. Intelligence agencies are singular in their understanding of the importance of these last two points, dedicating vast resources to ascertain both premium analytical and strategic prowess. There are some publicly available resources that expound on intelligence analytical methodology, but little has been written about the strategic process that determines if, how, and when the product of that methodology is delivered to the requesting customer while observing the prospective benefit of all parties involved – a *strategic calculus* necessary in order to manage manifold overlapping power-asymmetric relationships with no temporal delimitation. In lay terms: how to manage the sharing of secrets between many powerful people, all of whom know each other, when the relationships have to be

mutually beneficial for an unforeseeable length of time. This inexact practice, more intuitive art of intelligence retro-feeding than science, is of absolute importance in the long-term survival of the agency, its operational capabilities, and the well-being of the customer as well.

An intelligence agency's procedural methodology involves diverse methods of information gathering, expert analysis of the data gathered to be represented in terms relevant to the customer's requirement, and finally strategic filtering to ensure the well-being of all parties involved. This last element boils down to determining not only what the customer requires but what he should be entitled to receive regardless of its availability to the agency so as to determine how this action will affect long-term operational capabilities, regional dynamics, agency visibility, and whether the effect the actionable intelligence will have on the customer's ultimate decision is desirable both in terms of the customer's and the agency's priorities.

An incomplete appropriation

Private company security research teams studying the malware remnants of espionage campaigns are unwittingly following a similar procedural methodology. When it became popular to market APT research in the form of cyber-threat intelligence reports, 'infosec' companies took a firm step forward into the arena of intelligence brokerage. This move has proven profitable and shows no signs of relenting as the market becomes increasingly populated by pay-per-report, subscription feed, and partnership arrangements offered by companies young[6] and old. However, few teams (if any) have fully embraced their role beyond the purely technical in order to properly operate in an intelligence brokerage space. As the product becomes 'cyber-threat intelligence', researchers have begun an awkward, piecemeal emulation of the production cycle of an intelligence agency without recognizing it as such. Without that necessary self-awareness, researchers lack the intention and leverage to curb public relations and marketing department priorities that shape the company's ultimate offering of the research product. This has led to the repeated perpetration of cardinal mistakes in intelligence distribution.

To explore the shortcomings of the privatized intelligence production cycle, a fundamental asymmetry between the players involved must be acknowledged. Though private research teams and intelligence agencies will follow similar intelligence production cycles[7], we must not conflate their attributes. (1) Intelligence agencies benefit from cover for action, meaning that other governmental institutions do not find the agencies' intelligence production activities suspect. (2) Agency employees enjoy legal protections, even those involved in network exploitation activities. And finally, (3) their work is shielded from political blowback or geopolitical incongruousness. Each point is inversely applicable to security researchers and thus sets the tone for the power asymmetry:

1. Security researchers enjoy no cover for action for their production of intelligence reports into what may or may not constitute legitimate intelligence operations. Private sector researchers display an academic interest

---

[5] The numbers are inflated by the breadth of areas of analytical expertise commonly required, which often include: politics, legal and regulatory matters, military and diplomatic affairs, finance and economics, diverse technical specifics like nuclear and biochemical equipment, and regional dynamics between legitimate players and political parties as well as shadow players like gangs, paramilitary and terrorist groups, and narco-agents.

[6] Many of these lack the visibility and technical expertise to provide any such offering but do so from their inception specifically to serve an over-eager market.

[7] RGASD cycle: Request, Gather, Analyse, Strategize, and Deliver.

in understanding advanced uses of malware for espionage operations and also derive career benefits from partaking in attention-grabbing announcements. In the eyes of a counter-intelligence branch, that alone does not justify 'snooping' into and potentially disrupting government sanctioned operations. The issue is more clearly displayed when divorced from the realm of 'cyber': if private company employees decided to tail detectives as they in turn surveilled the targets of an FBI investigation, then subsequently notified the targets and published or sold this information under their company brand, this activity would find little protection from counterintelligence, political, and legal blowback. The same lack of cover for action applies to the realm of digital intelligence research but is made less clear by an illusion of divorce from severity granted by the medium and its inexact terminology.

2. Security researchers are afforded no explicit legal protections for the grey areas regularly visited throughout the course of an investigation. For example, is the researcher committing a crime when accessing a command and control server (with or without the need for exploitation or overt manipulation beyond the use of credentials stored in a malware sample)? Similarly, is the researcher in possession of stolen property when accessing data exfiltrated by the attacker from his targets? Though common sense views side favourably with researchers, legal ambiguity remains with little solace afforded.

3. The companies too lack a cover for action and are in no way insulated from the political blowback that arises from the public disclosure of sensitive operations. They suffer from a further dimension of 'guilt by association' as research into sensitive operations and subsequent reporting is misconstrued as an act of geopolitical aggression when the victim and perpetrator are involved in any form of international tension. A global company P associated with country X publishing on country Y's operations is considered a proxy aggressor, regardless of the extent of company P's international operations or research integrity. Prejudice finds easy expression in the form of sanctions, barring from public contracts, and disincentives for local companies to trade with company P. Similar prejudice takes actionable form on an individual scale as company P's researchers originating from or legally residing in country Y are antagonized, with natives labelled traitors and legal residents facing threat to their legal status and livelihood in the absence of the protection of native citizenship. The effect of this industry Balkanization is well underway at this time.

Having delineated the intrinsic asymmetry between private research teams and intelligence agencies, I can proceed to elucidate the innate stumbling blocks that shaped how private industry researchers emulate the intelligence production cycle. Private research procedural methodology emerges from the peculiar circumstance of hunting for indicators to interpret from the perspective of the perennial outsider. The following is a rundown of the private sector's emergent parallel to the RGASD cycle:

Request – Private research begins either with a vague request for incident response to assess the nature of a suspected

breach or with the analysis of a decontextualized sample. In some cases, entire campaigns are uncovered on the basis of readily available malware deposited in watering hole sites or collected in malware repositories like *VirusTotal*. Here we encounter the first significant difference between private and IA cycles – a lack of adequate scope definition. The IA's production cycle relies on the customer's request to define the scope of the inquiry and the ultimate shape of the result. The researcher's entry point lacks similar clarity of purpose. Incidence response could be fulfilled by disinfection, technical indicators, and an understanding and remediation of the initial infection vector so the intelligence-gathering activities that arise from here are not necessarily aimed at serving the initial breached customer. The production cycle will enrich the research team's situational awareness and may thus serve future customers. In that vein, the actionable potential of the results of the investigation remains undefined.

Gather – The gathering stage is a largely technical practice varying per company and research team. In contrast with the agencies, closed HUMINT sources are seldom centrally employed. The process usually entails reversing samples and protocols, exploring command and control infrastructure, profiling victimology and data exfiltrated, and correlating distinguishing peculiarities both with open sources as well as data from other campaigns suspected of shared provenance.

Analyse[8] – Despite recruiting a wide array of practitioners, dedicated analyst positions in threat intelligence are still largely reserved for candidates with strict computer science backgrounds. Despite the arguable importance of this qualification in the gathering stage, its broader imposition is a detriment for the purpose of analysis. As a direct result of a lack of appropriate speciality segmentation, analysis in the threat intelligence space is politically weak and oversimplified. Hypotheses are treated as fact; countries are discussed as entities of singular composition with predictable motivations. Analytical prowess is further hampered by sectoral pandering which dissuades researchers from veering too sharply in the direction of hypotheses unpopular or uncomfortable to desirable sectors of potential customers. As a direct result, attribution problems abound as companies either attribute with comical predictability[9] or abstain from the matter of attribution altogether even in the face of overwhelming proof or admission by the perpetrators themselves. At this time, the process of analysis is limited to sparsely substantiated motivational conjectures.

Strategize – Departmental outsourcing means that the strategic stage is neglected or non-existent in the private sector cycle. Strategic prioritization is outsourced[10] to sales, marketing, and public relations departments whose quantitative metrics cannot account for the possibility that a sale may not have a positive effect. The IA cycle involves the exercise of agential regency over the customer on the basis of the value of the intelligence promised relative to the action

---

[8] Companies whose offering is limited to a 'threat intelligence feed' circumvent this step altogether as the content is typically the result of extracting indicators through automated static and dynamic analysis that is charitably considered part of the previous gathering stage.
[9] Partially as a result of the cases they choose to investigate in the first place.
[10] Further strategic outsourcing enlists governmental priorities as sensitive investigations are foregone despite customer requests because of pressure from presumably unrelated government institutions.

potential of the customer in light of the findings. The decision to leave the customer request partially or entirely unfulfilled is reputationally problematic for the IA but ensures longevity of operational viability. Similarly, foregoing a sale is anathema to the short-term business priorities of the private company but strategically ensuring mutually beneficial effects to every venture contributes to its long-term viability. At this time, the closest example of strategic thinking applied to the production of intelligence reports is the partial withholding of sensitive details[11] from public reports.

Deliver – The final stage is truly problematic, as the private sector regularly commits grievances unthinkable to their intelligence agency counterparts by way of public releases. For the IA, the final delivery stage means handing over directly to the partner a product that reflects the omissions resulting from the strategic calculus and accurately adjusted to the necessities delineated by the scope of the customer's request. Furthermore, the IA can then approach trusted partners with elements resulting from this investigation that will further advance these strategic partnerships (with no immediate *quid pro quo* expectation).

In the private sector, though investigations may pertain to specific customer requests, the resulting reports are often promulgated to a wider reach including (a) established (but unwitting) victims and (b) potential tasking targets in the affected sector or discerned tasking pattern. Going further, once the report is expunged of details that may identify the original customer, it is then (c) released for public consumption. The intended purpose is a PR-coup to both attract new customers for closed-release intelligence reports as well as garner brand recognition and industry respect for formidable findings.

The previous labels (a, b, c) point to grievances in intelligence relations that amount to an act of aggression mounted by the private firms against the perpetrating IAs. Translating these common-place private sector actions to their equivalents from the perspective of the IA follows:

(a)   The already problematic partial loss of visibility of an operation is expanded unnecessarily to include other targets, perhaps to its furthest extent of shutting off all visibility for a given operation or multi-tasking[12] platform. This also involves the partial exposure of the operation to the target itself, a move with potential legislative and geopolitical repercussions that threaten the standing of the perpetrating agency.

(b)   Spooking other members of the affected sector or tasking pattern amounts to proactively shutting off access to what may be necessary targets as the operation unfolds.

(c)   The sanitized public release is a way of not only burning the investment in the attack platform but also seeking to embarrass the IA, whose very discovery is a failure and every detail exposed (tools, methods, victims, etc.) a strike against its competency. Adding

---

[11] These include specific victimology and key malware features or infrastructure details that serve as leads for future investigations.
[12] To clarify, the term 'multi-tasking' refers to the use of a single malware platform for the tasking of multiple operations in a manner largely indistinguishable to the researchers that encounter malware of the same family and provenance.

to this problematic circumstance is the possibility that the intelligence customer will be exposed (perhaps simply from inference based on the selected targets). Exposing the customer is the ultimate intelligence grievance.

The cardinal grievances herein described threaten the operational viability of the IA as an institution. This adversarial perception created by the choices of delivery in private research will directly contribute to the following discussion of the perils that accompany private research at the hands of the antagonized IAs.

This larger rundown illustrates the private industry intelligence production cycle for investigating malware-based espionage operations. Shortcomings are the product of employing 'jack-of-all-trades'-style intellectual resources and endemic strategic failures commensurate with a naive lack of understanding of the players perpetrating the attacks under investigation. These circumstances have created a perilous and ethically tense situation for security researchers.

FOR LACK OF A GRAND STRATEGY

The outsourcing of decision-making to departments with purely quantitative priorities creates a strategic vacuum that proliferates varying perilous and ethically challenging situations for private research teams, both as individuals and companies. I will first discuss these perils and ethical tensions and then describe a position to fill that power vacuum.

Living in fear…

Researchers in the thick of sensitive research know that they often tussle with truly powerful actors. Where research into cybercrime stands on solid ground as an informative and welcome extension of law enforcement, espionage research has little precedent and, as previously discussed, flimsy cover for action and questionable legal protection. If researchers receive personal insults and death threats for disrupting the operations of Brazilian cybercriminals, what might nation-state actors of varying scruples be willing to do to stop a loud, boisterous, profiteering public nuisance? Researchers face a variety of perils as they carry out their work. Their companies too face a different set of difficulties as the result of hindering these actors. As such, we can divide the perils faced into two categories: those arising from the nature of the research as an individual's intervention into an always already sensitive operation and those arising from the awkward and unwelcome intervention of a company in the geopolitical space. These two levels are best discussed separately:

Individual

The researcher as a private individual faces unique challenges when in the cross-hairs of a nation-state actor determined to enact some form of retribution. The operator of an espionage campaign is not a common criminal nor a simple citizen and his resources are truly manifold. As a special class of government insider responsible for a sensitive operation, the attacker can go so far as to legitimize special recourse in order to neutralize the threat posed by the meddling security researcher. The options available slide relative to the nature of the attacker, ranging from civilized to unscrupulous, and include: subtle pressure, patriotic enlistment, bribery, compromise and blackmail, legal repercussions, threat to

livelihood, threat to viability of life in the actor's area of influence, threat of force, or elimination. Though a process of escalation may be evidenced, there is little incentive to obey a specific progression between these rungs. Pragmatic actors will enact the path of greater assurance to their goal of continued deniable operational viability. To clarify, the intent is not that of demonizing all intelligence services but rather to discourage the rampant assumption of operational homogeneity amongst these complex and variegated institutions. Intelligence services vary across continents and are shaped by the requirements of their geopolitical context – the unacceptable measures of some are the elegant solutions of others.

Some of the threats described above deserve further attention, in particular *compromise*, *threat to livelihood and living conditions*, and *elimination*. The topic of compromise and blackmail has received greater attention in recent years throughout the infosec community and the conference circuit. It is usually accompanied by imperatives of a need for greater operational security (OPSEC). The premise is essentially that blackmail and pressure on the basis of secrets, debt, and shameful proclivities and missteps are inexpensive ways of 'owning' a person commonly employed by intelligence services. They also have the added benefit of enlisting the reluctant victim into convincing surrounding onlookers of the intended deception, motivated by the fear of their compromise becoming known. This type of compromise is in some cases related to the threat to livelihood as private information security companies have displayed a more or less strict moralism in their hiring practices, often preferring practitioners untainted by publicly known blackhat tendencies. This means that the compromised victim is in a double-bind since the security services can then ruin the researcher-cum-victim's career by 'leaking' knowledge of their collaboration. Where the researcher has not been compromised, association can be insinuated or fabricated to varying degrees of effectiveness.

When the researcher resides in a space within the sphere of influence of the campaign operators, threat to living conditions is a viable option and an effective display of the powers of the attacker-as-government entity. If the researcher is a non-native legal resident, security services can threaten to (and effectively) revoke that resident status, thus forcing the researcher to relocate, in some cases separating families or forcing a return to dreadful conditions. With inflexible employers this may also involve a threat to livelihood. If the researcher is a native citizen, these threats often involve betrayal to one's country, being barred from government work and clearances[13], or worse. In certain countries, citizenship is only a protection from overt and legal repercussions but processes without oversight are the main playing field of security services. Vague threats carry weight in this space.

Finally, thanks to popular culture portrayals, threat of elimination is most often associated with the work of security and intelligence services. However, in this scenario, threat of elimination does little to solve the campaign operator's problems. Only in the case of brutal services with little creativity will elimination really amount to a solution. In that particular case, it must be noted that the *threat* of elimination carries no weight, elimination itself is the ill-conceived

solution[14]. The larger point does not rest with the particular method chosen by the attacker but rather to impress the degree of peril carried by the situation that the researcher inhabits as he inconveniences powerful players in their legitimate operational space.

Company

An institution is not influenced or deterred with the same methods as an individual researcher but in many cases the threats are analogous. Threats to the company focus on impacting operational viability, revenues, ongoing and potential contracts, strategic partnerships, PR value, as well as regulation-based financial repercussions. The security services are tasked with ensuring the viability of the financial, political, and diplomatic exchanges of the nation. When the actions of a company with little cover for action other than opportunism are considered as standing in direct opposition to that national viability, that hostility merits any effective measures available. Provocation occurs in two scenarios: first, where the company's research causes political, diplomatic, or military tensions to flare between nations in an already escalated posture. Secondly, when the company's public disclosure (or private offering provided directly to sensitive targets) endangers the reputation of the intelligence agency itself or worse yet comes close to revealing or endangering the requesting customer. The former scenario is undesirable; the latter scenario is unacceptable[15].

In responding to these sensitive situations, political, financial, and regulatory repercussions against a company are employed with ease. Companies with government contracts will see these contracts dangled and unrelated vital strategic partnerships may suddenly become unstable or entirely unavailable. When international companies are involved, unsubstantiated but well-placed insinuations may suffice in closing off entire crucial market sectors and, if not, threats of loosely applied embargoes can destroy the most meticulously built business. At the end of the day, reducing or entirely eliminating the financial viability of a for-profit business is tantamount to assassination as the entity itself cannot subsist without what are often key markets or market demographics like Fortune 500 or defence-contracting companies. These threats will not be taken lightly by a company with multiple offerings for whom threat intelligence may very well be the product *de rigueur* and not a lofty pursuit worthy of sacrificing hard-earned standing in the larger market.

Finally, there is a more ephemeral threat to both individual researcher and company alike as it more specifically targets their product: piggy-backing operations. For the capable espionage operator with a prospective interest in the product of skilled threat researchers, the products of their research is a boon that entails visibility into an untold number of intelligence operations across different sectors and regions. This visibility is enviable to any intelligence agency with global standing, and as such, merits investment into an operation in its own right. As researchers are bound to be

---

[13] Effectively barring a threat researcher from a large swatch of employers in the US.

[14] These and other threats are extensively treated by the author in a forthcoming paper : Guerrero-Saade, J.A.; Pontiroli, S. Double-0 Status: The Perilous Transformation from Security Research to Intelligence Brokerage.

[15] As previously hinted, this is perhaps the one inviolable tenet of intelligence work: exposure or harm cannot befall the customer as a direct result of their involvement with the intelligence agency.

more cautious, standard espionage operations may be attempted for short-term visibility, well-placed moles may further extend this visibility where compartmentalization is not properly implemented, but nothing beats the access afforded by well-placed advanced malware. In these cases, we may see agencies pull out 'the big guns', designing truly sophisticated malware intended to operate under the very noses of the most capable malware researchers. As private sector threat intelligence becomes more valuable and researchers become more skilled than teams in many (but not all) intelligence agencies, these operations will become increasingly desirable and present favourable return on investment. As it stands, the visibility afforded by a single successful piggy-backing operation on one top-tier private threat intelligence firm would provide incomparable situational awareness. When considering the exorbitant resources required to stage operations capable of providing even semi-comparable access into such diverse areas, the most rudimentary cost-benefit analysis will favour the former by an exponential factor.

…And anxiety

Researchers' causes for concern are not entirely self-centred. There is also an element of anxiety as ethical tensions commonly arise during research. Ethical concerns can be subdivided into *passive concerns* hoisted onto the researcher by the very awareness of an operation and those *active concerns* that arise from contemplating or enacting measures to identify or disrupt those operations.

Passive ethical tensions are inevitable companions of this type of research. Conducting an investigation over a computer screen allows for an easy disconnect from the subject of the investigation. Certain cases remove this divide and force the researcher to reassess motivations, ethical commitments, and boundaries. An obvious trigger is the involvement of terrorism and radical ideologies, evidenced either by the targets of the espionage operations (where the attacker's intention is likely monitoring and curbing their activities or tasking for more aggressive 'intervention') or by the perpetrators[16] of malware-based espionage to embolden their kinetic operations. The former creates a greater conflict as researchers whose stated intention is to secure information systems are faced with espionage operations whose legitimacy is not immediately questionable. A layer of difficulty is added by the common practice of mixed use, where campaign operators will deploy the same malware to infect radical targets along with a diverse swathe of questionable targets (like research institutions in adversarial countries or opposition politicians within their own borders). The question of whether to turn a blind eye becomes more complicated as victims in clear need of protection are mixed in with extremists and the researcher's available option remains binary: to detect or not to detect?

The ever-looming threat of duplicity and ingenuity on the part of legitimate targets casts a further shade of grey. Unscrupulous agents use professions and affiliations to mask illicit movements; shell companies and research institutions are put in place as cover for questionable activities. It is well within the job description of intelligence and security services to investigate these cases and determine the true essence of actors and institutions alike. Despite greater visibility gains into an espionage operation, the researcher's insight into the operations targeting is always superficial. Yes, a politician, an academic, a businesswoman and a journalist were targeted, and at face value these seem unacceptable. But special circumstances can (and *have*) legitimated each of these scenarios, as may be the case with a corrupt politician, a radical academic, a money-laundering businesswoman, or a terrorism-facilitating journalist. The researcher has no awareness of the circumstances that invalidate the legitimacy of their respective positions, and proceeds to condemn the operation and shut off the necessary access for the espionage operation[17] to continue.

A functional school of thought has arisen to combat this ethical tension by holding steadfast to the commitment that maintaining system integrity is worth any cost. The argument insists that regardless of the attacker's intention, malicious code abuses the essential trust that enables technological progress and cannot be evaluated on the basis of case-by-case intentions. If legitimate operations lose visibility in the process, so be it, they can resort to conventional espionage methods that do not require violating the integrity of a system on which millions of users rely for valid and benevolent use.

On the other hand, the case for active engagement has been gaining increasing traction, particularly in US government contracting circles. The idea is marketed under terms like 'active defence', 'offensive security', or simply 'hacking back'. The arguments seek justification in the ultimate payoff, which is either access to otherwise unavailable information about the attackers or as a matter of retribution and deterrence. The common discomfort with the idea of hacking back arises from two interrelated concerns: first, 'what distinguishes us from attackers if we employ the same methods to compromise their systems?', and second, 'what about cases of misattribution where the wrong individual or institution is targeted in "hacking back"?' The result may very well be a firm jump away from the role of security researcher to that of vigilante blackhat or simply confused attacker.

Eliminating the incidence of all the previously mentioned scenarios requires a larger structural change, beyond the scope of the researchers themselves, that institutionalizes the mediation of a global strategic view into the production cycle itself, such that handling ethical tensions and perilous engagements is a part of the process and not an afterthought. That inclusion will in turn birth a new period for the threat intelligence market altogether.

THE MATURATION OF THE THREAT INTELLIGENCE MARKET

Describing a more mature threat intelligence market will require the adoption of new ideas and the willingness to let go of certain entrenched practices. If companies incorporate the position of a strategist, both capable of employing a broad-view strategic calculus and empowered to enact the decisions resulting therein, we can reasonably expect an end to the practice of releasing intelligence reports to the public. That market, catering solely to need-to-know partners, will in turn

---

[16] For example, Desert Falcon – where the perpetrators displayed both radical ideologies and an interest in physical security companies operating in the Middle East [9].

[17] Gauss is an operation that comes to mind as targeting specific financial information from Middle Eastern banks, perhaps for their nefarious use [10].

encounter a validation vacuum where third-party mechanisms will have to serve in place of the skeptical infosec community to endorse or delegitimize findings.

From OPSEC to a strategic operational calculus

In recent years, the topic of operational security (OPSEC) has gained greater traction among information security experts. OPSEC is seen as a necessity for proactively maintaining personal integrity in the face of all types of attackers and the perception of poor practice is considered a standard for scoffing criticism. The need for consistent and rigorous OPSEC is undebatable in an industry dedicated to curbing those with the knowhow and willingness to pilfer data that does not belong to them. That said, the most cunning threats aimed at security researchers are in many ways beyond the scope of mitigable peril for a solitary private-sector researcher. The researcher is part of an entity and as such is dependent on its structure and complicit in ventures and overtures over which he has little strategic control. His tasking is most often the result of company-wide stances and product-driven releases. Before a researcher is worthy of specific targeting, the company will be in the crosshairs of the interested intelligence service and it is at this level that the attacker's interest must be managed.

The demands of staying off the radar with the intention of evading professional surveillance enhanced by nation-state capabilities are simply incompatible with day-to-day business practices. A researcher cannot disappear on a whim, fail to report to upper management, communicate in person only, provide the product of work through indirect means like dead drops, etc. Similarly, the researcher has no access to disposable identities with the necessary paperwork nor are they in any way enhanced by law-circumventing protections. Demanding that a researcher be capable of thwarting the targeting of an intelligence service under these constraints is simply unrealistic. Moreover, companies that fail to acknowledge their role in the intelligence brokerage space are unlikely to provide the wealth of resources required for this activity to be performed at even an intermediary level. The deficiencies in this respect span a lack of both basic and extraordinary resources, as well as a lack of training, supporting infrastructure, and established security protocols.

The decisions necessary to manage IA tasking interest and thus mitigate interference have to take place at the level of proactive strategizing, enhanced by solid backchannel relationships, that establishes a mutual self-interest. The targeting must be pre-empted by creating a mutually beneficial balance that pre-empts further intervention for fear of loss. Readily available negotiation vectors are business (these same IAs are often the eager recipients of threat intelligence), visibility (IA investigations are enhanced by collaboration or sometimes handed off entirely to private contractors), and continued operational viability (considerably lowering the IA's expensive malware platforms' risk of exposure). Additional unorthodox means of negotiation must be weighed against the cost of solidifying the perception that the company is in fact a more nefarious entity than just a threat intelligence firm.

Finding a balance across the aforementioned vectors cannot be entrusted to separate departments with irreconcilable priorities. The need must be consolidated into a single position responsible for reclaiming the previously relinquished decision-making process and imbuing it with a broad-view strategic prerogative. The decisions made therein are often likely to make little 'business sense'; in other words, these decisions are unlikely to yield immediately quantifiable returns. Instead, the yield is a priceless asset for the company: the assurance of both its continued business viability and positive relations with strategic partners (regulators and legislators as well as the intelligence and security services of different territories). Moreover, the strategist's job is not only preventive. At times, aggressive unified strategies are required to offset incoming coordinated aggression through the media, the severing of strategic partnerships by means of misinformation, attempts to stifle research, and other common means of subversion against businesses. In the absence of a 'point-person' to marshal the company's vast but disjointed resources, the aggressor finds an easy victim.

An entirely new dimension of prospective decision-making is heralded by the addition of a strategist, that of *agential regency*. Due to the different forms of information asymmetry involved in the intelligence production cycle, the customer cannot define the scope of an inquiry into a piece of malware found in their systems about which they know little. The research company responding to the request eventually becomes the only player aware of both the nature of the infection and the victim's willingness to know. The company is in a position to decide whether the customer is in fact a capable recipient of the intelligence generated by the investigation. The strategist must decide on the customer's behalf whether they can, in fact, benefit from those results.

As a product intended to aid decision-making, worthwhile intelligence must be actionable. Actionability not only reflects upon the intelligence product alone, it is determined in relation to the customer's capability for action. As an unexpected actor intervening into the established dynamic between attacker and victim, the company researching a malware-based espionage operation benefits from a unique vantage point of information asymmetry from which to determine the relative action-capability of the customer in relation to an attacker the customer knows next to nothing about. This means that the research company, and not the victim-cum-customer, must decide whether to provide said customer with full or partial findings, or perhaps no information at all, for the customer's own well-being.

While the relationship with the customer may strain, the company's alternative gains are twofold:

The company remains the sole recipient of greater intelligence that can be leveraged in further investigations or delivered to more appropriate customers. The company can also avoid becoming part of the causal chain that leads to a strained relationship between 'giants', as may be the case between a large corporation and a country or, worse yet, between two nations. When caught in nation-level turmoil, the research company is inevitably out matched and its public involvement in a diplomatic fiasco can only cause problems. It is the job of the strategist to foresee these situations and gracefully sidestep them altogether.

Once companies accept and bolster this strategic emphasis into their intelligence production cycle, the threat intelligence brokerage space is likely to shift away from a front-facing business that requires splashy PR moves antithetical to the purposes of its true clients. The new business model takes

place entirely out of view, in the form of strategic partnerships that require investigations with a defined scope of inquiry. Properly embracing cyber-intelligence brokerage will entail an entirely new set of challenges in the calculus of agential regency that subsequently create a new vacuum, no longer strategic but validatory in nature.

A hidden industry in need of validation

In a hypothetical scenario where the majority of companies in the threat intelligence market accept their role as intelligence brokers and bolster their production cycle to include the strategic execution of agential regency, we can foresee an entirely new vacuum in need of address. As both the threat intelligence providers change their relationships with customers and trusted partners alike and the well of publicly available threat intelligence runs dry, customers will come face-to-face with a validation deficiency. In an information asymmetrical market, where closed-source intelligence is being traded as part of large transactions restricted by non-disclosure agreements, a series of important questions arise: what is the value of the intelligence offering? Is the threat real or fabricated by the vendor? Are the technical and political analyses accurate? If so, has the research company provided all of their visibility? And, in that case, how does this particular vendor's visibility compare with the rest of the market? More succinctly, what evaluative metrics are available to the customer[18]? If threat intelligence is to become a genuine requirement in assuring the continued integrity of the recipient company's systems, an evaluation of the intrinsic value of the intelligence offering must enter the equation. Due to the recipient companies' lack of true calibre in-house talent, this evaluation will have to be handed off to a third-party mechanism with market interests of its own.

External validation as a requirement is far from inconceivable. All it would take for customers to second-guess the value of expensive threat intelligence is a fancy appliance's failure to catch an elementary (but serious) threat to intellectual property, like an attack by the infamous APT1/Comment Crew, or being sold a vendor's entirely fabricated threat to cash-in on geopolitical tensions. In the current market, a few Tweets would give rise to an uproar of vendor disparagement and offers of greater visibility. But in the closed-door intelligence brokerage market this recourse will be unavailable. Instead, we can expect market innovation to birth a third party in one of three likely configurations:

Delegated authority – the simplest (and weakest) form of validation is that of an authoritative third party that provides 'trustworthy' vendors with a certificate of approval. The certification process may focus on the reputation of previous work or specific enquiries and will likely codify the status of already well-known vendors while ideally shutting out known charlatans who fabricate threats or purposefully misrepresent their data.

Independent verification – the more involved form of validation requires a qualified (perhaps academic) team to independently validate the vendor's findings behind closed doors. An option to avoid excess intelligence sharing is to provide the independent team with the samples, disk images, or memory dumps that began the investigation and nothing

else so as to constitute a sort of control study. An alternative function of the independent team may be to produce a geopolitical analysis of their own as well. This could be a way to involve university research teams to provide a validatory service.

Systematic intelligence escrow – a hands-off but thorough approach is for a third party to manage an automated cryptographic escrow system designed to receive threat intelligence from different vendors that can then be compared to establish a context of vendor visibility. The exchange can be based on hashes or other cryptographic means of sharing and processing information without conveying its exact contents. The specific value added by an intelligence escrow is the creation of a shared context of threat intelligence that can provide the customer with an assessment of the value of the vendor's visibility on a particular threat and also settle some of the disputes that arise from inexact or vantage-point-specific naming conventions[19]. The motivation for vendors to participate will be a requisite stamp of approval or visibility score required as part of the customer's acquisition process.

These means of addressing the validation vacuum that arises in a closed-door threat intelligence market reflect deficiencies in different stages of the intelligence production cycle. Delegated authority, independent verification, and systematic intelligence escrow validation models respectively address gathering, analysis and delivery stage failures.

CONCLUSION

The current threat intelligence market is in the midst of an identity crisis. As companies transition from plain IT security to intelligence production, the relevant methodology of intelligence brokerage must be embraced in order to stand a chance against the supernatural market tensions that are the product of meddling with the operations of diverse intelligence agencies and enraging their respective governments. The transition to intelligence brokerage proper is encouraged as a means of survival for threat intelligence producers facing escalating geopolitical tensions. By empowering the producers to strategically control their offerings, these tensions are relieved or entirely sidestepped and the market can flourish away from the limelight. A foreseeable result is the transference of this encumbrance onto the intelligence customers themselves whose innate failure (of not already having adequate in-house capabilities) will have to be resolved by yet another market player whose sole responsibility is to provide validatory assurance of the value of the threat intelligence offering.

There too exist cyber-Leviathans[20], *temporary* outliers to current market conditions. These are companies whose main offering and profit-incentive is tied into the common digital well-being of large swathes of users by virtue of their

---

[18] ...now that the skeptical infosec community is no longer a part of the equation.

[19] An issue currently plaguing our industry with an inscrutable menagerie of code-named operations and string-remnant team names. Some of the disputes can be cleared up by a shared context of visibility, and some cannot. The latter cases point to a desperate need for a more discerning framework of what constitutes an operation, a campaign, an attack team, its tools, etc. As more operators share tools and infrastructure, and as portions of these operations are outsourced to '*-as-a-service', the shared ground for resolving these disputes will only grow thinner.

[20] May the reader forgive the noxious appropriation of Thomas Hobbes' work on social contract theory.

(non-intelligence) offerings. As such they are capable of withstanding the aforementioned tensions for longer than dedicated threat intelligence outfits. In what may ring as a self-serving observation on the part of the author, anti-malware companies fit into this outlying category as they often produce threat intelligence as a matter of academic situational awareness with the added indirect benefit of being good PR. Only within similar companies are intelligence-producing research teams safe from direct financial constraints or sabotage by virtue of an alternate revenue stream. However, this insulation is not perfect. While the research itself is less likely to be affected, the overall company offering is all the more susceptible to PR smearing, nationalistic politicization, market Balkanization, and deadly regulatory restrictions that can bring even the greatest enterprise to its knees. The Internet's cyber-Leviathans will be faced with similarly difficult choices in the new market of threat intelligence brokerage.

REFERENCES

[1]   Machete. https://securelist.com/blog/research/66108/el-machete/.

[2]   https://securelist.com/blog/research/70504/the-mystery-of-duqu-2-0-a-sophisticated-cyberespionage-actor-returns/.

[3]   https://securelist.com/files/2014/11/darkhotel_kl_07.11.pdf.

[4]   https://securelist.com/blog/research/69114/animals-in-the-apt-farm/.

[5]   https://securelist.com/files/2014/11/Kaspersky_Lab_whitepaper_Regin_platform_eng.pdf.

[6]   https://securelist.com/files/2015/06/The_Mystery_of_Duqu_2_0_a_sophisticated_cyberespionage_actor_returns.pdf.

[7]   https://securelist.com/blog/incidents/64107/miniduke-is-back-nemesis-gemina-and-the-botgen-studio/.

[8]   https://www.bluecoat.com/security-blog/2014-12-09/blue-coat-exposes-"-inception-framework"-very-sophisticated-layered-malware.

[9]   https://securelist.com/files/2015/02/The-Desert-Falcons-targeted-attacks.pdf.

[10]   https://securelist.com/blog/incidents/33854/gauss-nation-state-cyber-surveillance-meets-banking-trojan-54/.