In wolf's clothing: Complications of threat emulation in contemporary cyber intelligence practice

JD Work
Marine Corps University
Quantico, VA
jw3646@columbia.edu

***Abstract*— *The dramatic expansion of threat intelligence in private practice has created a challenging multi-stakeholder environment in which cyber incidents now play out, including incidents which are not hostile intrusion or attack but rather the simulated modeling of such events for the purposes of Red Team exercise and assessment. The increasing complications of the interactions between these legitimate functions, against the backdrop of evolving adversary capabilities and deception tactics, gives rise to a number of potential consequences - including degraded warning posture. Mature organizations' intelligence, digital forensics / incident response (DFIR), and threat emulation functions must be prepared to address these issues in the course of time sensitive crisis management.*

Keywords*— *Cyber intelligence, threat emulation, warning, deception

I. INTRODUCTION

The growing recognition in recent years of the importance of cyber security as a vital function for organizations has created a substantial market for cyber intelligence services that provide insights into threat capabilities and intentions. These offerings support early warning, alignment of information security procurement spend, and the detection and mitigation of intrusion and attack incidents. The complexities of this emerging market challenge a number of enterprise activities, including longstanding security assessment practices that themselves are also now changing under regulatory pressure and improving industry benchmarks to incorporate a greater degree of fidelity in the representations of current and anticipated near term tactics, techniques and procedures (TTP) that network defenders may encounter in the wild. The collision of intelligence-led threat emulation practices in penetration testing and related Red Teaming activities with novel commercial cyber threat intelligence collection and analysis capacities has produced a number of incidents that raise difficult questions about the nature of exercise visibility, hostile deception, adaption in adversary tooling, and ultimately the manner in which risks of conditioning and other negative consequences may be mitigated by professionals within a common community of practice.

II. THREAT INTELLIGENCE AND THE COMMERCIAL MARKET

The rapidly expanding market for cyber threat intelligence services has resulted in a proliferation of vendors providing subscription finished intelligence reporting products, along with a number of intermediate intelligence offerings with varying scope of collection, processing and exploitation intended to be consumed both by decisionmakers as well as organic intelligence functions within supported enterprises. These privatized intelligence offerings compete across varying features and qualities, which are measured against other vendor portfolios, against the internal production of organic intelligence functions within private industry firms, as well as against the reporting provided through government and public-private information sharing efforts. These privatized intelligence services function against the backdrop of constantly shifting landscape of open source information provided by traditional journalism, "new" and social media, and gray literature; and develop their own capabilities for collecting technical intelligence from malware artifacts, network traffic, endpoint telemetry, and deception defense solutions. In some shops, options for human intelligence and its subset of virtual persona operations, communications intelligence, imagery intelligence, and other traditional collection disciplines may also be pursued. Intelligence offered by these solutions providers is leveraged for assistance in warning, detection, hunt and mitigation / remediation in vulnerability and threat management. Intelligence services are also frequently called upon to support digital forensics and incident response, penetration testing / Red Teaming, and other information security operations functions beyond the traditional network defender roles [1-9].

The dynamics of competition in this market have created difficult challenges for privatized intelligence providers in this space. A number of recent trends may be noted, in which a growing number of market entrants and a shift away from functional or geographic specialization have created new pressures to effectively capture client budget, even as overall market size has continued to increase. More than 45 independent services were noted as having entered the market by 2018 (although the number of active vendors does however fluctuate based on the frequent pace of merger and acquisition activity in the sector, as well as business exits for other causes). Industry market research has estimated that by 2020, approximately 15% of all large enterprises will rely on cyber threat intelligence solution. This competitive landscape is complicated by the constant tension required to prove value of external consulting and subscription service offerings over and above capabilities offered by organic intelligence functions or information sharing participation. Further complexity is introduced by the fact that most consumer organizations subscribe to multiple intelligence vendors – an average of more than four providers; with top dozen more significant vendors each supporting more than 75 enterprise clients on a recurring basis in 2018 [10], [11], [12],[13],[14].

The fiscal stakes of this market – dominated to no small degree by start-up firms – are not trivial. Operational pressures also result by the nature of the mission, where failure is not merely costly but results in damages from espionage or destructive effects that are increasingly recognized at the boardroom level, if not in major headlines. These drivers therefore place a premium on certain qualities

of intelligence – speed and timeliness, relevance, actionability among them - that are further emphasized by individual vendors and industry luminaries as competitive differentiator. This is further reinforced by priority emphasis reported in surveys of intelligence consumers, to which vendor communities are known to pay attention [15], [16].

III. False alarms triggered by live threat emulation

Incident	*Date*	*Outcome*
APT28 threat emulation	November 2016	False alarm, warning distinguished, but persistence of error
APT28 threat emulation	August 2018	False alarm, warning in error
APT29 resurfaces	November 2018	Warning issued, delayed attribution
US .mil threat emulation	January 2019	Deconfliction after delay
COBALT STRIKE threat emulation	February 2019	False alarm, delayed deconfliction
COBALT STRIKE abuse disclosure	February 2019	Warning issued without deconfliction

Fig 1: Cases of threat emulation triggered false alarms

Penetration testing is frequently a mandated function for regulated industries, imposed by government oversight or by standards accreditation processes [17],[18],[19]. Current praxis recognizes the limitations of simple automated vulnerability scanning and other generic assessment techniques that bear little resemblance to the sophisticated, subtlety, and variety of adversary behaviors observed in the wild. As a result, recent emphasis has been placed on intelligence-led threat emulation tradecraft, relying upon the best available picture of capabilities and intentions available to specific intrusion sets to craft an offensive concept of operations tailored to the targets' equities, exposures, and potential attacker interests. Intelligence informed threat emulation in penetration tests and other security assessment vehicles as a contemporary best practice has in turn been emphasized by savvy regulators, including the Bank of England's CBEST framework and the Honk Kong Monetary Authority's (HKMA) Cybersecurity Fortification Initiative; along with a similar newly released UK government wide standard [20],[21],[22].

The increasing prominence of this mission has mandated that Red Team offensive security functions therefore become keen consumers of intelligence in order to ensure proper threat representative operations. In many cases, useful patterns and recurring trends may be derived from specific intelligence reporting on adversary tactics, techniques and procedures (TTP) that are recognizable both to the defender as well as of operational value to the attack simulation. Current cyber threat intelligence emphasizes representation of these TTP within structured models, originally the Lockheed Martin originated Kill Chain and its derivatives, and now increasingly the Mitre originated ATT&CK matrix [23],[24],[25].

Observation of realistic, threat-representative emulation artifacts in the wild naturally creates certain challenges for intelligence services to distinguish these events from live intrusion or attack incidents. The increasing prevalence of such offensive security assessment efforts is likely to exacerbate this problem. In a number of real world cases, higher order effects have resulted.

The APT28 / IRON TWILIGHT / FANCY BEAR / STRONTIUM intrusion set has emerged at the center of several relevant incidents. Cyber intrusion, espionage, and associated political warfare actions by this activity group have been attributed by US and allied governments to specific units of the Russian Main (Intelligence) Directorate of the General Staff [26],[27],[28],[29]. The adversary activity in the wild has been particularly aggressive in targeting government and industry sectors across campaigns spanning years, and as such is well documented by multiple commercial intelligence and information sharing reporting sources. There is further a generally accepted consensus regarding the features by which one distinguishes this activity from other similar Russian origin intrusion sets and by which the intrusion set is recognized and clustered - based on characteristics of tooling, infrastructure, and other operational choices made on a recurring basis by adversary operators. This is rather significant as such a consensus is not a given where multiple privatized intelligence services and governments track a specific problem set. Characteristic features have remained relatively stable over time, and the set as a whole has received substantial attention in popular media due to attribution in high profile political espionage incidents. Among the many iterated adversary problems, this coverage results in a not insubstantial degree of "brand" recognition compared to other intrusion set cryptonyms among non-technical audiences – a particularly important aspect for executive interaction. As a result, the APT28 intrusion set has been seen as offering great utility for prospective Red Team functions as a representative adversary. Indeed, efforts to formalize set playbooks for threat emulation purposes have been increasingly made public [30].

APT28 has observed to target NATO associated equities for many years, consistent with known espionage requirements of the originating hostile service [31],[32]. In November 2016, commercial cyber intelligence services identified infrastructure whose social engineering characteristics were consistent with previous APT28 operations targeting NATO. This incident resulted in warning intelligence of a pending or active campaign in initial phases of new operations. However, sensitive source reporting available to one vendor distinguished this incident as Red Team activity. Despite this, infrastructure associated with the testing event continues to be attributed to APT28 through at least early 2019, due to persistent circular reporting where aging indicators of compromise are disseminated without further evaluation [33],[34],[35],[36]. The incident demonstrated that substantial realism of threat representation in offensive security assessment can create challenges in attributing an event as testing once detected.

Unfortunately, it appears that lessons of this incident were not understood as widely as would have been preferred. A repeated similar incident reached national prominence in August 2018, when a mobile security startup detected phishing infrastructure associated with targeting of technology infrastructure used by Democratic National Committee and other Democratic Party interests. Public warning was issued, and national media headlines resulted [37]. Although not confirmed, one may presume that intense interest from intelligence community elements involved in

the National Security Agency / USCYBERCOM Election Security Small Group and associated activities [38]. However, upon further scrutiny this incident too turned out to have been threat emulation based security assessment, under contract to state and local Democratic Party elements that had reportedly not coordinated the planned test events with national Party leadership or any other organizations [39].

This false alarm likely impacted the initial evaluation of another incident not long after. In mid-November 2018, a wave of spearphishing attempts was detected by multiple cyber threat intelligence services and other cyber security defenders. This campaign leveraged tooling and infrastructure that bore a high degree of similarity to prior known APT29 / IRON HEMLOCK / COZY BEAR / YTTRIUM operations [40,][41],[42]. However, sufficient differences – including the use of commodity COBALT STRIKE implant tooling also frequently used by Red Team functions – coupled with the substantial time elapsed between detected large scale actions by the APT29 operators led to not only initial caution, but robust debate between multiple researchers over attribution analysis in this incident. To no small degree such debate is necessary, expected, and the sign of a healthy community of intelligence professionals that refuse to uncritically accept raw information without aggressive consideration of alternative hypotheses and drivers. However, this debate was almost certainly informed by awareness of earlier uncoordinated threat emulation incidents in which small incongruities were the only evidence of third-party impersonation. While the exchanges did result in formal finished intelligence publication - including widely distributed public versions of this reporting with further subsequent media pick up and amplification of the story - the incident is further notable due to the speed and degree of openness with which the debate reached social media [43],[44]. One may further presume such discussion occurred on equally if not more robust basis within private working groups and trusted information sharing channels. Ultimately, the question appears to have been somewhat resolved given the inclusion of the incident in subsequent amended complaints as part of ongoing DNC legal action against Russian General Staff for damages related to ongoing intrusion activity against the Party [45].

The commodity tooling identified in the November 2018 APT29 case represents a special problem case for the cyber intelligence community. This is not the first nor the last instance of abuse of dual use capabilities, developed originally for penetration testing purposes, repurposed by adversary operators as low cost and lower probability of unique attribution implant options in sustained campaigns [46]. Abuse of COBALT STRIKE has been linked multiple suspected state programs including Iranian origin intrusion sets tracked variously by multiple vendors as Copy Kittens / ROCKET KITTEN; Vietnamese origin intrusion set APT32 / OCEAN LOTUS, as well as China origin intrusion sets APT19 / DEEP PANDA / Codoso and APT40 / Temp.Periscope / Leviathan. Additional abuse of the COBALT STRIKE tool by criminal operators, including FIN7 / Temp.METASTRIKE / CARBON SPIDER, has been identified [47],[48],[49],[50],[51],[52],[53]. Adversary repurposing of open source and commodity penetration testing tools is of course not limited to a single capability family. Additional abuse of tools including Core Impact, Pupy, PowerShell Empire, and the venerable Metasploit have been reported across multiple intrusion sets [54],[55],[56].

This widespread abuse of stolen, cracked and otherwise illegitimately obtained versions of dual use tools complicates assessment of new intrusion incidents, such as when characteristic artifacts and network traffic associated with implants or command and control (C2) infrastructure are identified. Such cases are further complicated as it is not enough for ethical hackers working within authorized Red Team engagements to simply brand their tooling or infrastructure in connection with the entity they represent, or to operate from acknowledged network infrastructure associated with the sponsoring organization. Adversary subversion of victim organization naming in infrastructure and in malware metadata is extremely common, often as part of social engineering efforts. Further, hostile operators may frequently leverage compromises of otherwise legitimate organizational servers for lateral movement, secondary payload hosting, exfiltration, or other later stage intrusion activity in sustained campaigns. And given patterns of abuse, the potential that adversary operators may be re-deploying a previously captured implant against new targets must always be considered – especially where immature (or lazy) adversary developers fail to remove legacy C2 configurations. As a result, even openly tagged features are insufficient to establish that a detected incident is indeed authorized behavior vice hostile presence. Coordination mechanisms are almost always unclear when required to address such incidents where observed by parties that are neither the attack nor the defender.

Evidence of such challenges in coordination surfaces only in occasional glimpse. In February 2019, an experienced senior researcher working for Chronicle, the threat intelligence play founded by the Google X research and development activity, identified a suspected COBALT STRIKE command and control infrastructure in connection with an unknown incident. The incident involved infrastructure associated with the Japanese technology firm NTT, observed in beaconing behavior of a malicious implant sample shared through a public malware repository. Alerting in this case was passed publicly through social media, which occasioned an unusual response in which the responsible penetration tester acknowledged the incident and indicated that the offensive operation in question had been stood down [57]. While the potential for Red Team involvement in the event was immediately recognized and discussed, resolution of the incident conclusively required some time due to differences in geographic location and working days of the involved parties.

More difficult cases are posed when the Red Team activity involves entities who have a track record of not responding well to outside inquiry by researchers, particularly when those entities are used to enjoying a perceived protection of security classification around their activities. This does not of course serve to eliminate the potential that their activities, however benign, may be observed by commercial intelligence services. But such incidents do bring into potential sharp relief the differences between cultures and organizations. Handling of these events is a delicate matter that often involves substantial self-censorship and / or restraint on the part of the commercial service. One such incident also surfaced in January 2019,

involving C2 infrastructure detected within the networks of a non-US military [58].

The abuse of COBALT STRIKE prompted a highly unusual action by the legitimate developer firm, who chose in late February 2019 to publicly release a list of identified C2 infrastructure associated with suspected misuse of the tool, including particularly a number of cases where unregistered or cracked instances of the software were illegitimately in operation [59]. This action however angered a number of other parties within the cyber security community, in part likely due to the fact it burned some Red Team activities through an alleged abuse of trusted vendor information [60]. One may also note that private intelligence holdings regarding live adversary infrastructure, and previously detected threat emulation exercise operations, were immediately devalued by widespread public release. Aside from the loss of intelligence advantage in a competitive marketplace, public disclosure of infrastructure under hostile control leads to changes in ongoing operations and can reduce future probability of detection. While such disclosure is exceptionally common, where commercial intelligence firms and independent researchers offer public disclosures out of differing motivations, there remain ongoing debates over the appropriate approach to coordination of intelligence gain / loss equities considerations. This incident clearly illustrates such debates are nowhere near to being universally resolved.

IV. False alarms triggered by range incidents

Such frictions may be inevitable when threat-representative attack emulation is conducted live on the global Internet. Operations designed in this manner are of course the most realistic test conditions of both defender posture and the unknown inherent vulnerability of the target attack surface. Internet scale security assessment is able to work against objectives leveraging complexity, scope and depth of options that simply cannot be replicated under other conditions. Nonetheless, there are certain circumstances in which test objectives are not intended to be carried out against live production environments. In these scenarios, threat emulation may be focused on delivering effects that would be out of scope due to potential consequences, risk of collateral damage, legal factors, or other unknowns that would constrain competitive behaviors in ways that are of less value to participants and observers. Similarly, the operators and planners involved may also wish to explore threat representative actions under more controlled conditions than the full scope of the global Internet, in order to test specific decisionmaking under pressure, evaluate offensive or defensive capabilities through a formal munitions effectiveness process, experiment with new technology integration options, or train new operators on either side under realistic conditions but where mistakes may be permitted in order to accelerate learning impact. For these purposes, multiple national private elements have long pursued capabilities for testing in cyber range environments [61],[62],[63],[64].

The control procedures to segment, air-gap or otherwise isolate cyber ranges from the wider networked environment and spectrum require close attention to detail – an exercise that may in some cases be likened to handling unexploded ordnance as operators work with developmental, recovered, or less than fully engineered offensive capabilities. On some occasions, these controls have failed. Responsible test and evaluation programmes will constrain malware so that it does not deliver damaging or destructive effects in the wild absent specific targeting. This is not however always assured. Even in cases where mishandled malware inflicts no collateral damage, the resulting escaped samples may produce false alarm incidents.

Incident	***Date***	***Outcome***
Shamoon variant range leakage	March 2017	False alarm, delayed deconfliction
XENOTIME / Temp.Veles intrusion	June-August 2017	Possible delayed attribution due to prior false alarms
USAF classroom demo leakage	February 2019	Unfounded belief despite deconfliction

Fig 2: Cases of false alarms triggered by range events

In March 2017, a modified variant of the Shamoon destructive implant family was uploaded to a public malware analysis system. The sample was found to be identical to an earlier observed instance, however the hardcoded detonation date / time had been changed – with less than 24 hours on the clock at the time of discovery. Initial warning was passed from at least one unidentified security vendor to a regional information security professional in the Middle East, who subsequently also alerted multiple other researchers in Saudi Arabia and Kuwait [65]. The incident sparked no small degree of alarm due to prior major destructive attacks in campaigns attributed to the government of Iran against targets in Saudi Arabia and Qatar on multiple occasions since 2012, including in November 2016 and what was then most recent re-strike in January 2017 [66],[67],[68]. In reality, this incident was not a new wave of attacks but rather the result of a mishandled sample used in an academic cyber attack and defense exercise, which had been modified to include a prominent disclaimer to this end – a banner likely not read by malware analysts relying on automated behavioral and signature detection systems vice static analysis reverse engineering techniques [69].

The timing of this incident was also particularly unfortunate, as it came only days after unauthorized disclosure of an alleged classified US government program that reportedly examined publicly known malware variants in order to consider techniques for potential adoption in future operations. Third parties sought to advance the narrative that this activity included deliberate deception efforts towards misattribution objectives. The widely reportedly claims were based on purportedly leaked documents stolen from the Central Intelligence Agency. A specific driver component – RawDisk - whose functionality is by the Shamoon malware family was claimed to have been a specific focus of the alleged program [70]. Although never independently validated, these claims were widely accepted without critical evaluation and continue to contribute to the widespread misunderstanding of concepts of “false flag” cyber attack. This complicated narrative also seriously burdened legitimate analysis, especially given the generally limited capacity and maturity of regional intelligence functions. Repeated re-strike by Iranian origin intrusion sets delivering Shamoon payloads and other associated variant malware tooling would continue to be observed through December 2018 [71].

It is further unclear the extent to which this inaccurate false flag narrative – exacerbated by false alarm warning - may have played a role in undermining trust and cooperation in the wake of later intrusion incidents targeting similar critical infrastructure networks in Saudi Arabia between June and August 2017. These incidents were reportedly marked by serious disconnects in expected communication between impacted parties, including the infrastructure operator, industrial control systems vendor, and incident responders - each also supported variously by their own separate intelligence functions and government liaison engagements. Initial attribution offered by commercial services offered conflicting analysis, with some researchers speculated Iranian involvement, apparently based solely on victimology and associated geopolitical factors [72]. Subsequent intelligence was developed linking the incident to campaigns by XENOTIME / Temp.Veles, Russian operators with a suspected state-nexus [73],[74],[75].

Such disconnects were particularly challenging as they also served to separate extant incidents from prior intelligence warning that had been circulated regarding ongoing Russian attributed intrusion operations pursuing earlier targeting against critical infrastructure networks in Saudi Arabia and elsewhere in the Middle East, with specific focus on energy, oil and petrochemical industry sectors. While this initial targeting did not indicate the degree of tailored destructive capability later seen deployed against operational technology (OT) networks in XENOTIME incidents, the activity was highly suggestive of future adversary intentions [76],[77]. However, some leading cyber threat intelligence figures have strongly suggested caution when drawing inferences regarding the intent of intrusion operators based merely on technical observations – and is perhaps often warranted given collection gaps and other limitations on incident responders [78]. This view of the appropriate scope and purpose of intelligence support may have played a role in shaping the reception of intelligence by consumers in host nation authorities and / or impacted infrastructure operators. Industry warning in this case also followed earlier reported US and allied government efforts to assist the Saudi Arabian government security services in defending against sustained intrusion activity suspected to be of associated Russian origin in prior years – efforts which allegedly foundered due to immaturity of the host nation cyber intelligence and network defense capabilities [79]. One may presume therefore that appropriate engagement to offer counterweight or corrective may not have been present.

While these issues existed prior to, and were driven independently of the problematic “false flag” narratives, incidents are not addressed in a vacuum and the challenges of communicating timelines of events to decisionmakers under stress – especially for those audiences who may have entered the discussions with pre-existing beliefs that had been shaped by false alarms, black propaganda, and initial impressions formed due to inaccurate analysis lacking appropriate rigour. It is unfortunate, as from such poisoned wells many wild conspiracy theories may have their origin.

Inaccurate impressions formed on the basis of inconclusive observations, aggravated by corrosive propaganda narratives, likely have had much greater sustained impact over time on incident response assessment than may have been previously considered. Such influences can impact even otherwise skilled professionals with extensive experience in the space.

In February 2019, a malware sample was identified in public virus information sharing holdings. This sample displayed interesting obfuscation and evasion techniques used to disguise infection delivery through Microsoft Office software. The payload however was clearly intended as a proof of concept demonstrator, delivering only a “hello, world” equivalent effect. Metadata associated with the sample identified an individual officer cadet assigned to the US Air Force Academy in Colorado Springs. This payload had apparently leaked from classroom range environments through unknown pathways, with some delay, and was subsequently noticed by a researcher during routine technical intelligence exploitation drawing upon public malware libraries [80]. This in turn resulted in another uninvolved third party expression of concern within private security information sharing working group(s), in part driven by an almost certainly unfounded belief that this artifact might represent an previously undisclosed program for development of alleged US government offensive capabilities [81]. This misinterpretation of the public observables was clearly shaped by beliefs formed upon the basis of alleged leaked documents whose veracity remains questionable, and for which the applicability of any information therein to the specific case here was more than dubious. However, due to this bias, a relatively small incident was postulated to be a potentially major revelation in exploring the scope of unacknowledged relationships within the US cyber warfighting establishment – despite much prior, highly public discussion of education and testing activities conducted with simulated capabilities within a cyber range environment within the institution in question [82], [83],[84].

V. IMPLICATIONS AND OUTLOOK

The competition between Red Team and Blue Team functions within defensive cyber operations in the enterprise is a longstanding race of “bullet versus armor”, played out against the backdrop of a cooperative common objective. However, the growth of the commercial intelligence market and the increasing militarization of cyber security incidents by hostile programs pursuing espionage and sabotage outcomes greatly complicates what was previously a “game” with only a smaller number of players and bystanders. At the working level, catching and reacting to penetration testing activity is routine task for security operations center and other incident response staff. These actions are pursued with varying degrees of existing coordination, but are generally bound by specific rules of engagement, temporal schedules, and oversight mechanisms that will allow resolution of detected testing events quickly within the context of the exercise construct. The growing capacity of multiple commercial cyber intelligence services to collect against, and direct analytic interest towards, the artifacts of such exercises that may become visible to uninvolved parties changes these interactions. As recent cases illustrate, false alarms triggered by threat emulation events can rapidly create substantial consequences event within what remain relatively small communities of professional practice.

Repeated such false alarms risks conditioning intelligence services and consumers, where such desensitization may further attenuate already challenging

warning conditions. Conditioning occurs when recurring presentation of events or associated reporting influences analysts and / or consumers to no longer view these indications within a threatening context [85]. Far too many recent examples of warning failure suggest strong reasons for concern when considering factors that may contribute to continued future failures [86],[87]. Such conditioning may be particularly exploited by deliberate adversary action as part of a pre-planned deception campaign, or may offer opportunistic benefit to less sophisticated or organized attackers.

Absent a robust commercial intelligence market, threat representative fidelity would likely not be available to Red Team organizations, nor recognized as desirable and therefore emphasized as best practice by regulators and other oversight bodies. A robust market with multiple organizational providers, prominent key talent, and a wide range of consumers however creates new dynamics that have not previously been required considerations in the development of effective warning and accurate assessment of adversary capabilities and intentions. The dynamics of this commercial intelligence market pose what may potentially be intractable new challenges derived from complexity, exacerbated by near real time interactions of multiple players who lack organic communications connectivity.

Such interactions may well provide critical feedback channels to adversary operators and planners, upon which history has shown the success of deception operations depends [88]. The extent of adversary access to communications between these participants will differ across the varying mechanisms through which these conversations play out. However, the extent to which these matters feature more rapidly in public discussion, now often in real time, offers additional advantage to the attacker and deceiver.

However, such public and semi-public communications now serve critical functions for ad hoc, emergent coordination between organizations that did not (or could not) recognize potentially impacted equities a priori. The agility and reach of contemporary connectivity options – particularly where trust-signaling measures may be concurrently offered – acts as a vital counterweight to the complexity of geographic, organizational and professional distance. This is especially important during rapidly developing incidents that often bring together individuals and entities who may have had no prior reason to have ever met or interacted, let alone around the kinds of sensitive matters that must be considered within the scope of a significant cyber intrusion or attack.

Despite potential to trigger false positive warning from multiple intelligence services, there is little prospect that formalized deconfliction solutions would be successful. Given the size and diversity of the professional communities of practice involved, it is unclear what body might offer both sufficient trust as well as sufficient span of reach to effectively function in such a designated role. Deconfliction involving threat-representative emulation of adversary operations and specific TTP that remain known only within specialized intelligence sharing structures further raises substantial complications. Imposing requirements for deconfliction under mandates by regulators, insurers, or other oversight bodies would likely trigger a substantial additional burden – at a time where initial adoption of intelligence-led, threat emulation security assessment practices very much remains a newly emerging area for many enterprises.

Deconfliction through existing information sharing structures and trust working group relationships is a far more practical outcome. However, this fundamentally shifts the purposes of intelligence sharing around observed or anticipated adversary incidents to a blue force tracking function, and is a less than ideal fit for many such organizations. Difficult competitive dynamics are also raised by any such potential coordination for deconfliction. Where communities of enterprises competing in a given market or sector may cross organizational lines for the good of all where oriented on countering adversary impact to one or more members, stronger disincentives exist where cooperation is merely intended to avoid potential unintended consequences of “friendly fire”. A shift of this nature may also exceed the legally permitted boundaries of a given information sharing channel’s charter, especially in industries where regulated enterprises must carefully weigh any potential anti-trust implications resulting from interactions with competitor firms. Informal exchange within the margins of these trust structures may be therefore represent the best possible outcome that may be hoped for and encouraged.

It thus remains that proper recognition, and adequate measures established for the handling of such incidents within the commercial intelligence services, threat emulation security assessment service providers, and consumers of these services will continue to be core to addressing these events as and when they emerge. Indeed, the consideration and reasoned action on such contingencies should be considered a fundamental criterion in evaluation of programmatic maturity of intelligence and incident response functions.

REFERENCES

[1] S. Roberts and R. Brown. *Intelligence-Driven Incident Response.* O'Reilly Media, 2017

[2] A. Liska and T. Gallo, *Building an Intelligence-Led Security Program*. Syngress, 2015.

[3] R. Fanelli, "On the Role of Malware Analysis for Technical Intelligence in Active Cyber Defense,"*Journal of Information Warfare,* vol. 14, no. 2, pp. 69-81, 2015.

[4] A. Kornmaier and F. Jaouën, "Beyond technical data - a more comprehensive Situational Awareness fed by available Intelligence Information," in *6th International Conference on Cyber Conflict*, Tallinn, Estonia, 2014: NATO CCD COE.

[5] A. Mohaisen, O. Al-Ibrahim, C. Kamhoua, K. Kwiat, and L. Njilla, "Assessing Quality of Contribution in Information Sharing for Threat Intelligence," in *IEEE Symposium on Privacy-Aware Computing (PAC)*, Washington, DC, USA, 2017.

[6] W. Tounsi and H. Rais, "A survey on technical threat intelligence in the age of sophisticated cyber attacks," *Computers & Security,* no. 72, pp. 212–233, 2018.

[7] V. E. Urias, W. M. S. Stout, and H. W. Lin, "Gathering Threat Intelligence through Computer Network Deception," presented at the IEEE Symposium on Technologies for Homeland Security (HST), Waltham, MA, USA, 10-11 May, 2016.

[8] J. Robertson *et al.*, "Darknet Mining and Game Theory for Enhanced Cyber Threat Intelligence," *Cyber Defense Review,* vol. 1, no. 2, pp. 95-122, 2016.

[9] I. Deliu, C. Leichter, and K. Franke, "Extracting Cyber Threat Intelligence From Hacker Forums: Support Vector Machines versus Convolutional Neural Networks," presented at the IEEE International Conference on Big Data, Boston, MA, USA, 11-14 December, 2017.

[10] R. McMillan and K. M. Kavanagh, "Technology Overview for Security Threat Intelligence Service Providers," Gartner. 16 October 2013.

[11] R. McMillan and K. Pratap, "Market Guide for Security Threat Intelligence Services," Gartner. 14 October 2014.

[12] C. Lawson and K. Pratap, "Market Guide for Security Threat Intelligence Products and Services," Gartner. 20 July 2017.

[13] J. Zelonis, "Vendor Landscape: External Threat Intelligence, 2017," Forrester. 29 August 2017.

[14] J. Zelonis, "Forrester New Wave: External Threat Intelligence Services, Q3 2018," Forrester. 7 September 2018.

[15] R. M. Lee, R. Holland, K. Nickels, K. Dennesen, and K. McConkey, "CTI 101: A Crash Course in Cyber Threat Intelligence Basics," in *DFIR Cyber Threat Intelligence*, Arlington, VA, 2019.

[16] R. Brown and R. M. Lee, "The Evolution of Cyber Threat Intelligence (CTI): 2019 SANS CTI Survey," SANS Institute. February 2019.

[17] *Cybersecurity Requirements for Financial Services Companies,* NYDFS, 2017.

[18] C. Bosch, "Securing the Smart Grid: Protecting National Security And Privacy Through Mandatory, Enforceable Interoperability Standards," *Fordham Urban Law Journal,* vol. 41, no. 4, pp. 1349-1406, 2016.

[19] W. Knowles, A. Baron, and T. McGarr, "The simulated security assessment ecosystem: Does penetration testing need standardisation?," *Computers & Security,* no. 62, pp. 296–316, 2016

[20] Bank of England, "CBEST Intelligence-Led Testing," 2016.

[21] Hong Kong Monetary Authority, "Cybersecurity Fortification Initiative," 2016.

[22] CREST. (2019). *GBEST*. Available: https://www.crest-approved.org/gbest/index.html

[23] E. M. Hutchins, M. J. Cloppert, and R. M. Amin, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," Lockheed Martin. October 2010.

[24] M. J. Assante and R. M. Lee, "The Industrial Control System Cyber Kill Chain," SANS Institute. October 2015.

[25] B. E. Strom, A. Applebaum, D. P. Miller, K. C. Nickels, A. G. Pennington, and C. B. Thomas, "MITRE ATT&CK: Design and Philosophy," July 2018.

[26] Department of Justice, "Grand Jury Indicts 12 Russian Intelligence Officers for Hacking Offenses Related to the 2016 Election," 2018.

[27] UK National Cyber Security Center, "Reckless campaign of cyber attacks by Russian military intelligence service exposed," 2018.

[28] UK Foreign and Commonwealth Office, "Minister for Europe statement: attempted hacking of the OPCW by Russian military intelligence," 2018.

[29] M. Galeotti, "We Don't Know What to Call Russian Military Intelligence and That May Be A Problem," *War on the Rocks,* 19 January 2016.

[30] R. Falcone, "Sofacy 2018 and the Adversary Playbook," presented at the ATT&CKcon, McLean, VA, USA, 23-24 October, 2018.

[31] N. Mehta, B. Leonard, and S. Huntley, "Peering into the Aquarium: Analysis of a Sophisticated Multi-Stage Malware Family," Google. 5 September 2014.

[32] FireEye, "APT28: A Window into Russia's Cyber Espionage Operations," 27 October 2014.

[33] Crowdstrike, "CrowdStrike Identifies Activity Likely Meant to Emulate FANCY BEAR Operations," 7 November 2016.

[34] Tr1adx, "Domain IOC's associated with APT28 campaigns," 28 December 2016.

[35] AlienVault Open Threat Exchange (OTX), "Bear Spotting Vol. 1: Russian Nation State Targeting of Government and Military Interests," 16 January 2017.

[36] M. Stampar, "Maltrails," 12 February 2019.

[37] M. Murray, "Lookout discovers phishing site targeting DNC," Lookout. 22 August 2018.

[38] J. Lynch. (2018, 5 November) Midterms Security Watch: Quiet Election Day early sign of cyber policy success. *Fifth Domain*.

[39] L. Matsakis. (2018, 23 August) Why the DNC Thought a Phishing Test Was a Real Attack. *Wired*.

[40] M. Dunwoody, A. Thompson, B. Withnell, J. Leathery, M. Matonis, and N. Carr, "Not So Cozy: An Uncomfortable Examination of a Suspected APT29 Phishing Campaign," FireEye. 19 November 2018.

[41] D. Goodin. (2018, 19 November) Russia's Cozy Bear comes out of hiding with post-election spear-phishing blitz. *Ars Technica*.

[42] L. H. Newman. (2018, 20 November) Russia's Elite Hackers May Have New Phishing Tricks. *Wired*.

[43] A. Thompson, N. Carr, J. Slowik, R. Beitlich, *et al*. (2018, November 18) "As promised, I'm going to share some of the hypotheses we considered…". Twitter. https://twitter.com/qw5kcmv3/status/1064649362578370560

[44] C. Glyer and N. Carr, "Holiday APT Spectacular," in *State of the Hack*, 11 December 2018.

[45] N. Perlroth, "D.N.C. Says It Was Targeted Again by Russian Hackers After '18 Election," in *The New York Times*, 2019.

[46] R. Gold, "Threat Actor's Use of Cobalt Strike: Why Defense is Offense's Child," Digital Shadows. 2018.

[47] Clearsky and Trend Micro, "Operation Wilted Tulip," July 2017.

[48] N. Carr, "Cyber Espionage is Alive and Well: APT32 and the Threat to Global Corporations," FireEye. 14 May 2017.

[49] FireEye. "Advanced Persistent Threat Groups." 17 February 2019

[50] Proofpoint, "Leviathan: Espionage actor spearphishes maritime and defense targets," 2017.

[51] FireEye, "Suspected Chinese Cyber Espionage Group (TEMP.Periscope) Targeting U.S. Engineering and Maritime Industries," 16 March 2018.

[52] N. Carr, S. Mohankumar, Y. Londhe, B. Vengerik, and D. Weber, "FIN7 Evolution and the Phishing LNK," FireEye. 24 April 2017.

[53] S. Miller, J. Nuce, and B. Vengerik, "FIN7 Spear Phishing Campaign Targets Personnel Involved in SEC Filings," FireEye. 7 March 2017.

[54] G. Evron and T. Werner, "Rocket Kitten: Advanced Off-the-Shelf Targeted Attacks Against Nation States," presented at the 31st Chaos Communication Congress (31C3), Hamburg, Germany, 27-30 December, 2014.

[55] Secureworks, "Iranian PupyRAT Bites Middle Eastern Organizations," 15 February 2017.

[56] FireEye, "FIN10: Anatomy of a Cyber Extortion Operation," June 2017.

[57] S. Cutler, F. Roth, R. Sayfiev, *et al*. (2019, 10-11 February) "CobaltStrike Beacon using C2…". Twitter. https://twitter.com/silascutler/status/1094789601372319744

[58] A. Thompson, J. Case, *Drunk Binary*, *Sapphire*, *et al*. (2019, 19 January) "TFW when it's a confirmed C2 server..." Twitter. https://twitter.com/QW5kcmV3/status/1086716349102125056

[59] Strategic Cyber, LLC. "Cobalt Strike Team Server Population Study". 19 February 2019.

[60] ArmitageHacker, A. Thompson, et al. (2019, 21 February) "The manner in which the person discussed...". Twitter. https://twitter.com/QW5kcmV3/status/1098510007434444800

[61] D. L. Bergin, "Cyber-attack and defense simulation framework," *Journal of Defense Modeling and Simulation,* vol. 12, no. 4, pp. 383–392, 2105.

[62] G. Cowan. (2019, 18 January) Cyber ranges: Training and testing for virtual weapons. *Jane's International Defence Review*.

[63] M. A. Gallagher and M. C. Horta, "Cyber Joint Munitions Effectiveness Manual (JMEM)," *American Intelligence Journal,* vol. 31, no. 1, pp. 73-81, 2013.

[64] SHEN X. (沈雪石), "The analysis on the trends of cyberspace attack and defense technology (网络空间攻防技术发展动向分析)," *NATIONAL DEFENSE Science & TECHNOLOGY,* vol. 38, no. 4, 2017.

[65] M. Aldoub, *Meshal*, and *Hamoud et al*. "Just received from security vendor, an alert..." Twitter. 17 March 2017. https://twitter.com/kingmeshal1/status/842848896917225472

[66] R. Falcone, "Shamoon 2: Return of the Disttrack Wiper," Palo Alto. 30 November 2016.

[67] Symantec, "Shamoon: Multi-staged destructive attacks limited to specific targets," 27 February 2017.

[68] K. Albano and L. Kessem, "The Full Shamoon: How the Devastating Malware Was Inserted into Networks," IBM X-Force Incident Response and Intelligence Services. 15 February 2017.

[69] Crowdstrike, "Mishandled Sample of Shamoon Malware Leaks from Cyber Defense Exercise," 17 March 2017.

[70] K. Zetter. (2017, 8 March) Wikileaks Files Show The CIA Repurposing Hacking Code To Save Time, Not To Frame Russia. *The Intercept*.

[71] Symantec, "Shamoon: Destructive Threat Re-Emerges with New Sting in its Tail," 14 December 2018.

[72] E. Kovacs. (2018, 15 December) Iran Used 'Triton' Malware to Target Saudi Arabia: Researchers. *SecurityWeek*.

[73] J. Gutmanis, "TRISIS, TRITON, Hatman on the ground," presented at the S4x19 conference, Miami, FL, USA, 15 January 2019.

[74] S. Lyngaas, " TRISIS investigator says Saudi plant outage could have been prevented. ," in *Cyberscoop*, 2019.

[75] FireEye, "TRITON Attribution: Russian Government-Owned Lab Most Likely Built Custom Intrusion Tools for TRITON Attackers," 23 October 2018.

[76] Crowdstrike, "BERSERK BEAR Reconnaissance Against Middle East Oil & Gas Sector Targets," 6 April 2015.

[77] Crowdstrike, "Suspected BERSERK BEAR Compromise Activity, Broadened Targeting," 10 April 2015.

[78] S. Caltagirone. "'Intent divination' in conclusions has to be the most common analytic issue I see...". Twitter. 12 February 2019. https://twitter.com/cnoanalysis/status/1095375732027678720

[79] CCDP, "Incident response to suspected Russian origin intrusion activity targeting Saudi Arabia," April 2016.

[80] F. Roth. "Testing a new character obfuscation #YARA rule...". Twitter. 9 February 2019. https://twitter.com/cyb3rops/status/1094145047291924480

[81] Private communication with author, "Caught building cyberweapons…," 10 February 2019.

[82] M. Carlisle, M. Chiaramonte, and D. Caswell, "Using CTFs for an Undergraduate Cyber Education," presented at the USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE '15), Washington, DC, USA, 11 August, 2015.

[83] K. McCaney, "D-FENSE: Army nips Air Force in NSA's cyber competition," in *Defense Systems Update*, 21 April 2014.

[84] T. Chuang, "College hackers converge in Denver for cybersecurity competition," in *Denver Post*, 9 March 2017.

[85] D. T. Moore, "A Short Primer on Deception and What to Do About It," *American Intelligence Journal,* vol. 32, no. 2, 2015.

[86] J. D. Work, "Escalating consequences: Offensive cyber operations and intelligence failures in the Putin era," presented at the Cambridge Intelligence Seminar, Corpus Christi College, Cambridge University, 11 May 2018.

[87] J. Slowik, "Meet Me in the Middle: Threat Indications and Warning in Principle and Practice," presented at the DFIR Cyber Threat Intelligence, DFIR Cyber Threat Intelligence, Arlington, VA, 2019.

[88] B. Whaley, "The Process of Deception," in *The Art and Science of Military Deception*, H. Rothstein and B. Whaley, Eds.: Artech House, 2013.