APT REPORTS AND OPSEC EVOLUTION, OR: THESE ARE NOT THE APT REPORTS YOU ARE LOOKING FOR

Gadi Evron
Cymmetria, Israel

Inbar Raz
Perimeter X, Israel

Email gadi@cymmetria.com; inbar@perimeterx.com

ABSTRACT

With the advancement of defensive cybersecurity practices and the regular release of reports exposing toolsets used in APT attacks, advanced threat actors have had to adapt. However, while APT reports should have threat actors scrambling to keep up, in reality they are providing APT actors with the information they need to implement new operational security practices and technologies that have defenders working as hard as ever to protect their networks. Not only are attackers adapting; they are evolving at a faster rate than defenders. So what are we, as defenders, doing wrong?

The fact is, many public APT reports suck. Even though they tend to be long and technical, they are often not full reports, but rather a commentary on the attack platform(s) and deployment technique(s) used, intended for PR purposes. This results in an asymmetry – an information gap – that benefits the attacker. Current APT reports basically act as free Q&A for APT actors, providing them with valuable information about defenders' insights into their tools and actions. As a result, APT actors are able to adapt their own OPSEC practices and technologies in order to stay one step ahead of defenders. APT reports in their current state are more beneficial to attackers than defenders.

Currently, most APT reports provide abundant information on indicators of compromise (IOC), C&C set-up and malware used. This presentation examines actual techniques that can be used to re-engineer the entire attack process, including how attackers decide what information is valuable to target, where that information can be found, how they create a target report and then attack plan, and the ongoing concerns of an attacker during lateral movement (i.e. OPSEC, intelligence gathering, and keeping their identity hidden). Based on these techniques, we discuss specific defensive countermeasures that can be used. If APT reports included more actionable intelligence that defenders could use to create better defence practices, their value would then be greater to defenders than to attackers. We discuss how intelligence on the attack vector of an APT, or on what information was compromised, is actually more valuable to a defender than the information that currently dominates APT reports (malware analysis, IOCs). APT reports with more actionable intelligence would afford us the ability to publicly re-engineer specific attacks, consequently rendering useless certain attack techniques that are currently not available for public knowledge.

The cybersecurity sector needs to demand earlier reporting of breaches (or at least a heads up to the security community), actionable public information sharing, and a move away from our current fixation on attribution. We need to make hackers spend significantly more time, effort, and resources in order to succeed. By producing better APT reports, not only can the security community increase attackers' costs and cause them to constantly be on guard, but also significantly disrupt the attackers' operations and make it difficult for them to rebuild their attack infrastructure after being compromised and exposed.

The bottom line is: in order to counter the evolution of APTs, we need APT reports that provide a more wholesome view of an attacker's motivations and chosen vector in addition to an analysis of his techniques. This shift in focus can give security professionals more tools to successfully re-engineer an attacker's methodology.

> *There is an information gap in that APT reports are more helpful to the attacker than to the defender – there needs to be a shift.*

Cybersecurity is a constantly evolving field, with protection technologies shifting to keep up with the transformation of cyber threats that are growing in complexity. Operational security, otherwise known as OPSEC, is a collection of processes both for developing software with risk management in mind, as well as carrying out a mission using said software. In the case of a network breach, OPSEC also includes the measures used by the attacker to avoid detection, and ensure success. Contrary to what some might believe, operational security is useful for both the attacker and the defender. When considering operational security from the perspective of an attacker, OPSEC can be utilized in order to achieve the attacker's goals while preventing detection. There are instances, however, when OPSEC is a hindrance. In terms of the time needed to achieve success, it often necessitates slow movement on the attacker's part in order not to be compromised. OPSEC's scalability can also be problematic, as can the ease of deployment.

Over the past 10 years, a specific type of threat known as an 'advanced persistent threat', or APT, has become increasingly common. As a result, defence technology has been transforming in order to meet these kinds of threats. Some of the most notable breaches and hacks in the past few years have been a result of APT campaigns, such as Stuxnet and Flame. These two campaigns actually varied in size and specification. Stuxnet was a computer worm, believed to be a joint American-Israeli project, that compromised Iranian centrifuges. In this instance, the worm was smaller in size but designed to target specific vendors. Flame, on the other hand, was much larger – 20 megabytes in comparison to Stuxnet's 500 kilobytes – but much less target-specific. Its main purpose was to infect a system and look for intelligence, by taking screenshots and recording audio and then sending that information back through an encrypted channel.

When APTs are identified and blocked, it is often the result of a hacker being compromised by faulty operational security. On the attacker's side, this requires a transformation in product. For instance, after Stuxnet and Flame were identified, new advanced persistent threats became stealthier and more complex. One example is Gauss, which had tight execution constraints which only allowed the malware to execute on certain targets, resulting in it evading capture for significantly longer. Even now, there is still a piece of the malware, detailing which computer it was meant for, that remains encrypted. Another APT that emerged as the actors learned from their mistakes was APT3, a piece of malware that, when first active, used only small and disposable tools before bringing out the more complex and intense tools, in order to prevent their loss in the event of it being caught earlier on.

On the defender's side, the first significant sign of hope in defeating APTs came in early 2013 when the first major APT report was released by *Mandiant* [1]. Its disclosure of the advanced persistent threat campaign named 'APT1' made a huge impact thanks to the extent of the exposure. The report showed that the campaign had compromised 141 companies across 20 countries. The attack methodology was to establish access to a network and visit the network every so often to steal new data. With a sizable staff, APT1 was able to steal incredible amounts of data, remaining in one particular system for as long as four years and 10 months, and stealing a whopping 6.5 Terabytes of data from just one organization over the course of 10 months. *Mandiant*'s report outed APT1 publicly for the first time and, since then, APT reports have become a common outcome after an attack, produced with the intent of aiding other defenders by detailing the attack methods and preventative measures. Unfortunately, many of these APT reports are failing to achieve just that. The major problem with these reports is not ill intent, but a gap in information which causes the details to favour the attacker rather than the defender.

This disparity occurs for a number of reasons, but all with the same consequence. Most APT reports are lengthy, causing readers to have to sift through a great amount of detail in order to get to the most important piece information. Sometimes, that information is contained within the report, but sometimes it is not and the report is not disclosed in full. In these cases, the part of the report that is publicized is done so for PR purposes and the full report is available only to paying customers. This is a clear cut case of when the attacker is disproportionately benefited because the part of the report that *is* available often only contains deployment techniques. Deployment techniques, for other attackers reading the report, represent nothing more than an inside look at 'what the other guy can do', with less relevance to any other defenders.

An in-depth analysis of the information commonly contained in APT reports produced the observation shown in Figure 1.

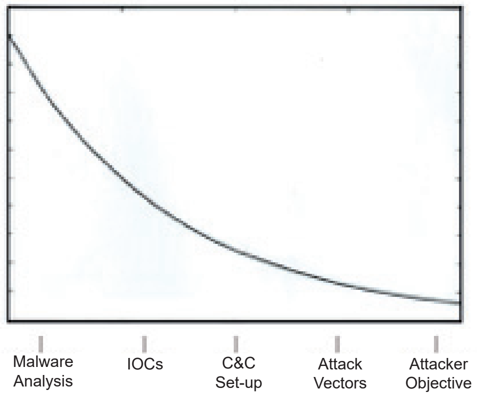

Figure 1: The relative amounts of different types of information commonly contained in APT reports (y axis: information).

Most of the information is malware analysis, which makes for interesting reading for fellow malware researchers but has very little value to an IT security person trying to defend a network. After that come indicators of compromise, which actually *are* useful for detecting and mitigating threats but have a very short relevance term, since the rise of OPSEC. Coming next, the C&C set-up has more relevance as its life span is usually longer (though less and less relevant to other targets or other campaigns). Attack vectors are very important because they are harder to change, are relevant to a larger group of targets, and are actionable – you have something to defend against – but these are not often included in the reports. And the most important piece of information – attacker objectives, the one thing that lets a defender focus and concentrate the defences in the right place – is rarely shared.

Moreover, preventative strategies are released, which become in essence a free lesson for the actor responsible for the advanced persistent threat, showcasing what should be done differently next time, with a detailed explanation of the defence software, essentially allowing the attack mechanisms to evolve and infiltrate the system more successfully.

For other attackers, these reports represent a free QA process, leaking both other attacking technologies as well as how they are defended against, resulting in APT evolution. There are several instances of seemingly unrelated campaigns having been carried out using similar code, demonstrating that there was some spread of information on the actors' end. In the case of APT1, different parts of the malware were registered under the same name and email. This lesson was clearly absorbed by the Russian programmers who created Turla, an APT which was far more sophisticated and used satellite connections to hijack networks and steal data, without any locations disclosed.

In order to fix this information asymmetry, one of the best solutions is to use 'reverse engineering'; this essentially means looking at an attack from the attacker's perspective in order to determine what can be done on the defender's part. This is where the *Cyber Engagement Process* comes into play. This is a simplified model that details the generalized main steps of an attacker: composing intelligence requirements; creating a target list; engagement, whether it be pre-breach, post-breach, or ongoing; and finally folding and retreat.

The first step, composing intelligence requirements, requires a deep look at the attacker's end as to what kind of information he is looking for. This step is quite problematic for defenders. In the event the defender manages to determine the objective of the attacker, some precautions can be used. For instance, in the case of Operation Aurora, security experts were able to determine that the goal of the attackers was to gain access to particular source code repositories at certain high tech and security companies. However, in most cases, there is no indication of the attackers' objective or what information is being sought, making defensive strategies much more difficult to select. Even when the objective is known, it can be difficult to understand the mindset of the attackers or why they would want certain information – as in the case of the Office of Personnel Management breach, when hackers targeted personal information such as Social Security numbers. For this step, with little knowledge of the attacker's intentions, one should perform risk assessments, to ascertain the potential impacts.

The next step in the Engagement Process is for the attacker to compile a target list. This is directly related to the intelligence requirements, as it is nothing more than a list to determine who has the relevant information as per the intelligence requirements: 'Who has the answers to my questions?' This can take shape in specific targets, or it can be more broad. From the defender's perspective, this again is difficult to act against, for numerous reasons. Pattern recognition – one of the solutions – requires more time than is often available. One measure that can be taken is to constantly be aware of other attacks that are happening, particularly in companies that are similar to one's own. In fact, there are many instances in which one company is hacked and then, a few days later, a number of other similar natured companies are breached. In the case of the *Target* breach, point-of-sale devices were hacked; any other companies with similar point-of-sale devices should have been on the lookout, making active changes to make sure they would not be hacked in the same fashion. To prevent this, threat assessments should be carried out often and breaches reported as soon as possible so that other defenders can take notice and be forewarned. However, not all actors will include this step as part of their engagement plan, as this is just a generalized cycle. One example is the Sofacy APT, which did not have a compiled target list, but seemed to latch onto whatever information it could find in order to monetize it later.

The third step of the cycle, intelligence gathering, involves many substeps. Beforehand, there is generally a pre-prepared target report outlining everything that is known about the target to be used during engagement. This helps in the creation of the attack plan, which might be tailored for the target, with target-specific tools and target-specific methods. In terms of attack prevention, publicly available sensitive data poses a huge problem. Even if the information is not public, some people may inadvertently make the information easily accessible, for example by using security questions with answers that can be found on social media accounts. When there is such a lack of security awareness, companies can become susceptible to probing, either automatic or manual. This is merely the pre-engagement stage of intelligence gathering. Before even infiltrating an actual network, the attacker can have a large amount of information at hand. While this seems disheartening, once an actor is 'engaged' and inside the network, there are actually many opportunities to intervene. The engagement process is not a one-time event; it is ongoing, which is why layered security is so important. The key in this step is to put as many obstacles as possible in the way of the attacker, so that it takes longer to reach the desired data, resulting in longer periods of exposure and a greater chance of getting caught.

For the attacker, once inside the network, the main goal is to remain undetected and move around stealthily, making its way to the desired data. This is commonly known as 'lateral movement'. In order to remain undetected within the network, more intelligence gathering needs to take place. This is from a different perspective, having already infiltrated the network, but this is where operational security needs to be revisited by the attacker. This begins with mapping out the defences of the target and continues with looking for other malware. It is also important to recognize that, in this day and age, there are many other players and many other APT actors and so there might be others lurking in the same network that need to be recognized.

With this awareness comes the possibility of having to abort while inside the network. Even if an installation is in the midst of occurring, if there are signs of detection, the safest thing to do is leave the network immediately. This can be seen in the case of the threat known as 'Hurricane Panda' where the threat actor detected the presence of *CrowdStrike* and backed away.

This example is closely related to the last step, folding and retreating, which varies from one attacker to another. Some fold after report publications: Red October was a malware program operating for five years before its discovery, upon which it dismantled. Another program that quit and ran in a similar fashion was The Mask, which left the system only four hours after a blog publication. Nevertheless, there are counter-examples that demonstrate that some malware actors do not bother to fold and retreat, such as APT12 or the Gaza Hacker team. For the defender, one of the worst consequences of an attack could be when an attacker chooses to destroy all the information he gets his hands on as a form of folding and then retreating. For this reason, sometimes it can be of more use to have back-ups of data rather than regular monitoring of the system.

With this Engagement Process in mind, it is apparent that more information needs to be communicated in APT reports – and not just more information, but the *right* kind of information. One shift that needs to occur is a reduction in focus on attribution. While previously, attackers used to try to steer clear of any possibility of attribution, that no longer seems to be the case. Instead, attackers seem to have little shame and are not afraid of attribution. In order to aid not just the actor and other similar attackers, reports should be focusing on attack vectors and an analysis of attack techniques. This kind of actionable information is the only way other defending companies can benefit from these reports. However, just as important as the actual report details is the *timing* of the reports; the public needs to know about breaches immediately after they happen in order for defenders to counter future attacks. With more prompt, actionable reports, the asymmetry of information sharing will cease and the disproportionate advantages for the attackers will be negated.

REFERENCES

[1] Mandiant. APT1: Exposing One of China's Cyber Espionage Units. February 2013. https://www.fireeye.com/content/dam/fireeye-www/services/pdfs/mandiant-apt1-report.pdf.