DRAW ME LIKE ONE OF YOUR FRENCH APTS – EXPANDING OUR DESCRIPTIVE PALETTE FOR CYBER THREAT ACTORS

Juan Andres Guerrero-Saade
Chronicle Security, USA

turla@chronicle.security

ABSTRACT

Words are the scaffold of our thinking. They allow us to climb conceptual heights, survey the land, to map concepts, and guide our understanding through conventional and unconventional pathways while retaining a semblance of structure and order. However, when it comes to the descriptive study of digital adversaries, we've proven far less than poets. Currently, our understanding is stated in binary terms: 'is the actor sophisticated or not?'. That is to say, 'is it respectable to have been breached by this formidable adversary, or were the defenders simply incompetent?'. This dichotomy of exceptionalism may have worked when the AV industry first began to encounter notable adversaries, hesitating to describe their vague features.

As the years go by, the menagerie of adversaries has become overpopulated and our familiarity with them has grown. It's time to expand our descriptive palette to include what intentions and capabilities we can surmise as present at the other end of the keyboard, to issue more fine-grained guidance on the nature of those dastardly attackers that have breached our walls. The intention of this talk is to move beyond 'sophisticated' (the pencil), to the observance of specific tradecraft (crayons), the study of intentions (watercolours), and what the observance of certain military concepts may tell us about the adversarial outfit in question (oil paints). This ambitious endeavour seeks to efface the oversimplified terminology of 'sophistication' in favour of a range of more nuanced descriptive language reflective of the TTPs researchers have been documenting for years.

Not only will the use of more nuanced language provide defenders with a better understanding of the forces they're actively engaging, but it should also allow us to better predict and understand the difference between a ragtag band of opportunistic crooks and a military outfit steeped in both real and abstracted conflict. When it comes to the latter, are we really satisfied by saying that they're highly capable? Or well-resourced? 'Sponsored by so-and-so'? Or may we be able to surmise that they're informed by previous military conflicts? By experience with counterinsurgency or counterterrorism? That their behaviour suggests constriction by the rule of law, restrictive legal frameworks, and overwrought societal concerns? Or are they perhaps emboldened by an existential struggle? Prone to fever-pitched decision-making? Even rewarded for unbridled creativity, easily confused for irrationality in terms of conventional warfare?

Let's move beyond finger-painting and get serious about our art.

INTRODUCTION

> 'The concept of seeing makes a tangled impression. Well, that's how it is. – I look at the landscape; my gaze wanders over it, I see all sorts of distinct and indistinct movement; *this* impresses itself sharply on me, *that* very hazily. How completely piecemeal what we see can appear! And now look at all that can be meant by "description of what is seen"! – But this just is what is called "description of what is seen". There is not *one genuine*, proper case of such description – the rest just being unclear, awaiting clarification, or simply to be swept aside as rubbish.' [1]

Research into digital espionage takes its cues from excruciating technical minutiae. High-level programming concepts translated into optimized assembly stitch together operating system APIs to orchestrate mundane functionality. A file is copied, a screenshot taken, a server contacted, a password exfiltrated. Piecemeal functionality is codified into a dense package. When characterized surgically, it's easy to view a malware infection as an impersonal operation – a fact of interconnected life. But that impersonal view glosses over the targeted nature of a subset of incidents, made transcendent not by the malware involved but by its tasking. A small but significant portion of the overwhelming amount of malware that traverses the Internet on a daily basis is meant for specific victims, intended for specific institutions, targeting specific verticals, as part of carefully crafted campaigns to fulfil intelligence requirements.

Private sector threat intelligence teams produce extensive breakdowns of discernible operations. These reports – often sold to eager customers as part of six-figure subscriptions – paint elaborate pictures far beyond malware functionality to delineate campaigns by well-resourced threat actors. The means of characterizing these threat actors differ from operation to operation, further variegated by the particulars of the research team's visibility, position and incentives. Sometimes attribution appears laughably simple, sometimes intentions are discernible, in some cases the fog of mystery remains undisturbed.

The industry is currently plagued by a counterproductive obsession with attribution. At times it's fuelled by overzealous customers intent on pointing the finger at their would-be attackers. In other cases, analysts recently surfaced from government and intelligence institutions forget that they no longer serve masters with recourse to retribution based on their attribution claims. And yet, others run away from attribution even when it's nearly certain, shirking politics via media-trained language laser dancing.

Regardless of whether we side with the diamond-studded attribution enthusiasts or the whodunit solipsists, a greater issue plagues us in the lacuna that divides technical breakdowns from attribution certainty – the hermeneutics of actor description. That is to say, when vision is not clear enough to make out the colour of our attacker's leather jacket, how do we instead describe their fuzzy contour? Do we resort to an estimation of height, colours and scents? Rather than provide the makings of a portrait, our collective discourse betrays a prevalence of stupefied

oversimplification. And no better word symbolizes our descriptive shortcomings than the exalted rank of 'sophisticated'.

Not to be misled by a dictionary denotation of 'cultured' or 'refined', the term 'sophisticated' as used to describe threat actors denotes placement in a hierarchy of technical capability and resources. While early uses may have arisen naturally from a researcher's genuine wonderment at the technical stratagems invoked by the attackers, the term has since been corrupted in the service of cheap PR and expensive professional services: the former hopes to grab headlines by suggesting that an attacker is so remarkable as to be newsworthy – though that's most often not the case. The latter turns incident response engagements into pay-for-exoneration schemes wherein a compromised company can claim that a mundane breach was the cyber equivalent of an 'act of god' against which no defence could have been mounted and thereby that no liability should be incurred.

While these terminological bastardizations are a seemingly unavoidable byproduct of the commercialization of threat intelligence (TI), allowing 'sophistication' to stand as the primary metric by which threat actors are measured in technical research is a limitation that threatens research fidelity. Rather than condemning the term and its colloquial proliferation, we must instead find a more fitting analogue to guide our research efforts. We will attempt this complex task in three stages:

First, an attempt to describe the features of the cyber domain and their epistemological implications. Then, we derive a functional understanding of behavioural profiling as is used in criminal investigations and adapt its underlying thesis to fit the domain of cyber operations such that we come to consider that *operational behaviour and tooling reflects adversarial configuration and imperatives*. With that in hand, we can go on to consider the possible adversarial configurations and intricacies behind what we call a 'threat actor'. And finally, we apply some of these profiling insights to past research in order to unearth unexplored facets that were originally glossed over or missed entirely at the time of discovery.

The purpose of this circuitous exercise is to inch threat intelligence research closer to a more rigorous practice of comprehensive and dynamic adversary profiling.

EPISTEMOLOGY OF THE FIFTH DOMAIN

Threat intelligence is a field of study that has arisen organically from different field investigations and 'just-in-time requirements'. As such, informed formal thought and its methodological fruits have proven rare. Without derailing our ultimate study of threat actor profiling, we would do well to first focus our attention on an epistemological question that predates – and possibly precludes – our endeavour. The main question is: 'what adversarial knowledge is intrinsically possible in the fifth domain?'. Or, '*what knowledge do the inherent characteristics of the fifth domain enable us to gather about attackers?*'. While this step back may appear as a devolution of attributory capabilities, it's meant to point out that in our immature stumblings to establish a new realm of study, we've reached an unjustified stance of certainty. With added experience, we'd do well to backtrack and study the foundations of the beaten path before laying concrete over it.

Major General Amos Yadlin[1] characterizes the cyber domain[2] as having 'unlimited range, very high speed, and [...] a very low signature'[3] That's an insightful beginning for a characterization and one that we'd do well to further flesh out. What does it mean to operate within a domain of 'warfare' that connects most targets almost instantaneously regardless of location? One where the materials for attack are infinitely replicable and the identity of both attacker and victim are largely unverifiable? Additionally, how is battlefield command-and-control changed by dependence on a medium chosen and maintained by the victim? Let's explore each of these concepts in their own right.

Signature

The simplest expression of signature in a domain of warfare is that of 'identifying marks'. When a shot is fired or a missile launched, there is an expectation that the trajectory of the weapon's deployment can be accurately traced. Similarly, an explosive device leaves identifiable traits of the materials employed in crafting the device. But the composition and deployment of a 'digital weapon' is subject to neither of these traits.

The issue of signature as a possible determination of provenance in the use of a weapon breaks down when it comes to 'cyberweaponry'. Whereas semtex is not readily available, both malicious and mundane-but-useful code is readily available online. Attackers of all calibres have a penchant for *borrowing* readily available code, to a lesser or greater extent. Where one might consider '*copy-paste dev'ing*' a low-skilled capability (as it so often is), the logic of its adoption in relation to a signature metric resembles a bell-curve:

- *Unskilled attackers* borrow code because they simply don't have access to better development resources – **Low Signature**
- *Well-resourced attackers* mostly prefer to develop in-house (for quality assurance purposes) but will still employ convenient snippets of code[4] or entire tools and libraries[5] so as not to unnecessarily reinvent the wheel – **High Signature**
- *Cunning actors with an interest in misleading researchers* will once again opt for a bulk of open-source code or even off-the-shelf malware in order to *hide in the noise* of the larger (unremarkable) use of these tools – **Low Signature**

That observation elucidates the cornerstone concept of signature in the fifth domain as that of scarcity – otherwise referred to as 'closed-source' or proprietary tooling. Where one may reasonably expect that any high-level functionality can be arrived at by different developers, specific implementations require intimate familiarity. Even deep study from the outside (the work of proficient reverse engineers) will not always reveal more obscure requirements like compiler configurations,

[1] Commander of Israeli Defense Intelligence (2006–2010).
[2] In his original quote, he refers to 'cyber' as the fourth domain – an Israel-centric departure from the US-centric model that includes 'space' as the fourth domain, thus displacing 'cyber' to the fifth domain.
[3] *Zero Days*. Alex Gibney (Director), Magnolia Pictures, 2016
[4] The Lamberts' use of a publicly available HTTP wrapper written in C++ for Pink Lambert functionality.
[5] Equation Group's use of SleuthKit.

unusual libraries, and coding styles. It's the closed-source nature of a proprietary codebase that serves as a form of signature in the fifth domain[6].

Many plausible scenarios affect this metric, some without overt indications. These include:

- Developers migrating to other teams or defecting to other countries.
- Repackaging the codebase or reselling to another customer-cum-threat actor.[7]
- A breach of the original attacker by another threat actor.[8][9]
- Hack-and-leak operations.[10]
- The original proprietor of a codebase choosing to open-source their own toolkit in order to hide their operations amidst the sea of new adopters.

As such, we can see that the proprietary nature of a toolkit is a presumption based on observations that can become outdated or inaccurate without notice. Perhaps we can view signature less materially – as a combinatory characterization of actor behaviours, tooling, tasking, etc.[11] – but that brings up more endemic problems in a medium that favours replicability.

Replicability

There is an inverse relationship between replicability and authenticity. In 1935, a Frankfurt School thinker, Walter Benjamin, came to tussle with the effect of replicability on the value and authenticity of art. His thesis was ultimately that the ability to infinitely reproduce a work of art devalues the work of art. That's to say that having something as tacky as a Mona Lisa mousepad detracts from the aura of encountering the mystery lady in her original oil-painted, Louvre-ensconced glory. This antagonism extends beyond the realm of art and will prove particularly problematic in relation to cyber operations.

To understand concerns over authenticity in the fifth domain, we have first to see how replicability permeates it. With sufficient observation, resources and access, anything in this domain proves functionally replicable. Beyond the signature issues previously discussed, replicability is a multifaceted feature with both positive and negative effects. It enables extreme capabilities in both defence and offence.

For defence, replicability is expressed in two forms: the ability to abstract defence and the ability to replay network attacks:

1. A *defensive abstraction* is visible in the form of signatures[12] (on disk, in memory, or at network level) that flag the presence of an undesirable component. Beyond rudimentary 'sigs' that only check for hash values, specific sigs enable the detection of smaller code components, heuristic or network behaviours, and generally suspect behaviours. This enables a more proactive defence of both things known and things unknown but malicious in a familiar way[13].

2. *Attack replication* is an advanced defence measure alluded to by Rob Joyce in an illuminating talk [2] describing measures that would make the work of TAO[14] hackers more difficult. He refers to it as the 'out-of-band tap', a defensive device set up to fully mirror and record network traffic within a perimeter and store it in a way that is not integrated with the rest of that network. This allows particularly cautious defenders to move beyond event logs and second-order indicators of infection to a complete replay of the actions undertaken by the attacker, the tools transferred, and commands communicated during the actual[15] attack. This is one of the most powerful (and least adopted[16]) tools in the advanced defender's arsenal.

From the perspective of the attackers, replicability has found notable expression in defeating authentication (by way of 'replay attacks') and fooling monitoring devices. The latter is best exemplified in the infamous Stuxnet attacks, where the normal operation of programmable logic controllers was recorded in 21-second intervals and replayed to assuage the concerns of the nuclear station operators while the sabotage was occurring[17][18]. This replay mechanism not only added to Stuxnet's ability to operate undetected for a prolonged period of time but also added a psychological warfare element by thwarting what may have otherwise proven to be competent assessments of the root cause of the centrifuge failures.

[6] A thesis which, if accepted, highlights the ethical dilemma of wide distribution of TTPs in the form of public blog posts and reports for indiscriminate consumption. After all, 'false flag' operations are only made possible by a shared awareness of the deceptor and deceptee of the features of the third party mimicked by the former.
[7] Where external developers (i.e. private contractors) are employed.
[8] This is a particularly insidious scenario as it's not only an unannounced breach but also a technological transfer to another actor with established capability and intent. The false flag potential is highest here.
[9] Increasing reports of attribution by means of counter-CNE and fourth-party collection among nation-state sponsored teams makes this seemingly implausible scenario a reality in need of accounting for.
[10] Increasingly common among reckless 'hacktivists' willing to dump sophisticated toolkits and ready-to-use zero-day exploits for the sake of momentarily inconveniencing an evil-doing provider.
[11] Commonly paraphrased as the umbrella acronym 'TTPs' (Tools, Techniques, and Procedures), for better or worse.
[12] The anti-virus industry uses this term: a 'sig' is a functional description of a portion of a file, its features, or a sequence of its behaviours that allows a system to recognize its presence and provide a general or specific diagnosis of that file.
[13] Signatures for 'suspicious' behaviour.
[14] Publicly known acronym for the National Security Agency's Tailored Access Operation unit.
[15] An important distinction, given that dynamic analysis often consists of executing malware in a mimicked environment without the luxury of operator interaction or live command-and-control infrastructure to show further attack stages.
[16] Due, in part, to network mirroring being a tragically underserved area of the defensive market. The availability of suitable hardware for out-of-band tapping at scale and the efficient storage and analysis of the voluminous output is currently sparse.
[17] *'Before the malware runs an attack routine, it records the centrifuges' normal operating frequencies and feeds this recorded data to the WinCC monitor program during the attack. The result is that the system shows normal operation instead of alerting personnel to the anomalous frequencies the centrifuges are actually running at.'* [3]
[18] The mechanism in the malware is detailed in *Symantec*'s Stuxnet Dossier [4] p.47, 'State 1: Recording'.

Furthermore, replicability in the fifth domain raises a further concern for analysts in that it disproportionately enables a level of *misdirection* that is impossible in other domains. Hard as extremists may try, they cannot simply conjure rare explosives or other weaponry identifiably used by another state or extremist group. In order to misdirect forensic investigators, the aforementioned extremist group would likely first have to engage in a prohibitive covert effort to source these materials. The same limitation does not apply in the cyber domain.

Identifiable tooling is more or less readily available in its completed and weaponized form. While proprietary source code may be inaccessible, binaries used in attacks can be acquired, lightly altered if needed, and redeployed. Without alteration, these binaries will not provide functional value to their adopters, but that's not to say that they can't provide value of a secondary nature by their mere presence or deployment. The aforementioned form of misdirection is not without precedent and has been observed in multiple instances in the wild. Both Turla [5] and Cloud Atlas [6] have deployed already-compiled malware[19] from unrelated threat actors in order to misdirect investigators at times of scrutiny.

Similarly, a more insidious use has already appeared with the discovery of a copied Rich header [7][20] in the Olympic Destroyer [8] attacks. This truly pedantic form of misdirection, consisting of copying a fragment of BlueNoroff[21] binary metadata onto Olympic Destroyer binaries, was so granular as to nearly go undiscovered. The chosen mechanism was likely the result of a misunderstanding on the part of the attackers of how automated code similarity systems work. The intention appears to have been to misdirect the researchers into pointing the finger at North Korea as the perpetrator of the attacks on the PyeongChang Winter Olympics and thereby achieve the tertiary effect of inflaming geopolitical tensions at a time of precarious diplomacy in the Korean Peninsula.

It's worth noting that this form of 'APT-on-APT' misdirection[22] is only made possible by the greater public awareness of cyber attack methodology and its mainstream politicization. The pursuit of a tertiary geopolitical effect entails the intention to benefit not just from the fruits of the intrusion itself but from its subsequent misinterpretation as the act of an unrelated third party whose tooling and/or infrastructure has been co-opted to some extent. While replicability enables blanket immunizations for herd protection and defensive monitoring to extreme fidelity, it's a feature of the fifth domain with devastating abuse potential.

At this time, those changes are hand crafted and thereby unscalable. However, the trend towards the automatization of operational trade-crafted exhibited by Project Sauron [10] (a.k.a. Remsec or Strider [11]) and evidenced in the Vault7 Marble Framework leak [12] should increase our collective level of concern. Attackers exhibiting a heightened level of structured operations security are likely to find great appeal in the ability to blend their targeted operations into the noise of mundane infections[23] or incorporate enough counterintuitive granular indicators to temporarily divert even expert attention towards a less newsworthy culprit. They are unlikely to thwart expert scrutiny in the long term but decision making at times of geopolitical pressure does not always count with the requisite luxury of time.

Much as Benjamin found art devalued by mechanical reproduction, we find replicability destroying the notion of a straightforward incident in the fifth domain. The lack of ultimate arbitration to wade between the early 'hottakes' and the results of expert scrutiny in high-profile incidents and settle on a definitive answer that invalidates the former has led armchair analysts and spectators into a sort of interpretive solipsism that views every attribution claim as unfounded, every attack as a false flag, and all functionality as intended to disguise some other unidentified effect[24].

Speed

The value of the fifth domain's seemingly immediate speeds is largely lost on a generation accustomed to fibre connections but in the eyes of the old military guard, immediacy is a truly attractive feature. Where land, sea and air missions are largely constricted by transportation logistics and speed limitations, cyber missions often benefit from immediacy at a virtually non-existent transportation cost. Speed is yet another measure that disproportionately favours the attacker in multiple ways, limited only by speed choke points like outliers in connectivity and configuration, ease of reconnaissance, and interaction dependence:

- The target's remoteness (i.e. oil platforms or vessels at sea) and configuration (air-gapped or segmented networks) will likely affect exfiltration, and command-and-control speeds and reliability. However, neither has proven prohibitive in the past.
- From an operational perspective, the most time-intensive phase of an attack is the reconnaissance and tool-preparation stage, specially where specialized tools are required for the first time. From the perspective of a military or time-sensitive intelligence operation, the need to prepare the specifics required against an unexpected target will prove the greatest speed limitation on the road to mission success.

[19] It's important to note that without modification, the malware's functionality did not serve the threat actors that deployed it. It simply served to mislead.

[20] Note that this is a truly obscure, largely undocumented feature of Portable Executable files generated with Microsoft Visual Studio.

[21] A subgroup of the Lazarus Group believed to be of North Korean provenance and commonly associated with large-scale financial heists like attacks on SWIFT-affiliated banks and casinos [9].

[22] Loosely referred to as 'false flags', despite their low performance standards as such.

[23] Who bothers to scrutinize a commonplace infection like Conficker or Zeus anymore?

[24] While the features of the fifth domain lend themselves to Kaizer-Sözeesque digital schemes, schizophrenic thinking about cyber attacks is likely the result of a lack of hands-on involvement rather than adversary cunning. Where one finds oneself tempted by these mental pirouettes, one should be reminded that celebrity hottakes are also capable of extensive undiscerning replicability and that the damage of these widely served misconceptions is hard to correct for a wide-eyed public without access to – or interest in – the technical minutiae of an industry primarily intended to keep video streaming uninterrupted and online banking unmolested.

- Finally, as most common infection vectors require a form of user interaction[25], something as simple as an office schedule can prove an irksome operational limitation for time-constrained attackers without recourse to user-independent infection vectors.

Despite these choke points, speed in the fifth domain disproportionately favours attackers[26]. Moreover, it can and has been used as an attack feature in itself:

- By purposefully limiting exfiltration speeds in order to operate for longer periods of time without arousing suspicion from network monitors.
- By degrading a victim's connection speeds in order to promote the use of fail-insecure protocols and habits.

Ultimately, while attackers get to prepare their attacks before setting foot on the premises, defenders are always already at the disadvantage of becoming aware of attacker tooling and cleanup requirements once the operation is ongoing[27].

Medium dependence

The fifth domain's peculiar feature of medium dependence goes hand in hand with that of speed. Cyber attacks are dependent on a medium controlled by the victim or an uninvolved third party in a variety of ways. This is unusual when compared to other domains. For example, satellite-based command-and-control is dependent on atmospheric conditions but these are not controlled by the target, nor third-party providers. Similarly, while air and sea vessels have to navigate the treacherous conditions of their respective medium, the favourability of the sea and sky cannot be altered at will. Without significant alteration to a victim's environment, the attackers are at the very least dependent for their command-and-control on the victim having paid the Internet bill.

Furthermore, command-and-control availability in the fifth domain is subject to medium-dependent conditions that include infrastructure controlled by a series of third-party providers[28] (i.e. domain registrars, VPS hosting providers and ISPs) and oversight by one or more foreign governments[29], and furthermore is subject to the general health of the Internet[30]. Additionally, implant availability is also subject to mundane constraints such as whether the victim system is shut down overnight to conserve energy, or the dwindling battery life of a mobile device that wasn't charged the night before. Where availability is critical, redundancy is required. However, an inverse relationship between redundancy and covertness should also be kept in mind.

[25] Lamentably, as simple as opening a link or attachment.
[26] In so far as we consider the 'attack' what Daniel Moore refers to as an 'event-based attack', the turnkey element made available after the reconnaissance and lateral movement (the 'presence-based attack') phase of the incursion has already been staged [13].
[27] And more likely, long after the operation started.
[28] Each with more-or-less qualified security teams and remits.
[29] Some of which shut down Internet access at times of political instability.
[30] Subject to redirections and occasional regional failures.

Range

Range limitations are largely non-existent for Internet-connected targets. Cyber weaponry[31] isn't subject to the laws of gravity, cloud conditions, and the fidelity of satellite imagery to determine whether it can, in fact, reach its target. Given standard connectivity, the target should be reachable. Even where medium hops are required (as in the case of air-gapped networks), the target can still be reached given poor transfer hygiene practices and enough actor ingenuity.

However, *target discernment* is not a given in the fifth domain. While a great deal of technological refinement has gone into accurate targeting in other domains like air – where certainty in hitting a bunker and not a nearby hospital is an absolute necessity – the same is neither true nor easy in the fifth domain. While the attackers are able to target a given set of network devices, their users are not likely immediately identifiable. Moreover, in some environments the relationship between user and device isn't necessarily static (as multiple users can employ a single device and vice versa).

Regardless of intended targeting, attackers will likely have to access multiple devices along the route before determining whether they have, in fact, reached a device that would enable them to carry out their operation. The inability to identify[32] or reach[33] targets with ease has also incentivized more promiscuous spreading mechanisms. While worms appeared to be extinct for nearly a decade, notable exceptions surfaced along the way, with the trend towards worms for 'ransomware' or destructive ends increasing in the past two years[34]. Obviously, the greater the aggressiveness of the spreading mechanism, the greater the likelihood of the malware being discovered in a shorter period of time[35].

Default discreetness

Another peculiar feature that sets apart the fifth domain from those that preceded it is the default discreetness of network intrusions. While an armed incursion or a bombardment is a visceral event, a network intrusion will likely go undiscovered for a prolonged period of time. It's possible for a cyber incursion to take place, be successful, and culminate without the victim being made aware. Victim awareness is subject to the availability, quality and configuration of their network security

[31] That is, cyber weaponry intended to be deployed via the standard connective tissue of the Internet. Given the ingenuity of modern defence contractors, that isn't the only connective medium at play.
[32] High-powered SIGINT agencies have spent sizable resources mapping the relationships between devices, cyber personalities, and users. Refer to the leaked slides for the NSA program 'TREASURE MAP' (classification warning) [14, 15].
[33] Reaching and operating within air-gapped networks has promoted worm-like features (notably used in Stuxnet promiscuously, Duqu 2.0 internally, and WhiteLambert as a module designed to control a large number of infections internal to a network).
[34] WannaCry [16] being the most notable case; followed by more limited self-spreading malware like NotPetya (a.k.a. Nyetya or ExPetr, [17]) and BadRabbit [18].
[35] As an example of extreme conspicuousness, WannaCry aroused researcher suspicion within less than a day of its infection attempts reaching scale.

products. It is also subject to the manipulation of the attacker, whose tools and operating procedures are likely designed to subvert or abuse some security measures.

While the overt antagonist to discreetness is loss of availability, the natural antagonist to discreetness is persistence. Persistence measures and an insistence on remaining on premise for an undefined period of time have proven the folly of most covert cyber intrusions. It's likely that a myriad stealthy attackers have *come and gone* through desirable networks without ever being spotted by virtue of focusing on a delimited task and embracing a quiet exit without unnecessary persistence.

The default discreetness of these cyber incursions allows for more interesting dynamics to play out between attacker and victim. This includes the potential for manipulation of data and sabotage of devices while keeping the attacker's hand invisible. It also allows lengthy on-premise operations to take place long before making the victim aware (by a loss of availability, as in the case of targeted ransomware and wiper attacks). Victim awareness is often something optional that the attacker can trigger if it fits their schemes.

However, not everything is under attacker control. Unlike other domains, the cyber domain invites rampant vigilantism and incentivizes unforeseeable interactions from third parties invested in securing the space and dismantling[36] its overall offence potential. While the attacker may operate covertly and the victim may not be in a position to discover the intrusion on their own, third-party defenders and researchers may still stumble upon artifacts of the operation and bring it to the attention of the victim, or the public at large.

Security companies – acting as third-party intercessors for the victim – find justification in either a specific relationship with the victim or in a larger remit to defend an essential service or cross-section of the Internet from attack. Once again, this quirky dynamic of the fifth domain doesn't find a precise corollary in others and complicates covert action. While all responsible governments have an interest in the overall wellbeing of the Internet, the interference of security companies with state-sponsored operations originating from 'friendly' countries remains a point of contention [19].

Decoupled identity

Finally, identity complications in the fifth domain extend beyond the issues with proprietariness and target-to-device mapping previously discussed. Just the same as we can't readily identify what vulnerable devices correspond to what real-world target personas, institutional identity is no easier to verify with certainty. Where location and overt signalling tend to suffice in identifying institutions in the real world, the fifth domain lacks this overt certainty and instead relies on cryptographic machinations to prove functional identity where necessary. In simpler terms, while anyone can find the *Microsoft* offices in Redmond, ascertaining that a domain, server, or piece of code legitimately belong to *Microsoft* is far more complicated.

The fifth domain's capacity for replicability is ostensibly stymied by proprietary access to certificates, keys, and sole control over 'legitimate' domains. These artifices are the sole signifiers supporting the claim that a server contacted or a piece of code executed comes from the stated source. This entails that the cornerstone of institutional identity comes down to the institution's ability to safeguard digital tokens equally susceptible to theft and replication.

The scenario boils down to: customers trust a company like *Microsoft* to create and support reliable, benevolent, and quality-assured code and services, so all we need is a means of making sure that we can extend that trust to the code or service in question by proof of provenance. If a cryptographic checksum proves that an executable belongs to a trusted developer, then that code should be trusted to the same degree that users trust the developer. Trust is the vaccine intended to inoculate the project of general-purpose computing turned promiscuous execution. By logical extension, that trust implies that we trust the developer to safely protect these private keys, certificates, servers and domains that underlie proof of ownership mechanisms.

However, not all developers and trusted institutions have proven so cautious in securing these precious trust signifiers. Cunning attackers have made away with code-signing certificates from development companies of all stripes, including: gaming companies [20], hardware and driver developers [21, 22], and the victims of campaigns designed entirely to steal these precious certificates [23].

After all, why fight mathematically secured trust-based models when they can be co-opted to serve the attackers? Examples of abuse range from the subverting of lax verification processes for issuing SSL certificates [24], to more brazen direct attacks on certificate authorities [25]. Perhaps the most reckless example is the subverting of both a cryptographic paradigm [26] and the bedrock update mechanism of the *Windows* operating system itself [27] in order to further spread an infection [28] within a victim network. The possibility of decoupling and abusing identity in the fifth domain renders trust-based execution models a qualified boon for attackers.

THE ART OF PERPETRATOR PROFILING

With the bleak realization that most of the features of the fifth domain disproportionately favour misdirection, deception and offensive practices, we must highlight the importance of adopting research methods with a greater propensity to create reliable and dynamic threat actor profiles. Two realizations shaped the current convention of threat intelligence publications: that well resourced attackers are seldom deterred by being publicly outed, and that outing a threat actor does more to degrade researcher visibility than to disable the threat actor's operational capacity. As such, rather than waiting for 'solid' attribution claims or naming normalization across vendors, mature threat intelligence producers embraced publishing technical characterizations of nebulous clusters of malicious activity as code-named threat actors.

While the practice of using code-name actor profiles has frustrated non-technical spectators who feel the stable of attackers is overpopulated with fanciful names, it has allowed defenders to track their attackers and coherently share

[36] For both selfless and self-interested reasons.

actionable information without engaging in unproductive flights of attribution fancy[37] any more than necessary. While equating Turla with the FSB or Equation with the NSA may bear fruits in second-order realizations about attacker intentions and operational legitimacy, it does little to enrich the technical indicators[38] themselves. Moreover, the bias generated by the supposition of familiarity with the perpetrator of the operation tends to misdirect research efforts long before they've reached a point of desirable rigour.

The ability to technically describe an undefined attacker in a timely fashion and provide the means for defenders to track that actor in their networks is of greater importance to an interconnected world under siege. Bike-shed arguments about the impropriety of animal names and the importance of information sharing come from an understandable place of frustration but they do not reflect the operational necessities of the primary intention of this work – to enable more efficient and effective defence. While academics, journalists and historians are encouraged to competently enter the space for their own studies and abstractions, most have yet to do so with adequate expertise. Pure technical research sources remain sparse and should not be diluted or held back in the service of other disciplines.

Rather than suggest better taxonomies or cumbersome frameworks, the notion of improving our current adversary profiling practices is to be rooted in an understanding of where those practices currently stand. The intent is to encourage deductive clarity and embrace dynamism. A desirable next stage research paradigm would do well to expand the circle of authorship to all those with a capability and a vested interest in order to match the continual state of flux of the adversary as object-of-study.

Addressing heuristic plasticity

Let's recognize that which the industry currently does well: generating narratives of threat actor campaigns and intended targets, sharing these relatively openly and sometimes freely, and thereby providing enough technical understanding of these nebulous clusters of otherwise uninteresting activity hidden amidst billions of indicators. As researchers, it's not the armchair assessments that we need to push back against as much as the practice of one-off, non-replicable research that limits the scope of our accomplishments against active threat actors.

Market forces incentivize researchers to move on to the next hot item, regardless of the superlative merit of the research they may have only just published. To entities affected or continually targeted by the adversary in question, this can feel like a hit-and-run – a sudden fervour of technical support that vanishes as quickly as it came. Given a similar propensity for fast flux in actor composition, tooling, targeting and tradecraft, research methods would do well to place greater emphasis on promoting the ability to track the particular artifacts that have served as bases for their assessments long after the original researchers have moved on to other projects.

Much the same as actors are dynamic entities subject to change, so must our assessments be non-static profiles friendly to correcting, extending and updating. Extensible research requires:

a. Access to the relevant data.

b. The clear identification of the bases in data on which assessments were made.

c. The ability to treat this data dynamically.

Though (a) may appear to be a given, that is not necessarily the case in an industry where access to data is sold and research data can come from sources limited by legal liability and the particulars of sharing agreements. While one cannot in good faith require vendors to give data away if they do not deem this beneficial, the industry would do well to insist that threat intelligence vendors selling reports provide access[39] to the relevant data to an extent that would enable in-house analysts to replicate the research efforts.

Secondly, (b) refers to the importance of drawing clear lines from the source datum that serves as the logical foundation of an inference. This may appear elementary but it's not always the case in current practice. Reports tend to cite a bulk of supporting data without necessarily pointing to the specific piece of that data that supports a specific conclusion. Binaries vary amongst themselves in build and configuration even within the same malware family. A telling screw up or misconfiguration in one binary won't be present in all others.

More importantly, second-order data like historic domain resolutions and registration information derived from specific third-party services is often cited without sourcing. Contrary to popular belief, even the best providers of this data don't contain the same information as their competitors[40], which makes specific leads hard to track down after the fact. Additionally,

[37] This instinct reflects the intuition that single-source (i.e. cyber-domain-specific) factoids cannot reliably point the finger at a real-world attacker, and even if it could, knowing the identity of a nation-state sponsored institution halfway around the world does little to defend a non-governmental network under siege.

[38] We'd be remiss not to also scrutinize the obsession with indicators of compromise. These indicators allow other researchers to verify the work described in a publication. They also allow defenders to watch for very specific components in their perimeter. In the latter case, this is easily subverted by tactically scrupulous threat actors. Without behaviour signatures or operational context, indicators alone are a weak auditive measure and unworthy of their exalted place.

[39] That is not to say that this data should be distributed indiscriminately. It's well established that attackers are ardent fans of research publications and would surely welcome input on how best not to get caught going forward. However, means of direct exchange with relevant customers are already available. The more we choose to divulge hunting and tracking methods, the more we may come to see the paywall as a needed buffer rather than the enforcement of a defender class-system. Perhaps a less controversial measure is the use of API key-dependent access, allowing defenders access to relevant data with existing subscriptions to shared repositories.

[40] For example, *DomainTools* and *PassiveTotal* are both excellent sources of data for profiling command-and-control servers. Various notable cases of OPSEC failures have interchangeably been discovered in one source and not the other. This speaks to the importance of researchers having access to as expansive a visibility as possible. And the need to reproduce this data (with adequate citations) during exposition.

information on the Internet isn't there forever. Reproduction with citations circumvents this issue and refers research teams to additional important sources.

Finally, (c) the need to embrace dynamic exposition is paramount and not trivial. The human brain is a fantastic tool capable of filling gaps and naturally enriching information in an accessible format to draw inferences. However, when it comes to references like hashes and IP addresses, the mind is often at a loss. Moreover, even when equipped with the greatest eidetic memory, information like domain resolutions and related domains and IPs change with great frequency and without warning. Similarly, trends in malware family adoption, deployment, and historic visibility change unexpectedly[41]. Our current methods of exposition tend to be very hard-copy-centric and do not embrace the capabilities enabled by the modern software medium.

This may appear as a simple matter of presentation style or adornment but when it comes to analysing data in bulk, highlighting real-time changes and the changing results of second-order queries has drastic tangible value. The research community would do well to embrace current promising attempts to do this[42] as a means of sharing information dynamically for the consumption of customers and research teams alike, and thus staving off the limited shelf-life of the research product. Noting drastic changes in adoption trends, developments in codebase, infrastructure registrations, redirections and configuration changes, and historic versus present domain resolutions better enables researchers to stay apprised of threat actors of interest.

Co-opting insights from criminal behavioural profiling

With the prerequisites for extensible research spelled out, we can address the larger issue of adversarial profiling methodology. Though we currently do a fine job of showcasing discernible elements of research into different campaigns and how these may relate to a context of other known campaigns, it seems our efforts right now are generally unstructured and 'artisanal'. Ideally, we should create a profiling methodology with clear guidelines for fruitful avenues of research that yield compatible, and thereby comparable actor profiles across vendors.

Though an initial hunch pointed in the direction of some sort of military profiling practice, this effort proved misguided. As compatible as that may be with the conception of cyber as a war-fighting domain, it would gloss over the larger set of possible adversarial configurations particular to fifth domain adversaries, thus tainting the object of study with unwarranted preconceptions. Instead, we'd do well to take some cues from a more open-ended form of profiling – that of behavioural profiling as developed for the purposes of investigating violent crimes – thus avoiding an over-regimented (and perhaps over doctrinarian) conception of who is on the other side of the proverbial keyboard.

Understanding the whole of behavioural profiling as a discipline (with its nuances and controversies) is beyond the scope of our current endeavour. However, some clear corollaries arise when we consider the desirable outcomes of this practice. It focuses on the notion that an analysis of both crime scene and patterns of behaviour enables profilers to formulate hypotheses to aid criminal investigations in narrowing down a pool of suspects, thus shaping the deployment of resources in approaching those suspects[43].

Given how television depictions have tainted our collective understanding of the practice of behavioural profiling, it's important to frame a functional understanding of the investigative contributions [29] expected of a behavioural profiler engaged in a criminal investigation.

The primary function of behavioural profilers is to investigate a crime scene(s) with the intention of *generating hypotheses* that can help narrow in on a likely perpetrator. Multiple hypotheses are noted along with their supporting materials so that they can be tested systematically throughout the course of the investigation. The practice of predictive profiling, most closely associated with behavioural profilers, involves not just a description of likely age and propensity for past criminality but also geographic profiling[44]. The predictive profile is meant to help nominal generation, a selection of pools of likely suspects from relevant civilian and criminal databases.

These pools can then be narrowed down with the use of prioritization matrices, which give individual predictions numerical rankings such that they can objectively be tested against the background and features of a nominal pool of potential suspects. Where there is a lack of material evidence linking potentially serial crimes, offence linkage analysis can help to determine the consistency or variability exhibited by the offenders during the course of committing a series of offences.

A favourable outcome of proficient behavioural profiling is not the immediate clairvoyant resolution of an ongoing investigation or a cold case but rather any or all of the following contributions:

- *Search advice* – using the predictive understanding of the likely offender to inform the search parameters of a forensic investigation, the collection of evidence, witness interviews and veracity assessments, and the discovery of likely body deposition sites.
- *Investigative suggestions* – based on a combination of previous experience and familiarity with criminal

[41] One of the reasons I (and my colleague Costin Raiu) have insisted on the importance of historical malware research – or malware paleontology – is that visibility increases unexpectedly over time. An increase in awareness and available signatures actually draws out more malware submissions and incident response information than was originally available. While the furore of initial discovery tends to draw the most research effort, more data becomes available over time and rewards a second look.

[42] Current notable solutions in this style include *Maltego*, *VirusTotal Intelligence Graphs*, *DomainTools Iris* and *PassiveTotal Projects*.

[43] The desirable outcome is noteworthy and will be discussed further as it applies to our own cyber adversary profiling.

[44] A notable example of material conditions serving as a solid starting point for profiling offenders is David Canter's *Circle Theory*, which entails an analysis of spatial patterns of criminality (in the investigation of arson, rape, murder and burglary) not only to predict possible locales of future crimes, but also the likely provenance, familiarity, and mode of transportation of a likely offender.

investigations, along with logical inferences based on the predictive profile developed, profilers provide suggestions for investigative avenues. These suggestions must be accompanied by a clear rationale and are developed on a case-by-case basis.

- *Risk assessment* – by extension of the offender profile, profilers can assist in determining the circumstances (social, environmental and interpersonal) that may increase the risk of repeat offences in order to help mitigate ongoing risk to relevant communities.
- *Media advice* – behavioural profilers can help maximize the utility of the media in cases where controlled releases of information and appeals to the public can be utilized not just to motivate public cooperation but also to predict its potential effects on the behaviour of the perpetrator.

While this is a cursory oversimplification of a complex and not uncontroversial field of study, the expected contributions of criminal behavioural profiling can help structure a functional profiling methodology for fifth domain actors as well[45]. It should also help mediate a more realistic expectation of the desired outcome of prolonged investigations into these cyber adversaries, as an agglomeration of hypotheses, observations and data intended to guide timely investigations, all the while enabling the management of perceived risk, propensity towards further incidents, and the intelligent deployment of limited resources to mitigate further exposure.

A new analogue for cyber threat actor profiling

A central thesis in criminal behavioural profiling is that 'behaviour reflects personality'. Our parallel operating thesis is that *operational behaviour and tooling reflects adversarial configuration and imperatives.* Though we know that many of the atomic indicators involved in our incident response engagements are ultimately fungible, a comprehensive understanding of the profile of the likely actor involved can help guide us out of the mire of misdirection, or at the very least structure our insights into testable hypotheses open to future observations.

There are epistemological limitations to fifth-domain indicators that mean that we cannot base a solid attribution claim on fifth-domain indicators alone. While a well-constructed criminal profile can narrow down a nominal pool of suspects, it alone does not finger a specific perpetrator. Similarly, a comprehensive cyber threat actor profile should narrow down a nominal pool of suspects for other researchers, academics, and spectators at large while allowing for the state of flux of an ongoing investigation almost certainly assured to witness other incidents by the same threat actor in due course.

In a domain characterized by unlimited replicability, researchers should not pursue the golden standard of identity attribution. What our actor profiling should yield is an approximation for the functional purpose of enhancing the ability of defenders to deploy their limited resources intelligently by predicting the likely behaviours of these undeterred repeat offenders. Meanwhile, it may also allow all-source investigators and relevant stakeholders to narrow a pool of suspects and correlate data beyond the domain of pure cyber to which they alone have access.

It's important to remember that we aren't dealing with the same object of study as criminal behavioural profiling. In our case, the criminality of the incident is disputable[46], the intent is not always discernible, and the punishment largely unenforceable. At this time, the only deterrence known to work in cyberspace[47] is that which cannot technically be accomplished. This largely unattainable standard is up to network defenders to approximate, by raising the annoyance and cost of an incursion into their perimeters.

Much like circle theory in behavioural profiling takes advantage of the expression of geographic and logistical familiarity in related criminal offences to profile a perpetrator, so can we take advantage of observable operational traits to profile the actor in question and describe what we may come to expect from them going forward. Criminal profiling finds some of its roots in interviews volunteered by perpetrators in their accounts of atrocious serial crimes. In our case, although we don't have the opportunity to reliably interview past perpetrators, we do have access to an abundance of case studies by researchers that have pieced together campaigns based on objective indicators. With an eye for promising traits, perhaps we could mine these past case studies for comprehensive profiles.

Promising profiling characteristics are already interjected in passing in most research publications and include:

a. Proficiencies and deficiencies

Proof of proficiencies or deficiencies in the tooling, infrastructure, and operator interactions.

- Programming languages used in development and deployment:
  - Do these include more difficult or stringent programming languages?
  - Do their coding conventions employ unusual or outdated libraries?
  - Are the languages themselves more popular with younger developers?
- Configuration choices and consistency in the registration of command-and-control infrastructure:
  - What is the operating system of choice?
  - How are these configured and secured?
  - Is the registration style consistent?
  - Is the infrastructure maintained well over time?
- Is the code utilized likely developed in-house, borrowed code that's been modified, or entirely copied?

[45] As with any suggestion of a new way of approaching a known problem, I expect that this functional comparison only begins to scratch the surface of the parallels that can be drawn between both disciplines, but I suspect that many of these points touch close to home for other practitioners.

[46] Depending on the territory, the material purloined, and the institutional remit of the perpetrator.

[47] Short of a well-placed hellfire missile.

- Are more complex or finicky components utilized?
  - i.e. rootkits, kernel components, memory injection, or novel means of subverting security solutions.
- Concern over security solutions and interactions with researchers:
  - Are measures undertaken to undermine anti-virus products, automated security solutions and dynamic analysis tools?
  - Are efforts employed to hinder or subvert manual analysis specifically?
  - In the case of attackers that have previously been exposed in wide-distribution publications, did their operations change?
    - › Was the change drastic and consistent over time?
- Capabilities exhibited in the use of common infection vectors:
  - What do the infection vectors tell us about the operators and the intended victims?
  - Are the operators familiar with their victims?
  - Are they familiar with the local culture?
  - What are their language skills?
  - Are they focusing their attacks on technologically illiterate or socially vulnerable victims?
  - Is there a discernible reconnaissance effort?
    - › Are they already familiar with the type of systems used by the victims?
- Operator practices:
  - Are their operator practices consistent?
  - Is there on-keyboard involvement? If so, are mistakes exhibited?
  - Alternatively, has the process been automated?
- Exfiltration practices:
  - Do they exfiltrate all at once and in bulk?
  - Is exfiltration staggered or made to conform to standard network operations?
  - Is the volume indicative of extensive backend storage infrastructure?
  - Does it entail a lack of familiarity with what may be considered valuable in the victim enterprise?
  - Alternatively, are items of interest filtered on the victim system before exfiltration?
    - › Does it represent a content-matter-specific or persona-specific constraint, or possibly a legal limitation?

b. Outfit configuration traits

Discernible traits that more closely define the composition of the adversarial outfit, its resources, and its tasking relationship.

- Regional traits:
  - Are operational time zones discernible from artifacts or connection times?
  - Do these exhibit work in shifts?
  - Do they regularly observe regional holidays?
- Operational scale:
  - Are there indications of the number of operators and developers involved?
  - What is the scale of the actor-controlled infrastructure?
  - How many victims are actively infected and maintained at a given time?
- Operational tempo:
  - How often is infrastructure changed or redeployed?
  - Are infections carried out in waves?
  - What is the tempo of retooling?
  - At what rate are new campaigns undertaken?
- Targeting style:
  - Are the intended victims well chosen?
  - Is an organization blanketed with infection attempts?
  - Is there an attempted phase of lateral movement to discover an intended system or victim?
  - If so, are other infections cleaned up or are infections rampant?
- Developmental prowess:
  - Is the tooling employed developed in-house or outsourced?
  - If outsourced, is the tooling commercial off-the-shelf (COTS), high-end third party, or contractor-sourced?
  - If in-housed, is the codebase actively maintained?
  - If COTS malware is employed, is it consistent with the resources exhibited in other operational components?
- Operational purity:
  - Are the operators, proprietary tooling, or infrastructure being employed in seemingly unrelated activities (i.e. criminal moonlighting)?
  - Are targeted infections mixed with small financial gain opportunities (like the deployment of cryptominers on victim premises)?

c. Outfit autonomy or dependence

Indications of the adversaries placement in a larger institutional structure or a delimited functional role.

- Cluster tool sharing:
  - Is a specific implementation of closed-source tooling being shared amongst different threat actors within a seemingly related cluster?
    - › Is there a ‘digital quartermaster’ [30] involved?

  - What is their exhibited level of care for rare and expensive tooling?
    - › Do they protect their zero-day exploits?
    - › When one of these is burned, do they come up with another shortly thereafter?
- Shared tasking:
  - Do multiple threat actors from within the same cluster attempt to breach (or effectively breach) the same target around the same time?
- Staggered operations:
  - Does one team share their access to a given victim with another threat actor?
  - If so, are these handoffs consistent with a functional separation between teams?
- Larger geopolitical circumstances:
  - Do visible changes occur at the advent of geopolitical events in relevant regions?
  - Does targeting shift with certain geopolitical events?
  - Does team composition shift with more traumatic geopolitical events?
  - Are there indications of changes to the adversarial outfit to match institutional restructuring in suspect organizations?

These are discernible traits visible in most investigations to a lesser or greater extent and they help us build a comprehensive and comparable profile. The latter should subsequently enable us to test hypotheses for second-order deductions and for the predictive deployment of limited on-premise resources. Fortuitous geopolitical changes will allow us to bolt down some of our operating hypotheses, but even when these events make institutional attribution seem certain, threat actor profiles should remain responsive to observable indicators.

ADVERSARIAL CONFIGURATIONS AND FIELD OBSERVATIONS

Finally, if the intent is to generate a predictive profile that'll help weed down a pool of suspects, we'd do well to become familiar with the possible 'suspects' that could populate that pool. While we too could settle for a nominal pool of likely adversarial institutions by name and reputation, a more productive approach from a defender's standpoint should focus on familiarity with the different possible adversarial configurations we are likely to encounter during our investigations. It's important to remember that the goal is not to 'nail the culprit' but rather to create a dynamic profile of testable hypotheses that will serve investigators long-term as the adversary continues to operate and change.

Configuration pool

While nation-state adversaries are the de facto boogeymen of cyberspace, nation-state adversaries are in no way standardized, nor do they operate as a homogeneous entity. They exhibit differing skills, operational prowess, intentions and capabilities. And their configurations are as diverse as one might expect of nations at different stages of socioeconomic development racing to adopt specialized weaponry. The main axes of variation are the internalization or externalization of operational functions. These material conditions will affect traits visible to researchers and investigators in the wild.

The nation-state moniker is loose or often downright inaccurate. Similar to 'sophisticated', a 'nation-state' attacker is meant to signify any combination of the following: the suspected investment of extensive resources, the blending of operations with the use of non-cyber resources, or the pursuit of areas deemed of sole governmental interest (such as politically motivated, strategic resource-oriented, law enforcement, diplomatic, or counter-terrorist targeting). None of these metrics alone definitively designate a threat actor as being an official part of a nation-state apparatus, which motivates the broadening of the term to 'nation-state sponsored'. So as to say, the threat actor may not be a part of the state apparatus but is its beneficiary and contributor.

From the perspective of the researchers, the exact dynamics at play are not always discernible through fifth-domain indicators alone. However, we can model configurations previously encountered, described, rumoured, or reasonably imagined:

a. Nation-state (internal): the operationally purest nation-state threat actor is one that has all functions in-house. That's to say, an official arm of a governmental institution with a cyber remit whose tooling and infrastructure are internally developed and managed. To say that tooling is internally developed may be giving these too much credit. In most cases, these teams don't have access to capable developers. Their tooling largely consists of commercial off-the-shelf (COTS) malware or poorly coded RATs, old exploits, and spear phishing.

From an institutional perspective, the idea of relying solely on in-house talent holds a greater promise of operational security. However, this configuration is prone to rudimentary mistakes, partly from lack of high-end talent but also a lack of proper resources. In some cases, even the COTS malware is pirated or generally unsupported. The infrastructure is no better managed. Observed infrastructure configurations include compromised private servers, the abuse of cloud services as staging servers, or the registering of servers of their own.

Careless testing practices or unsecured backend connections have resulted in teams with this configuration being traced directly back to physical governmental premises[48]. These public failures may result in the institution losing their operational remit; they can also result in increased budget and upgrades to 'avoid future failures'. It's important to note that the 'internal' configuration is likely temporary. As a government starts to test the waters of fifth domain operations for intelligence collection, successes are likely to carry added budgets and result in upgrades to external tooling.

[48] Dark Caracal is a notable recent example where researchers were led directly to buildings belonging to Lebanon's General Security Directorate [31].

b. Nation-state (blended): this is the most common nation-state configuration, as internal operators run campaigns with the use of externally developed and supported tooling, externally sourced exploits, and/or externally maintained infrastructure.

Great furore has erupted over the commercial providers of mid-range tooling for these configurations. Notable companies in this trade include Hacking Team, Gamma Group (known for their FinFisher suite), and more recently NSO Group. These companies sell a more complete package that includes operator training, backend infrastructure to manage collected data, ease of configuration for implant logic and compilation, automated anti-analysis measures, 'anonymizing' chains of infrastructure redirectors, and even zero-day exploit subscription services[49].

These services offer a convenient way[50] to increase the capabilities of diverse governmental institutions looking to step up their operations without previous experience or extensive development capabilities. However, the arrangement also exposes these institutions to be uncovered en masse. As researchers (and notable hacktivists) close in on the software suites themselves and their respective providers, the resulting research is likely to reveal entire swaths of customers whose operations had not otherwise been exposed.

It's important to note that, just as emerging threat actors may suddenly upgrade to using the improved tooling of these mid-range providers, evolution is not a one-way road. Subject to fluctuations in budget allocations and institutional favouritism, these subscribers may find themselves going back to lesser tooling (perhaps with the added experience of having had access to better tooling for a period of time).

A lesser-known higher tier of commercially supported tooling exists in the tight-lipped realm of defence contractors. These outfits provide far better tooling, including advanced anti-analysis measures, specialized functionality, and overall better quality-assured code. The resulting product is almost always a modular framework, designed to be upgradeable, adaptable and expandable. These frameworks are designed to specification and are thereby made to be forwards- and backwards-compatible with other similarly sourced frameworks. One can only wonder at the cost of these frameworks but the resulting product is almost always of an admirable build quality.

As these arrangements rely on operations being run in-house, other services are also provided to assure a greater level of operational security such as programmatic infrastructure registration and maintenance as well as services meant to anonymize connections[51] to servers and implants on the field. The best threat actors rely on a combination of all of these quality-assured services alongside extensive in-house talent to both design specialized components as well as design and implement the operations themselves.

It's important to note that some of the threat actors that benefit from these military/defence contracting arrangements also happen to abide by the most stringent legal frameworks. The involvement of lawyers and external oversight is often reflected in the implementation of stringent measures to avoid infecting the wrong target, to automatically disable infections within certain territories, and to discontinue functional operation or disinfect after a certain deadline when their legal approval presumably expires.

c. Nation-state (external): the rarest arrangement is that of an all-externalized operational remit. That is to say, nation-state contracting an external outfit to manage both tooling and operations. From a counter-intelligence perspective, few institutions should be comfortable with an external institution being aware of their tasking requirements, managing the means of infection, and collection. However, this also allows a level of plausible deniability. It also likely involves a 'murky' legal area. But given that in many emerging countries relevant laws remain largely unwritten, this configuration should not be written off entirely.

A more plausible implementation of an external configuration is the deployment of operators overseas to circumvent blockades, sanctions, and domestic infrastructure limitations[52]. Such cases may well involve the need to generate funds illicitly, to 'earn their keep' in a privileged position overseas. This entails possible evidence of official operations overlapping with criminal ops meant to generate revenue.

The acknowledgement that nation-states are capable of some level of comfort when it comes to outsourcing operations brings us to the contentious topic of mercenary configurations. That is to say, the co-opting of external outfits[53] to (wittingly or unwittingly) fulfil some degree of nation-state requirements.

Mercenary configurations (criminal): criminal outfits are exceedingly susceptible to co-opting by unscrupulous nation-state institutions that can offer protection from liability in exchange for carrying out operations with state interests[54]. This sort of arrangement provides a greater level of plausible deniability for nation-state interests that can simply cast blame on the criminals engaging in everyday criminal activities (like botnets and credential theft). All the

[49] This offering includes an interesting nuance – customers are often not sold the exploits directly. They are offered a length of coverage. The providers take care of weaponizing the desired lure and malware and supply the weaponized product through alternate infrastructure that is then served to the victims. This is a means of protecting against the carelessness of inexperienced customers and their tendency to burn these exploits.

[50] i.e. 'throwing money at the problem'.

[51] However, no system is foolproof. The failure of one such solution is noted in a leaked CSEC slidedeck on MAKERSMARK (i.e. Turla) operators under the fitting observation: 'Designed by geniuses, Implemented by morons', *Cyber Leads to CI Leads*, Slides 6-10, (Classification warning) [32].

[52] As is often rumoured of North Korea's hacking teams said to be operating abroad for the combined benefits of better connectivity and plausible deniability. Surely, the general availability of food is an added perk.

[53] A practice noted as early as 1986 [33].

[54] As suspected of wanted man Evgeniy Bogachev, alleged developer of GameOver Zeus [34].

while, the same malware and infrastructure can selectively be employed to collect more desirable information from the same infected systems.

Mercenary configurations (hacktivist): the lack of organization, anonymity, and 'democratic' decision making of hacktivist circles has made them fruitful grounds for nation-state operators looking to co-opt capabilities (largely by means of deception). The resulting operations are likely to serve as background noise to mask the main operations or as a symbolic 'groundswell' to support political measures.

Mercenary configurations (private sector): a radically charitable reading of recent involvement of private sector companies in information warfare operations at the behest of government actors suggests another blended configuration where nation-state threat actors co-opt (witting or unwitting) private sector partners to support wider operations.

Mercenary configurations (former/parallel): a largely undiscussed configuration is that of former nation-state operators and parallel investigation teams serving nation-state requirements away from their oversight or legal limitations. The result is a difficult to attribute mess of operations whose tasking reflects the interests of one or more governments as well as private sector and criminal interests[55].

Former agents are a difficult to manage byproduct of military and intelligence organizations[56]. They are highly trained and possess historical and institutional knowledge. Depending on their past operations or the conditions of their exit, these individuals may also have a hard time finding conventional gainful employment. Former operators of nation-state outfits form a fruitful recruitment pool[57] for mercenary hacking teams for unsanctioned nation-state interests.

These adversarial configurations are not exhaustive but should provide archetypes for the kind of entities we're actively profiling in our investigations.

Revisiting known 'perps'

In the interest of encouraging the research community to engage in this kind of adversary profiling, it seems fruitful to point at interesting features of past research that may have been overlooked or lost in the overactive shuffle of threat intelligence publications. In the following sections we look at interesting features of different notable threat actors, some still active; each actor is notable in its own right. In some cases, geopolitical events, indictments, and leaks will even provide the rare opportunity to peel back the curtain and test the accuracy of our assessments as external investigators. These research leads[58] hopefully not only highlight the importance of looking back but also the need to work with testable hypotheses. The goal is an adaptable, testable, responsive actor profile. Stopping short of pointing the finger at a definitive suspect is important when dealing with a dynamic adversary that can and will change configurations over time as well as in response to research publications themselves.

[55] Suspected of the notorious Wild Neutron (a.k.a. Morpho, Butterfly, or ZeroWing) [35].
[56] Reports abound of former Latin American operators blackmailing their way into working for multiple governments in the region as an unsanctioned mercenary team for anti-democratic and politically motivated collection operations [36].
[57] That is, in countries other than the United States where intelligence community alumni are exceedingly desirable to the private sector.
[58] Which any interested reader is encouraged to take on as a project of their own.

On cleanup

Equation Group: staggered cleanup

In 2015, *Kaspersky*'s *GReAT* team announced their discovery of the Equation Group. The extensive breakdown of Equation malware frameworks and infrastructure included among them a decade-old platform that the researchers codenamed DoubleFantasy [37] (and that, we now know, was internally referred to as 'VALIDATOR' [38]). Despite this being an early platform largely replaced by TripleFantasy [39] (or 'COMMONDEER' [40], a.k.a. 'CODE'), infrastructure remained active and some of it was sinkholed by the researchers. Sinkholing is a fairly obvious indication for careful threat actors that researchers are closing in on their operations and that, if any active infections remain, it's time to clean those up or risk not only exposing their tasking but other undiscovered parts of their toolkit.

Despite the extensive reporting on known infrastructure and the discovery of DoubleFantasy, the Equation Group operators did not, in fact, disinfect remaining targets. The *Windows* version of DoubleFantasy had been documented and its discovery led to the discovery of the *Linux* equivalent. Then an active infection calling out to a sinkholed domain[59] with a different User-Agent revealed the existence of a *Mac* version of Double Fantasy in a research institution in Latin America[60]. It's possible to consider this single instance an oversight in cleaning up a sprawling set of infections after a mass discovery.

However, in late 2016, following the unfortunate series of ShadowBroker releases, *Antiy Labs* [41] revealed[61] its discovery of multiple Equation Group infections in China. Among these were the *Linux* version of DoubleFantasy as well as yet another undiscovered variant designed for Solaris 10 (SPARC).

The 'staggering' of cleaning up a thoroughly burned platform by a systematic attacker could be interpreted in a variety of ways. The following is an example of the sorts of inferences we can draw from this case (and could go on to test with relevant data):

Hypothesis 1: The threat actor failed to consider architectural variants of the same platform as related and thereby did not consider unreported variants as compromised.

- **Pro:** Active infections reported were specific to unreported non-*Windows* variants.

[59] The implant was configured to use multiple command-and-control servers, meaning that the sinkholing of a single domain would not prevent the operators from cleaning up the infection.
[60] DoubleFantasy OSX was privately reported to *Kaspersky* subscribers in December 2015.
[61] Sadly, despite showcasing competent reversing work, the report failed to include adequate references to hashes or command-and-control infrastructure.

- **Con:** Active *Windows* infections after publication would undermine this hypothesis (relevant sinkhole telemetry specifics are not public).

Hypothesis 2: The attacker lost control of these specific implants before they could, in fact, be removed from the systems.

- **Pro:** The sinkholing of the command-and-control server may have made this more complicated in the case of the *MacOS* variant.
- **Con:** The *Antiy* report was released nearly 18 months after the *Kaspersky* publications.

Hypothesis 3: Despite a shared platform, operator crews are not under a single institutional umbrella. This entails that cleanup efforts would not be uniformly treated across different teams operating with the same platform.

- **Pro:** A leaked CIA post mortem[62] of the discovery of the Equation Group by *Kaspersky* researchers suggests that the Equation Group umbrella in fact included the operations of different organizations within the same government.
- **Con:** The same leak refers to the degree to which measures undertaken are standardized[63] within the relevant institutions, signifying that a cleanup effort would likely have been followed with great care.

Hypothesis 4: Having already suffered vast exposure at the hands of both researchers and leakers, the institution decided to get as much operational time out of their active infections as possible.

- **Pro:** Replacing active infections on systems that may already be under scrutiny (effectively turned into live honeypots) could expose an entirely new set of tooling to investigators.
- **Con:** Maintaining a clearly attributable infection on foreign systems could have geopolitical and diplomatic ramifications upon discovery.

Hypothesis 5: The actor does not actively monitor their infrastructure for loss of domain control or sinkholing.

- **Pro:** At least one active infection continued to beacon after sinkholing for a prolonged period of time and alternate means of control were not employed in disinfecting the system.
- **Con:** There was not a prevalence of infections at the time of the *Kaspersky* announcement, suggesting that the threat actor had already undertaken measures to clean up active infections and mitigate exposure. Alternatively, the threat actor may have turned onto improved tooling entirely and discarded the bulk of its operations upon the faintest indication of discovery.

Hypothesis 6: The actor (mistakenly) expected the malware to no longer be there.

- **Pro:** DoubleFantasy counts with a self-cleanup mechanism triggered if the command-and-control servers are unreachable for a certain period of time[64].
- **Con:** Hard-coded IPs were available in multiple samples, meaning that the original command-and-control server would likely still receive the beaconing of active infections in need of removal despite domain sinkholing.

Animal Farm and Careto: definitive cleanup

Threat actors display different attitudes when it comes to public attention. Animal Farm and Careto, two notable actors, responded to their respective discoveries by burning down their operations entirely. There are no reports of either of them having resurfaced since. Let's look briefly at each case:

Animal Farm is believed to have operated for at least seven years, largely targeting French-speaking countries and former French colonies. Despite previous privately reported discoveries[65], Animal Farm rose to public prominence with a publication in *Le Monde* [44] based on a leaked CSEC slidedeck. Subsequent technical research publications were released by *Cyphort* [45] (late 2014), *Kaspersky* [46] (2015), *ESET* [47] (2015), and *G Data* [48] (2016), detailing different components with whimsical names like Bunny, Babar, Caspar and Dino. Interestingly, despite a well established operational history and tempo, Animal Farm ceased operations after April 2014 and has not been seen publicly since.

Similarly, Careto[66] was also active for an approximate seven years of operations, targeting victims in more than 30 countries with multi-platform malware. Interestingly, *Kaspersky*'s discovery of Careto set off the burning down of the sprawling infrastructure before it had even been published. In January 2014, a month before publication [49], Careto operators began taking the infrastructure offline permanently. Their malware or operations have not been sighted since.

Particularly in the case of Careto, we see evidence of very tight monitoring of their infrastructure. The Careto operators likely discovered that they were the subject of active research due to suspicious activity contacting their servers and decided to burn down their operations before information had been widely distributed (and likely without definitively knowing that it would be published).

Now, what can we infer from these cases? At the very least, we can say that both threat actors place a high premium on limiting exposure. A profile derived from this inference might suggest the observance of military doctrine or institutional experience with highly regimented operational procedures necessary for

[62] 'What did Equation do wrong, and how can we avoid doing the same?' [42].

[63] This level of standardization actually proved the downfall of the Equation Group. By requiring the use of a specific encryption implementation (in order to avoid weak or poorly implemented encryption routines), the discovery of a single Equation sample allowed signatures to be crafted that would reveal more than a decade's worth of malware related to the same actor [43].

[64] This, of course, would fail to trigger in the case of sinkholing since the servers would, in fact, be reachable, just no longer operated by the threat actors themselves.

[65] By both the Canadian Communications Security Establishment (CSEC) in a classified report and *Kaspersky GReAT* researchers in a private report.

[66] Also referred to as 'The Mask', a likely mistranslation of the term 'careto' found in the malware, likely meaning an ugly face.

entrenched conflicts (like domestic terrorism or counterinsurgency) which require careful handling of burned embedded assets.

Does their disappearance mean that they are both actually gone? We can say with certainty that the vast number of signatures and countermeasures (as well as observance of their previous infrastructure practices) have not signalled their resurgence according to any public accounts. However, with operations of that level of well supported multi-framework tooling, zero-day exploits, and vast, well curated operations, it's hard to imagine that the relevant talent and their respective organizations would no longer be operative in some form.

It's possible that both teams have retooled, either by changing their frameworks entirely or by modifying their known frameworks to a new paradigm of stealth. The following is a case of the latter.

The Duqu Bet – a Phoenix from the ashes

The original Duqu operations were notorious as an accompanying salvo to the deployment of Stuxnet. Some code sharing between the Stuxnet and Duqu drivers [50] established their relationship further. Duqu [51], named for the prevalence of filenames with the prefix '~DQ', was discovered by *CrySyS Lab* [52] during an incident response engagement and extensively researched by *Kaspersky*'s *GReAT* [53–61] and *Symantec* [62]. Less than a month after its original discovery, Duqu operators began wiping their command-and-control infrastructure (some registered as early as 2009). With most samples compiled between 2010 and 2011, the complex malware platform appears to have been operative for a meagre two years before being shut down due to public exposure.

However, the investment in Duqu did not go to waste. Despite the appearance that Duqu had gone the way of Animal Farm and Careto, it turns out that the institution behind this campaign had opted to refactor. The resurgence of Duqu was discovered within the *Kaspersky* offices in 2015. The malware was believed to have been active for months and operated entirely in-memory. The redesign was extensive [63, 64] and not only included new countermeasures to avoid detection, worm the infection across a company network, and hide its means of control and exfiltration, but it also allowed for more than 100 plug-ins to be deployed from within virtual filesystems.

Interestingly, this is the only known instance of a disappeared APT resurging in this manner. It allowed a glimpse at what doubling down on the investment in a platform would look like – restructuring the malware to function entirely in a largely unmonitored element in personal computing (RAM) and to conduct its operations by means of hops between networked endpoints without care for persistence except in a central node.

A profile of this adversary would do well not only to focus on the retooling aspects of its history but also on the true unmitigated brazenness of its targeting. Duqu 2.0 was not only discovered in the offices of a reputable anti-malware company but also in the locations of the P5+1 meetings – where (presumably) allied governments met with Iranian officials to negotiate the Iran deal under the auspices of the Obama administration. Furthermore, these contentious attacks were carried out with the use of a well-known, and now directly attributable platform. If ever there was material to build a personality profile [65] of a cyber threat actor, Duqu's tasking should provide adequate material for speculation.

On development practices

Flame: just the ashes

Another amazing threat actor discovered around the same fruitful period as Duqu and Stuxnet was Flame[67], another extensible modular framework. Flame [66, 67, 68] is notorious in that it is perhaps the earliest discovered modular cyber espionage framework in the style we've become accustomed to with modern threat actors. That said, Flame was showing some of its age compared to Duqu, with the former likely operating since 2008 and relying on bulking binaries, close to 20MB in size. Flame was discovered by the Iranian CERT in collaboration with *Kaspersky* in May 2012, and by early June, the command-and-control servers were first redirected and eventually permanently shut down.

Unlike Duqu, it appears that Flame was essentially forced into a deserved retirement. What makes Flame interesting from a development standpoint is the comparative observation made by *CrySyS Lab* researchers that, when comparing the features of Flame and Duqu, it appears as if they were commissioned with the same specification [69]. The hypothesis is that two independent development teams were provided a similar specification to implement, resulting in the two different frameworks. To further fuel this speculation, there exists a transitive connection between the two precisely within the larger Stuxnet family. Where select Duqu drivers shared code with Stuxnet, so did a 2009 version of Stuxnet share code with Flame. Flame's resource 207 [70] provides a vague developmental link to Stuxnet and thereby transitively to Duqu as well.

The Dukes: erratic development

Managing in-house development over time is difficult. In the case of some threat actors (or threat actor clusters), that can mean making do with a revolving door of skill sets. For researchers interested in seeing a progression of spasmodic development, the Dukes[68] present a fascinating case study. Starting in 2011, the Dukes expanded their arsenal to include CozyDuke, MiniDuke, OnionDuke, CosmicDuke, SeaDuke and HammerToss. Many of these have similar functionality but are each developed in different programming languages, actively developing various tools at different times, and deploying them as needed complementarily. These include:

- MiniDuke, internally referred to as Nemesis Gemina [71], written in Assembly with techniques 'borrowed' from the well-known VXers[69] '29A' [72].
- CosmicDuke (or TinyBaron), written in C/C++ and built on a customizable framework referred to internally as 'BotGenStudio'.

[67] Also known as Flamer (*Symantec*) or SkyWIper (*CrySyS*).
[68] Also known as APT29, OfficeMonkeys, EuroAPT, or by the names of some of their malware families like MiniDuke or CozyDuke.
[69] Colloquial term for virus writers.

- CozyDuke [73] (or Cozer) includes payload DLLs written in C/C++.
- OnionDuke [74] is known for embedding a DLL payload written in C/C++ into executables crossing a certain malicious Russian TOR exit node.
- SeaDuke [75] comes into the picture in October 2014, with similar functionality to CozyDuke but written in Python.
- HammerToss [76] is comprised of .NET binaries capable of executing PowerShell commands.

At that level of developmental variegation, there's a lot of material for informed speculation regarding the internal configuration of the adversary in question[70]. These are multiple tools, some better maintained than others, regularly intermixed, but not always developed in a rising scale of complexity.

The Lamberts: parallel professional development

Acquiring a quality-assured attack framework is a costly endeavour. Threat actors in a position to acquire one of these are likely to both protect and leverage them to the greatest possible extent for an adequate return on their investment before attempting to source another. However, one notable threat actor consistently broke that paradigm by having multiple parallel development efforts. The Lamberts is the threat actor with the greatest OPSEC practices and general tradecraft of any discovered so far. This distinction was earned in part due to their practice of sourcing multiple advanced attack frameworks, leveraging and cycling them out without waiting to attract any kind of public attention. Amongst its arsenal, the Lamberts arguably counts at least four extensible modular frameworks out of the nine colour-coded families identified so far. Each of these frameworks counts on automated security measures, extensive quality assurance, and distinct functionality.

The Lamberts group has amassed a lavish toolkit over at least 13 years of operations. Covertness appears to be a paramount priority for the group, worthy of extensive investment. That imperative doesn't just factor into its spending practices but also into the design of its deployment tools, victim-box operations automation, and truly paranoid multi-stage encrypted droppers. This functional paranoia should factor as a central tenet in a comprehensive profile of this cunning adversary. Out of the many features of this complex threat actor[71] worthy of extensive study, we'll focus on its unique development practice.

Not only does the Lamberts group employ multiple frameworks at the same time, it appears that these frameworks are sourced from different development teams[72] and made compatible by a standardized specification:

- *Team A* – codes largely in Visual C++ with the use of custom libraries. Team A is responsible for the Green, White and Black[73] Lamberts[74]. Apart from the code overlap apparent from the shared use of these custom libraries, the families share a configuration format and encryption templates.
- *Team B* – a quirkier development team that codes in C++ with the use of the Standard Template Library (STL). STL predates the C++ Standard library and fell largely into disuse after the latter's rise in the late 90s. Team B is responsible for the Blue, Red, Gray and possibly Pink Lamberts.

The unlikelihood of such a well-resourced attacker, alongside two distinct and consistent development conventions, may lead spectators to surmise that this is the work of two distinct threat actors. However, decrypted configurations for samples from different families reveal that they're apparently being leveraged for the same code-named operations.[75]

These quirks suggest the presence of an overarching organization agnostic to the frameworks deployed. This gives us a glimpse of the externality of the development practices in the organization of the Lamberts. But perhaps it even suggests a division of labour between the organizers of the campaign tasking and the operators themselves. This would not be hard to believe with a threat actor that also displays the greatest discernment as to their tactical tasking, leaving no redundant or superfluous infections in their wake. Perhaps we are looking at an organization that has mastered a specialization pipeline: with the best developers doing only development, the best operators focusing on tradecraft and variegated deployment, and relevant subject matter specialists designing the campaign tasking itself.

On shared programs

A final oddity of possible threat actor configurations is that of shared programs – where a malware platform or series of campaigns is externally considered a single threat actor, but in fact reflects the joint efforts of two or more governments. While this joint ownership may be difficult to discern from a snippet of data, observing the 'threat actor(s)' over a period that includes exceptional geopolitical upheavals and revelations will, in fact, show indicative fractures and responses. The following are two examples that may reflect this unusual configuration, and the importance of maintaining a dynamic threat actor profile that can account for changes in the object of study.

RedOctober to CloudAtlas: geographic fracturing

GReAT announced the discovery of Red October in January 2013 [79, 80, 81]. The threat actor had been active for approximately five years and prolific in its tooling and

[70] To add a bonus nuance, there are indications that the two teams made famous by the DNC hack may have shared an inception, with APT28 and APT29 overlapping in 2011 [77].
[71] Or perhaps 'group of threat actors'.
[72] Perhaps working for different military/defence contractors.
[73] This may suggest they're also responsible for Brown Lambert, which is primarily used to deploy Black Lambert samples.
[74] This assertion is now further substantiated by code similarity analysis, which ties together the White, Black, and Brown Lamberts [78].
[75] One such example involves the use of the most prolific frameworks from each development team jointly for the same cryptonym operation, 'COD FISH' – Green Lambert (MD5: 4083139dc182495c450e2501fc601695) and Blue Lambert (MD5: 23df2b8320cd5954aa6700819cdb0faa).

infrastructure. Its victim spread included hundreds of victims globally. Given indications that the malware developers were russophone cemented an idea that the threat actor must be of Russian provenance. The operation was quickly shut down after the announcement.

In 2014, *BlueCoat* announced the discovery of the 'Inception Framework'[76], also known as 'Cloud Atlas' [82] due to its abuse of a cloud provider for its command-and-control infrastructure. Despite entirely new tooling in a new coding style, Cloud Atlas is considered a new iteration of Red October. However, though the operational overlaps are not in dispute, it appears that Red October and Cloud Atlas are not entirely the same.

The discovery of Cloud Atlas reportedly occurred in August 2014 when its operations were almost entirely focused on Russian victims. According to updated reporting by *Symantec* [83], Russia was the primary target of the Inception Framework for three years, followed by victims in Ukraine and Moldova. In another notable shift, the attackers now preyed upon Russian business executives, a deviation from Red October's vertical targeting and also a breach of the unspoken rule that presumably keeps Russian threat actors from targeting their domestic financial sector.

While this remains interpretative at best, we should acknowledge and test an underreported hypothesis: that Cloud Atlas is an operation related to the russophone Red October, as the research suggests, but that the organization behind the operations has changed its composition. Considering that Cloud Atlas was discovered operating in the aftermath of the Russian annexation of the Ukrainian territory of Crimea (February 2014), we can further hypothesize as to the cause for both the change in targeting as well as the change in tooling and coding style. While the breakup of a joint program would presumably leave behind shared experience and institutional memory, it would also compel the need to change operating procedures and tools so as not to be exposed to the knowledge of a departed partner-cum-adversary.

Regin: the many eyes of an umbrella framework

A more fascinating case is that of Regin [84]. Believed to have been active since 2003 [85][77], Regin was used to carry out espionage operations for nearly a decade. In 2014, it entered public discourse due in large part to the revelation of its involvement in the GCHQ-led hack of a Belgian telecommunications provider. 'Operation Socialist' [87] had the express intention of gaining access to routers that handle GPRS roaming in order to man-in-the-middle mobile devices.

Regin was a distasteful topic in the information security community for a variety of reasons, including: a post-Snowden climate of distaste for the perceived overreach of Western intelligence operations, the elementary techniques used to profile and reach the Belgacom admins[78], and the suspected lag between the discovery of the malware and its public reporting. The first named signatures for Regin were pushed on 9 March 2011, and *Microsoft* added it to its *Threat Encyclopedia* in April of the same year[79].

Extensive research went into reverse engineering Regin's multi-stage approach and understanding the extent of the operations. However, the climate and easy availability of classified documents served to oversimplify our understanding of Regin as one of the more unique objects of study in threat intelligence research. Rather than being a single threat actor, or as some claimed 'a tool of the US and British intelligence agencies' [88], it is in fact an agglomeration of implementations of a unified computer-network exploitation (CNE) platform for all of Five Eyes. This led to a series of misinterpretations and, despite extensive research, perhaps a series of misunderstandings as well.

The public's combined interpretation of external research alongside the revelations of leaked documents still did not yield an appropriate understanding of the dynamics at play with the Regin platform[80] and its layered intricacies:

For one, to what extent can we determine whose targeting was involved with specific victims? Given references to a process of deconfliction, there's the distinct possibility that this would not be possible from a purely external perspective[81]. Next, to what extent is 'Regin' a single unified malware family vs. a series of implementations of the same specification blended together by shared libraries for cross-compatibility and portability?

What was 'sig-ed' as Regin specifically may well be some combination of WARRIORPRIDE (CSEC, DSD[82]) and DAREDEVIL (GCHQ). Two other cryptonyms were haphazardly lumped into this cluster: UNITEDRAKE and STRAITBIZARRE. This was based on a misconception that these last two were NSA-equivalent names for the same suite or elements within it. However, that is not the case:

First, UNITEDRAKE is, in fact, a reference to the Equation toolset, with the different versions UnitedRake v3.x and v4.x corresponding to EquationDrug and GrayFish, respectively. It's important to remember that the Equation toolkit is extensive

[76] Sadly, the original *BlueCoat* report requires registration. *Symantec*'s updated research is available [6].
[77] A single sample may point to Regin being active as early as 1999 – 'APT Paleontology in the age of cyber', Costin Raiu, 5m40 [86].
[78] See: *LinkedIn*.
[79] As reported by *The Intercept.*
[80] A description that should be considered more in line with how the Tilded platform was described (as the framework involved in both Duqu and its contributions to Stuxnet).
[81] Without the (unsolicited) aid of classified documents.
[82] Now simply 'ASD', Australian Signals Directorate.

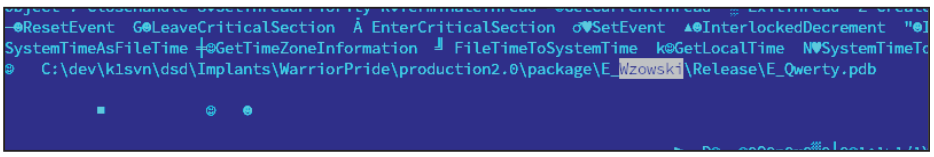

Figure 1: Unstripped QWERTY [91] sample (MD5: 40451f20371329b992fb1b85c754d062).

(having operated in some iteration or another for nearly two decades) and well established as its own operational cluster.

Secondly, it appears that STRAIGHTBIZARRE refers to a portable development library[83]. This library has also been referred to as the Wzowski API (under the apparently interchangeable references of 'CNELib v2.10' or 'WzowskiLib' [90].

A further observation may support the hypothesis that we don't yet have a fine-grained ability to accurately tell these platform-compliant families apart. It appears that, thanks to yet another unfortunate series of leaks, we may be in possession of some unadulterated version of the 'CNELib' referenced above. The ShadowBrokers 'Lost in Translation' leak [92] included x86 and x64 versions of a file named 'cnli-1.dll'[84]. With the advent of code similarity hunting at scale, we have come to find that code from this cnli-1.dll is baked into many Equation group samples (as we might expect) but also compiled into Regin samples[85]. Moreover, sigs based on versions of different bitness (32- and 64-bit) of the cnli-1.dll library match respective bitness versions of Regin.

There's a dual purpose in the adoption of a set of standardized libraries across different compatible malware frameworks. First of all, it allows different development teams to focus on the creation of specific plug-ins to suit specific needs and for the resulting plug-ins to be compatible and shareable across this threat actor alliance 'out-of-the-box'. Secondly, given that the library in question is essentially a wrapper for many of the mundane interactions with an operating system, it's likely we are looking at the component that enables ease of portability by allowing the functional calls to remain the same and letting a library implement the operating system specifics.

Ultimately, apart from a fascinating case study in the dynamism and complexity of threat actors in the wild, a retrospective study of Regin should be a humbling experience for threat intel researchers. It illustrates that, though our craft is improving, our cutting-edge developments are only just beginning to reach a level of granularity to understand some of these complex developmental interplays between organizations[86]. It once again highlights the importance of staving off definitive attribution in order to keep our hypotheses open and testable and our actor profiles responsive to the dynamic nature of the object of study.

While the United Kingdom has gracefully borne the brunt of attribution for all Regin attacks, it should be clear by now that they were not behind all of them. The NSA's TAO has similarly been dragged into this particular attribution battle, though it's uncertain whether they've ever employed that particular toolkit for their own operations[87]. While spectators that disapprove of intelligence operations in general may find this a pedantic distinction, it's an important distinction to make for the sake of respecting our craft as researchers and making sure we are availing ourselves of the fullness of our analytical capabilities without getting lost in indignant affectation.

[83] As in the reference: 'DROPOUTJEEP is a STRAITBIZARRE-based software implant' [89].

[84] cnli-1.dll x32 (MD5:a539d27f33ef16e52430d3d2e92e9d5c), cnli-1.dll x64 (MD5:07cc65907642abdc8972e62c1467e83b)

[85] Code overlap between cnli-1.dll and Regin – Attribution 2.0, Costin Raiu, AREA41 Conference, [93] (25m), [94] (Slide 35).

[86] Without being handed a crib sheet.

[87] And it isn't the only time this has happened. The earliest Project Sauron (Remsec/Strider) samples appear in June 2011, three months after the first Regin signatures effectively signal public awareness of the framework. Tweetable attribution claims appear to be similarly misguided in that case as well.

CONCLUSION

Threat intelligence in the private sector reached its current practices in an organic manner responsive to a combination of market forces, past expertise, ingenuity, and discovery. The research byproduct casually catalogues a fascinating global incursion into the fifth domain, for espionage and presumably for warfare. Though our methods have provided amazing research so far, there's room for improvement. As we push back against the more frivolous incentives malforming our research products and misguiding onlookers, we'd do well to further structure our research methods. A humble, self-aware, and concerted study of our failures and successes will help threat intelligence advance to a next greater stage of utility in the ongoing struggle to disarm the offence potential of the fifth domain.

REFERENCES

[1] Wittgenstein, L. Philosophical Investigations – Philosophy of Psychology: A Fragment xi, §160.

[2] Joyce, R. USENIX Enigma 2016 – NSA TAO Chief on Disrupting Nation State Hackers. https://www.youtube.com/watch?v=bDJb8WOJYdA, 22:15.

[3] Mueller, P.; Yadegari, B. (2012), p. 10, §2.6. https://www2.cs.arizona.edu/~collberg/Teaching/466-566/2012/Resources/presentations/topic9-final/report.pdf.

[4] https://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf.

[5] https://media.kasperskycontenthub.com/wp-content/uploads/sites/43/2017/10/20114955/Bartholomew-GuerreroSaade-VB2016.pdf.

[6] https://www.symantec.com/connect/blogs/blue-coat-exposes-inception-framework-very-sophisticated-layered-malware-attack-targeted-milit.

[7] https://securelist.com/the-devils-in-the-rich-header/84348/.

[8] https://blog.talosintelligence.com/2018/02/olympic-destroyer.html.

[9] https://securelist.com/lazarus-under-the-hood/77908/.

[10] https://media.kasperskycontenthub.com/wp-content/uploads/sites/43/2018/03/07190154/The-ProjectSauron-APT_research_KL.pdf.

[11] https://www.symantec.com/connect/blogs/strider-cyberespionage-group-turns-eye-sauron-targets.

[12] https://wikileaks.org/ciav7p1/cms/page_14588467.html.

[13] https://www.google.com/url?q=https://ccdcoe.org/sites/default/files/multimedia/pdf/Art%252005%2520Targeting%2520Technology.%2520Mapping%2520Military2520Offensive%2520Network%2520Operations.pdf&sa=D&ust=1533429954860000&usg=AFQjCNF48zrrgOvRnH3tTtHZU0TW2yc_QA.

[14] https://archive.org/details/nsa-treasure-map-new-der-spiegel-14-0914.

[15] https://archive.org/details/nsa-treasure-map-der-spiegel-14-0914.

[16] https://blog.talosintelligence.com/2017/05/wannacry.html.

[17] https://securelist.com/schroedingers-petya/78870/.

[18] https://securelist.com/bad-rabbit-ransomware/82851/.

[19] https://media.kaspersky.com/pdf/Guerrero-Saade-VB2015.pdf.

[20] https://securelist.com/winnti-more-than-just-a-game/37029/.

[21] https://www.wired.com/2015/06/foxconn-hack-kaspersky-duqu-2/.

[22] https://www.welivesecurity.com/2010/07/22/why-steal-digital-certificates/.

[23] https://www.symantec.com/connect/blogs/suckfly-revealing-secret-life-your-code-signing-certificates.

[24] https://www.comodo.com/news/press_releases/2017/07/Comodo-replace-symantec-certificates-facing-an-uncertain-future-in-chrome.html.

[25] https://media.kaspersky.com/en/Duqu-2-0-Frequently-Asked-Questions.pdf.

[26] https://trailofbits.files.wordpress.com/2012/06/flame-md5.pdf.

[27] https://blogs.technet.microsoft.com/srd/2012/06/06/flame-malware-collision-attack-explained/.

[28] https://securelist.com/gadget-in-the-middle-flame-malware-spreading-vector-identified-22/33081/.

[29] Gregory, A.; Rainbow, R. What Behavioural Investigative Advisers actually do. Professionalizing Offender Profiling (2011), pp.28-32.

[30] Moran, N.; Koel, B. The Italian Connection: An analysis of exploit supply chains and digital quartermasters. (2015). https://drive.google.com/file/d/0Bw35r_AUUldgMEZUdXUyWUR0T3M/view.

[31] https://info.lookout.com/rs/051-ESQ-475/images/Lookout_Dark-Caracal_srr_20180118_us_v.1.0.pdf.

[32] (Classification Warning) https://search.edwardsnowden.com/docs/HackersareHumanstooCyberleadstoCIleads2017-08-02_nsadocs_snowden_doc.

[33] Stoll, C. The Cuckoo's Egg: Tracking a Spy Through the Maze of Computer Espionage.Pocket Books (1990).

[34] https://www.nytimes.com/2017/03/12/world/europe/russia-hacker-evgeniy-bogachev.html.

[35] https://securelist.com/wild-neutron-economic-espionage-threat-actor-returns-with-new-tricks/71275/.

[36] https://citizenlab.ca/2015/12/packrat-report/.

[37] https://securelist.com/equation-group-from-houston-with-love/68877/.

[38] (Classification warning) https://edwardsnowden.com/docs/doc/image-583936-galleryV9-hpue.jpg.

[39] https://securelist.com/equation-the-death-star-of-malware-galaxy/68750/.

[40] (Classification warning) https://edwardsnowden.com/docs/doc/image-583969-galleryV9-ygjs.jpg.

[41] https://mp.weixin.qq.com/s?__biz=MjM5MTA3Nzk4MQ==&mid=2650170101&idx=1&sn=714546c757db8291d52b6a4332c364a4&chksm=beb9c1c789ce48d1d6a96ed51b2506feab47c344b50461e7a4f72b19d89879dc113669d47a33.

[42] (Classification warning) https://wikileaks.org/ciav7p1/cms/page_14588809.html.

[43] https://securelist.com/the-equation-giveaway/75812/.

[44] https://www.lemonde.fr/international/article/2014/03/21/la-france-suspectee-de-cyberattaque_4387232_3210.html.

[45] http://www-test.cyphort.com/evilbunny-malware-instrumented-lua/.

[46] https://securelist.com/animals-in-the-apt-farm/69114/.

[47] https://www.welivesecurity.com/2015/03/05/casper-malware-babar-bunny-another-espionage-cartoon/.

[48] https://www.gdatasoftware.com/blog/2015/02/24270-babar-espionage-software-finally-found-and-put-under-the-microscope.

[49] https://securelist.com/the-caretomask-apt-frequently-asked-questions/58254/.

[50] https://securelist.com/stuxnetduqu-the-evolution-of-drivers/36462/.

[51] https://theintercept.com/2014/11/12/stuxnet/.

[52] https://www.crysys.hu/publications/files/BencsathPBF12eurosec.pdf.

[53] https://securelist.com/the-mystery-of-duqu-part-one-5/31177/.

[54] https://securelist.com/the-mystery-of-duqu-part-two-23/31445/.

[55] https://securelist.com/the-mystery-of-duqu-part-three-9/31486.

[56] https://securelist.com/the-duqu-saga-continues-enter-mr-b-jason-and-tvs-dexter-22/31442/.

[57] https://securelist.com/the-mystery-of-duqu-part-five-6/31208.

[58] https://securelist.com/the-mystery-of-duqu-part-six-the-command-and-control-servers-36/31863/.

[59] https://securelist.com/the-mystery-of-the-duqu-framework-6/32086/.

[60] https://securelist.com/the-mystery-of-duqu-framework-solved-7/32354/.

[61] https://securelist.com/the-mystery-of-duqu-part-ten-18/32668/.

[62] https://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_duqu_the_precursor_to_the_next_stuxnet.pdf.

[63] https://media.kasperskycontenthub.com/wp-content/uploads/sites/43/2018/03/07205202/The_Mystery_of_Duqu_2_0_a_sophisticated_cyberespionage_actor_returns.pdf.

[64] https://www.crysys.hu/publications/files/duqu2.pdf.

[65] https://www.albany.edu/iasymposium/proceedings/2014/16-KaramanianSample%23.pdf.

[66] https://securelist.com/the-flame-questions-and-answers-51/34344/.

[67] https://www.symantec.com/connect/blogs/flamer-highly-sophisticated-and-discreet-threat-targets-middle-east.

[68] https://crysys.hu/publications/files/skywiper.pdf.

[69] https://www.rsaconference.com/writable/presentations/file_upload/br-208_bencsath.pdf.

[70] https://securelist.com/back-to-stuxnet-the-missing-link-64/33174/.

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

[72] http://virus.wikidot.com/29a.

[73] https://www.f-secure.com/documents/996508/1030745/CozyDuke.

[74] https://www.blackhat.com/docs/us-15/materials/us-15-Pitts-Repurposing-OnionDuke-A-Single-Case-Study-Around-Reusing-Nation-State-Malware-wp.pdf.

[75] https://www.symantec.com/connect/blogs/forkmeiamfamous-seaduke-latest-weapon-duke-armory.

[76] https://www2.fireeye.com/rs/848-DID-242/images/rpt-apt29-hammertoss.pdf.

[77] https://securelist.com/a-slice-of-2017-sofacy-activity/83930/.

[78] https://www.youtube.com/watch?v=jeLd-gw2bWo 29m30.

[79] https://securelist.com/the-red-october-campaign/57647/.

[80] https://securelist.com/red-october-diplomatic-cyber-attacks-investigation/36740/#8.

[81] https://securelist.com/red-october-part-two-the-modules/57645/.

[82] https://securelist.com/cloud-atlas-redoctober-apt-is-back-in-style/68083.

[83] https://www.symantec.com/blogs/threat-intelligence/inception-framework-hiding-behind-proxies.

[84] https://www.symantec.com/content/dam/symantec/docs/security-center/white-papers/regin-top-tier-espionage-tool-15-en.pdf.

[85] https://media.kasperskycontenthub.com/wp-content/uploads/sites/43/2018/03/08070305/Kaspersky_Lab_whitepaper_Regin_platform_eng.pdf.

[86] https://www.youtube.com/watch?v=UXgFqaqOzgc.

[87] Mobile Networks in MyNOC World. GCHQ, p.9 and p.20. (Classification warning) https://edwardsnowden.com/docs/doc/gchq-mobile-networks-in-my-noc-world.pdf.

[88] http://www.spiegel.de/netzwelt/netzpolitik/trojaner-regin-ist-ein-werkzeug-von-nsa-und-gchq-a-1004950.html.

[89] NSA ANT Catalog. p.29 (Classification warning). https://www.eff.org/files/2014/01/06/20131230-appelbaum-nsa_ant_catalog.pdf.

[90] https://medium.com/@botherder/everything-we-know-of-nsa-and-five-eyes-malware-e8eac172d3b5.

[91] https://securelist.com/comparing-the-regin-module-50251-and-the-qwerty-keylogger/68525/.

[92] https://steemit.com/shadowbrokers/@theshadowbrokers/lost-in-translation.

[93] https://www.youtube.com/watch?v=jeLd-gw2bWo.

[94] http://area41.io/slides/2018/AREA41_18_keynote_costin_raiu.pdf.