COLUMBIA | SIPA
School of International and Public Affairs

NAMED BUT HARDLY SHAMED

The Impact of Information Disclosures on APT Operations

Team: Matthew Armelli, Stuart Caudill, John Patrick Dees, Max Eager, Jennifer Keltz, Ian Pelekis, John Sakellariadis, Virpratap Vikram Singh, and Katherine von Ofenheim

Capstone Advisor: Neal Pollard

SIPA Capstone Project
Spring 2020

TABLE OF CONTENTS

EXECUTIVE SUMMARY

This project was completed to fulfill the capstone requirement for Columbia University's School of International and Public Affairs. It investigates the effects of information disclosures on the operations of cyber adversaries, and the implications of those observed effects for the U.S. Department of Defense's cyber strategy of persistent engagement and forward defense.

Information disclosures may lead cyber threat actors to become more sophisticated, resilient, and creative.

We examined the impact of disclosures on nine APT groups from five different contexts:

China	APT1/PLA Unit 61398 APT10/MenuPass Team
Criminal Groups	Cobalt Group
Iran	APT33/Elfin APT34/OilRig/Helix Kitten
North Korea	APT38/Lazarus Group
Russia	APT28/Fancy Bear APT29/Cozy Bear

Our research found that public disclosures generally failed to stop cyber actors' operations or cause long-term disruption, but that they do often impose at least short-term friction. We find that the disruptive effect varies significantly based on a number of factors, including the scope of the disclosure and the disclosing actor. However, disclosures may in fact also lead cyberthreats to become more resilient and creative because they need to retool, rebuild their infrastructure, or change their TTPs. The exceptions to this observation are China's APT1 and APT10, both of whom ultimately ceased operations following highly public disclosures and U.S. Department of Justice indictments.

Given these findings, disclosures are somewhat useful in achieving the objectives of persistent engagement by imposing costs and increasing the resiliency of networks. However, the level of cost imposed by disclosure events is simply not high enough to significantly change the decision calculus of most adversaries conducting cyber activity. Disclosures must be care-fully targeted and used in combination with other elements of power.

Lastly, private cybersecurity vendors hoping to use information disclosure offensively should consider the geostrategic context in which they operate, and they should target disclosures more effectively to counter individual groups.

INTRODUCTION

GENESIS OF THE PROJECT

This report was completed in order to fulfill their Capstone graduation requirement for the Master of Public Administration/Master of International Affairs program at Columbia University's School of International and Public Affairs (Columbia | SIPA).

Ken Wolf, a 2016 graduate of Columbia | SIPA and Senior Manager on Standard Chartered Bank's cyber threat intelligence team wrote the project's original Terms of Reference (TOR). In designing the TOR, he sought to combine the cyber threat intelligence needs of Standard Chartered Bank and other financial institutions with the research that scholars of cyber conflict are conducting to understand the implications of the United States Department of Defense's cyber strategy of persistent engagement and forward defense.

A team of nine postgraduates worked together throughout the Spring 2020 semester to conduct the research needed to satisfy the TOR. Eight of the nine are students of international security policy, and the remaining student international finance and economic policy. All have taken courses on cyber threat intelligence and cyber conflict. The project's faculty advisor, Adjunct Professor Neal Pollard, oversaw and guided their work, providing feedback and assistance where necessary.

RESEARCH QUESTIONS

The project's original TOR states that there have been leaks of information and tools associated with numerous advanced persistent threat (APT) groups in Iran. The TOR asks the capstone team to answer two questions in order to provide actionable information to Standard Chartered Bank and scholars of cyber conflict:

- ♦ *What is the impact of leaks and information disclosures on adversary operations?*
- ♦ *How can the answer to the first question inform the U.S. strategy of persistent engagement and forward defense?*

We found ourselves unable to comprehensively respond to the TOR without expanding its scope. In order to fully address both questions posed, we chose to include adversary groups from outside of Iran, including criminal groups and APTs from other countries, which are relevant to understanding both the threats that financial institutions face and to understanding persistent engagement and forward defense. We also considered the impact of different types of disclosures, which can contain different types of information, and which come from different sources, such as the private sector and government agencies.

We then posed two additional questions:

- *Under what conditions is a disruption likely to succeed?*
- *What evidence indicates that a disruption has successfully contributed to the strategy of persistent engagement and forward defense?*

INDUSTRY SITUATION

Standard Chartered Bank is a member of the financial services sector. This sector is an important part of both U.S. and global critical infrastructure, and Standard Chartered Bank is one of thirty global systemically important banks (G-SIB).[1] The sector is comprised of depository institutions like Standard Chartered Bank, insurers, investors, credit and financing organizations, and financial utilities providers. It is subject to an increasingly large and sophisticated number of cyberattacks.[2]

Financial services institutions face cyberthreats from both criminals and from nation-states.

Financial services institutions face cyberthreats from both criminals and nation-states. From criminals, they face traditional cybercrime carried out by criminal organizations that specialize in attacking financial institutions. Cybercriminal organizations are becoming increasingly capable. They have developed sophisticated methods to bypass anti-fraud systems and measures, and to attack ATMs. They are challenging the sector's ability to implement strong multi-factor authentication protocols, since authentication factors can be bypassed via creativity, social engineering, and vulnerabilities in biometrics algorithms. They relentlessly steal and monetize millions of consumers' credit card information.[3]

The nation-state threat to the financial services sector is multifaceted. Nation-states can target financial institutions for political reasons, carrying out disruptive attacks on institutions to cause their adversaries economic pain. Political considerations factored into Iran's Operation Ababil and North Korea's DarkSeoul, both of which were DDoS campaigns against U.S. and South Korean financial institutions respectively. States can also spy on financial institutions to gain intelligence on politically sensitive customers or the operations of the institution and its corporate clients.[4] Finally,

[1] Financial Stability Board, *2019 List of Global Systemically Important Banks (G-SIBs)* (22 November 2019): https://www.fsb.org/wp-content/uploads/P221119-1.pdf

[2] 'Financial Services Sector', *CISA*: https://www.cisa.gov/financial-services-sector

[3] Y. Namestnikov & D. Bestuhev, 'Cyberthreats to Financial Institutions 2020: Overview and Predictions' (3 December 2019), *Kaspersky*: https://securelist.com/financial-predictions-2020/95388/

[4] FireEye, *Cyber Threats to the Financial Services and Insurance Industries* (2016), 1: https://www.fireeye.com/content/dam/fireeye-www/solutions/pdfs/ib-finance.pdf

states may target financial institutions to steal money, as seen in North Korea's attacks on the SWIFT banking transfer system.[5]

METHODOLOGY

We approached the abovementioned question qualitatively. First, APT groups were selected from different geopolitical contexts: four countries and criminal enterprises. We chose notorious and otherwise prominent threat actors which have received a significant amount of public attention from cyber threat intelligence firms, and which have been the subject of numerous public disclosures. In order to control for the groups' differences in geopolitical situation, and to discern whether (or which) contextual similarities effect a specific response, we endeavored to analyze two APTs from each country. We had originally omitted APT38 (Lazarus Group) given its lack of a suitable North Korean correlative, but later expanded our selection criteria to include it at the client's request.

These disclosures force APTs to change their infrastructure, tools, and tactics following their release.

The countries and contexts covered in this report appear in alphabetical order. Within each country, APTs are listed in numerical order according to FireEye's APT-labeling system:

China	APT1/PLA Unit 61398 APT10/MenuPass Team
Criminal Organizations	Cobalt Group
Iran	APT33/Elfin APT34/OilRig/Helix Kitten
North Korea	APT38/Lazarus Group
Russia	APT28/Fancy Bear APT29/Cozy Bear

After the selection process, we mapped out public disclosures from cyber threat intelligence firms and government indictments (where applicable) according to the thresholds designated on the Pyramid of Pain. We included public releases of network and host artifacts, tools, or TTPs as a disclosure; and we added the public revelation of individual hackers and their cyber personae to both the Pyramid of Pain and our list of possible disclosure events. These events were selected because they impose significant costs on threat actor groups in terms of the changes that must be made to their infrastructure, tools, and tactics following a disclosure. Disclosure events below our threshold–domain names, IP addresses, hash values–provide actionable intelligence to network defenders, but

[5] W. Carter, 'Forces Shaping the Cyber Threat Landscape for Financial Institutions' (2 October 2017), *Swift Institute*, 6: https://csis-prod.s3.amazonaws.com/s3fs-public/171006_Cyber_Threat_Landscape%20_Carter.pdf?UWqJEbDm.dBKSLEIFTyYs1IxJaExh9Y7

fail to cause enough pain to the adversary to warrant inclusion.

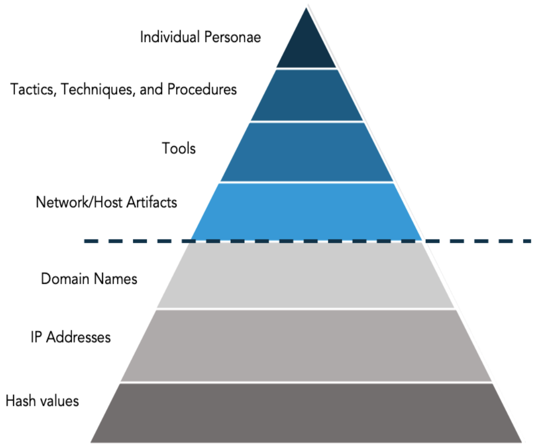

We then assessed the impact of different disclosure events according to several factors:

1 The use of public versus custom tools: The release of information on bespoke tooling should cause more pain to a threat actor than the release of publicly available tools because the development of custom software requires more time, talent, and money than the simple purchase of off-the-shelf malware. Outfits reliant on custom software would therefore, *ceteris paribus*, exhibit greater signs of disruption in the wake of a disclosure.

2 Domestic political pressure: Disclosure events which generate considerable publicity in the home-state or region of the group should cause more pain than low-publicity disclosures because public awareness can lead to demands that the group ceases its activities.

3 Government versus private sector disclosure: Government indictments should cause more pain than cybersecurity vendors' disclosure reports because the former entails internationally enforceable consequences. Individual private sector firms cannot, perforce, oblige others to follow rules; governments can.

4 Reach of disclosing state's law enforcement: Disclosures from a state that has an extradition treaty with the threat actor's own home state should cause more pain than the inverse because extradition agreements enhance the likelihood that a cost will be imposed on the wrongdoer. Furthermore, if the disclosing state's government maintains extradition treaties with many different countries, it should proportionally reduce the threat actor's ability to travel abroad.

5 Publicity of disclosure: The more publicity behind a disclosure, the more pain it should cause because more network defenders are likely to read it. The more network defenders to read a disclosure, the greater the chance that networks will be prepared against future actions taken by the threat actor.

6 Specificity of target: Cyberthreats that target specific entities should feel more pain from disclosures than cyberthreats that cultivate a broad range of targets because a single

target can react more quickly and comprehensively than a group of organizations or an entire sector. Further, if a group has a specific target set, it has built expertise and capabilities specific to that set (i.e. SCADA, ICS, oil and gas). When such a group is disrupted, it should have more difficulty reconstituting that capability than a group which conducts general computer network exploitation against a wider range of targets.

7 **Maturity of other cyber capabilities:** Disclosure events should cause more pain to newer and immature threat actors because, perforce, they have less architecture, fewer tools, and fewer TTPs to rely on. Mature actors might belong to larger threat actor organizations with advanced and capable arsenals, as well as numerous campaigns occurring simultaneously. Importantly, mature state actors with several APTs should be less disrupted in the long term because they likely have the organizational flexibility to incorporate a disrupted group into another. A less mature state actor without backup or incidental groups on which to fall back should experience more pain from disruptions.

We collected data in two steps. First, we conducted open-source research into disclosure events to create a timeline of each group's activity and related disclosures. This research enabled us to collect initial evidence on the impact of disclosures on each of the aforementioned APTs and to hypothesize whether or not disclosure events had a significant impact on them. We only had access to information that is publicly available online, so we had no insight into these APTs' actions in the times between one public disclosure and the next. Accordingly, we interviewed numerous experts in the private sector and government who track these groups and work in the cybersecurity field to determine whether our hypotheses were correct and to identify gaps in our information. We also used interviews with government experts and academics to explore the relationship between public disclosures and persistent engagement, in order to address the second part of our research question.

GENERAL COMMENTS ON INTERVIEWS

The team carried out interviews throughout the month of March and the first half of April. Our access to experts was slightly hindered by the COVID-19 pandemic and quarantine, and two people declined interviews, citing the need to manage the effects of the economic downturn and the security threats that arose from a wide-scale move to remote work. We did, however, speak to members of industry, government, and academia; and our interviewees included operations-oriented and policy-oriented professionals.

Although many interviewees consented to on-the-record interviews, some asked that we withhold their names from our final report so that they could speak freely.

APT1 (CHINA)

Timeline: Activity Levels and Disclosures

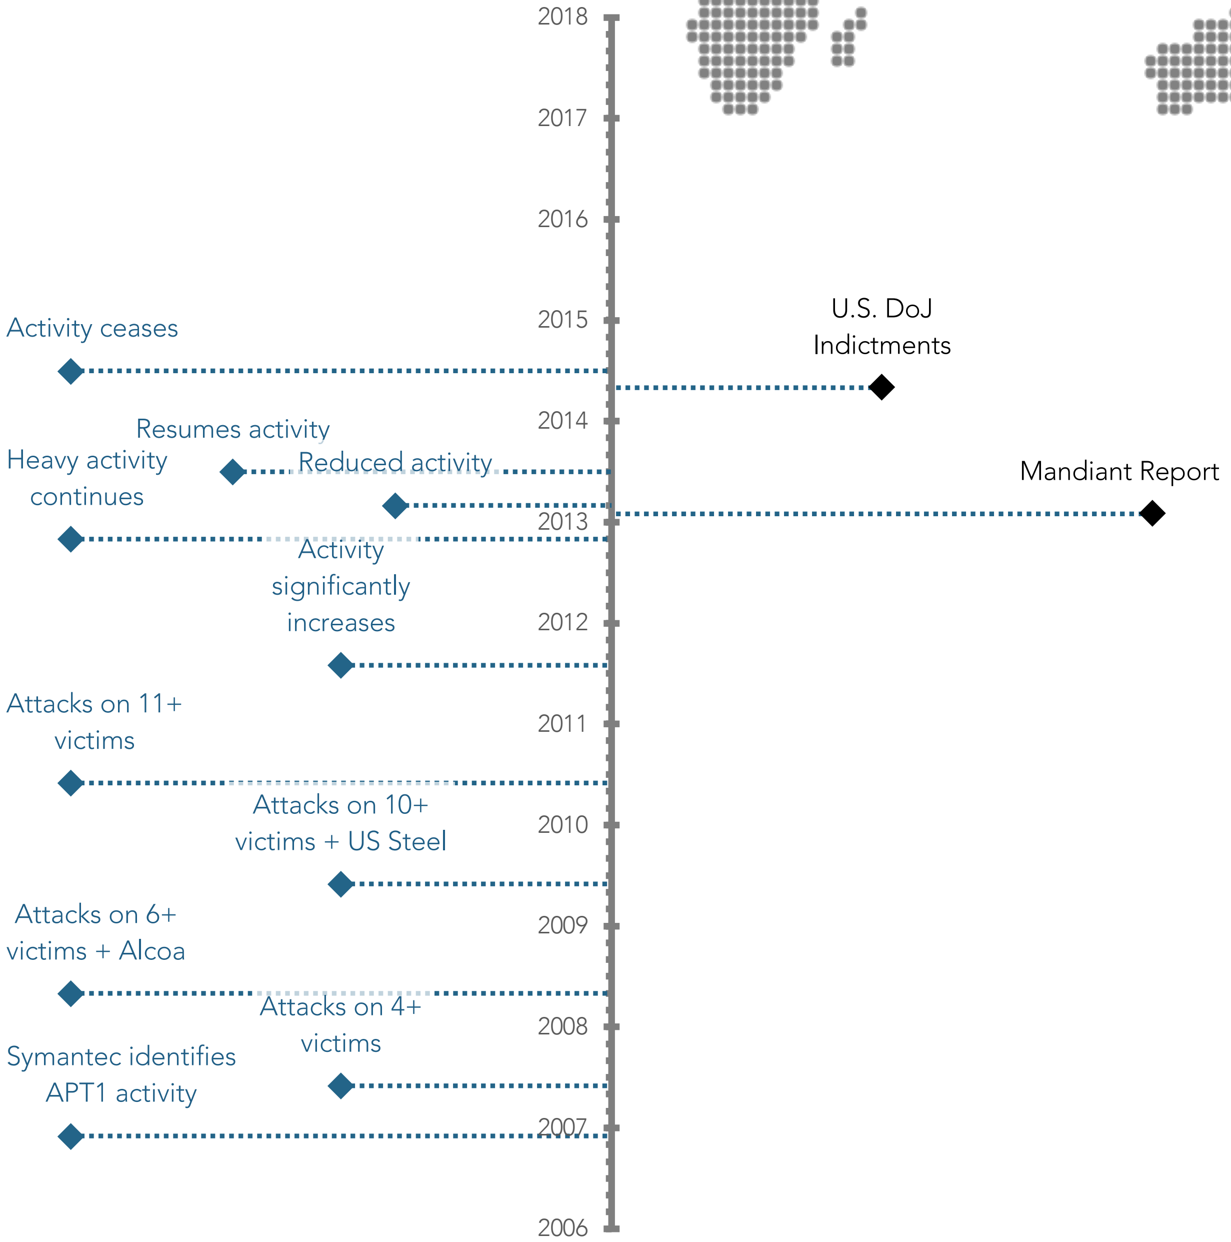

IMPACT OF DISCLOSURES

SUMMARY

APT1 (Comment Crew, Comment Group, Comment Panda) was a Chinese military cyber espionage group active from 2006 to 2014. APT1 gained widespread attention in 2013, when Mandiant released a major report that attributed the group to the People's Liberation Army and described its operations in detail.[1] Following this disclosure, APT1's operations were significantly disrupted and did not return to normal levels for 153 days. Fifteen months later, the U.S. Department of Justice indicted five members of APT1, after which all recognizable activity ceased. APT1 provides a clear case of significant and lasting operational disruption caused directly by public disclosures. Its extensive use of custom tooling, the international political environment surrounding Chinese economic espionage, and the coercive power of indictments contributed to the success of these disclosures.

THE GROUP

APT1 has been attributed to China's People's Liberation Army Unit 61398, which was assigned to the 3rd Department, 2nd Bureau of the PLA General Staff Department. The Unit was primarily responsible for targeting English-speaking countries.[2] The manpower needed to perform APT1's large operations was significant, and the group likely consisted of hundreds of operators supported by teams of linguists, developers, intelligence analysts, and operations managers.[3] Perhaps unsurprisingly, then, Mandiant described APT1 during the interval of 2006 to 2014 as 'one of the most prolific cyber espionage groups'.[4]

Mandiant found that APT1 exfiltrated data from 141 organizations in various industries. Nothing indicates that APT1 conducted offensive cyber operations: espionage appears to have been its sole responsibility, and in that work, it was notably persistent. Mandiant calculated that it spent an average of 356 days on a network, though in certain cases it exfiltrated data for years.[5]

TIMELINE

2006	Symantec first identifies APT1 activity in late 2006.[6]
2007	APT1 compromises four victims.[7]

[1] 'APT1: Exposing One of China's Cyber Espionage Units' (19 February 2013), *Mandiant*: https://www.fireeye.com/content/dam/fireeye-www/services/pdfs/mandiant-apt1-report.pdf
[2] Mandiant 2013 (n. 1), 8 f.
[3] id., 5.
[4] id., 2.
[5] id., 20 f.

[6] A. L. Johnson, 'APT1: Q&A on Attacks by the Comment Crew' (19 February 2013), *Broadcom*: https://community.broadcom.com/symantecenterprise/communities/community-home/librarydocuments/viewdocument?DocumentKey=f1265df5-6e5e-4fcc-9828-d4ddb-bafd3d7&CommunityKey=1ecf5f55-9545-44d6-b0f4-4e4a7f5f5e68&tab=librarydocuments
[7] Mandiant 2013 (n. 1), 20.

2008	APT1 compromises six victims, including Alcoa.[8]
2009	APT1 compromises ten victims, including U.S. Steel.[9]
2010	APT1 compromises over eleven victims.[10]
2011	Activity significantly increases.[11]
2012	APT1 continues heavy activity. Digital Bond reports that it was spear-phished by APT1.[12] Brian Krebs identifies APT1 as responsible for an intrusion at a Canadian electric company.[13] APT1 begins targeting SolarWorldAG.[14]
2013	**Jan:** Targeting of USW ends.[15] **Feb**: Mandiant APT1 Report released on 19 February. On 20 February APT1 changes the registration information for five malicious domains.[16] The Mandiant report disclosed four of these domains. **Mar**: Mandiant observes APT1 resume activity on 25 March.[17] **May**: Mandiant reports that APT1's activity decreased following the release of their report in February. They observed APT1 continuing intrusion activity but forced to retool and burn infrastructure that was disclosed.[18] **Jul**: Mandiant observes APT1 resume activity at normal levels on July 22.[19]
2014	**Mar:** ThreatConnect identifies APT1 activity using domains that had been disclosed in the Mandiant report and elsewhere as early as 2011.[20] **Apr**: APT1 member creates malicious domain.[21] **May**: U.S. DoJ indicts 5 members of APT1 on May 19.[22] **Aug**: Lockheed Martin reports an immediate reduction in malicious

[8] ib.; 'United States of America v. Wang Dong, Sun Kailiang, Wen Xinyu, Huang Zhenyu, and Gu Chunhui' (19 May 2014), *United States District Court, Western District of Pennsylvania*, 7: https://www.justice.gov /iso/opa/resources/5122014519132358461949 .pdf. Hereafter referenced as DoJ Indictment.

[9] Mandiant 2013 (n. 1), 20; DoJ Indictment (n. 8), 6.

[10] Mandiant 2013 (n. 1), 20.

[11] ib.

[12] id., 26.

[13] ib.

[14] DoJ Indictment (n. 8), 4.

[15] id., 26.

[16] id., 35.

[17] 'M-Trends 2014 Annual Threat Report: Beyond the Breach' (April 2014), *Mandiant*, 20: https://www.fireeye.com/current-threats/annual-threat-report/mtrends/rpt-2014-mtrends.html

[18] D. McWhorter, 'APT1 Three Months Later – Significantly Impacted, Though Active & Rebuilding' (21 May 2013), *FireEye*: https://www.fireeye.com/blog/threat-research/2013/05/apt1-months-significantly-impacted-active-rebuilding.html

[19] Mandiant 2014 (n. 17), 20.

[20] ThreatConnect Research Team, 'Old Habits Die Hard: Iterative Intelligence & Comment Crew Activity' (22 March 2014), *ThreatConnect*: https://threatconnect.com/blog/tag/apt1/

[21] DoJ Indictment (n. 8), 35.

[22] id., *passim*.

activity targeting their networks after release of Mandiant APT1 Report.[23]

TYPOLOGY OF ATTACKS

APT1 generally began its campaigns with spear-phishing emails that impersonated individuals known to the target.[24] After the victim had clicked the malicious link or attachment included in the email, their device downloaded bespoke malware that provided APT1 with backdoors and covert communication capabilities.[25] The constant evolution of this custom malware, coupled with the organized deployment of these upgrades, suggests that it had its own internal development capability.[26] In contrast, APT1 largely used public tools for privilege escalation.[27]

APT1 provides a clear case of significant and lasting operational disruption caused directly by public disclosures.

DISCLOSURE EVENTS

Two major disclosures revealed APT1's activity: namely, Mandiant's exposé of 19 February 2013 and the U.S. Department of Justice indictment of five members of APT1 from 19 May 2014.[28] Although reports on APT1's activity had previously been published, these disclosures were the first to detail TTPs and individual members and the first to attribute the activity to the Chinese government.

Mandiant's report disclosed IP addresses, domain names, malware families, TTPs, and the identities of three members. The impact was immediate and clear. Within twenty-four hours of the disclosure, a member of APT1 changed some of the domain registration information for four domains disclosed in the report.[29] Furthermore, Mandiant observed the group abandoning much of its infrastructure and beginning to retool. APT1 nonetheless resumed its intrusions on 25 March, albeit at reduced levels;[30] and its activity returned to previously seen levels by 22 July 2013.[31] Mandiant's report thus appears to have stopped APT1's intrusion activity for thirty-three days and to have had a larger disruptive effect for 153 days.

The U.S. Department of Justice indictment occurred fifteen months after Mandiant's report. It disclosed domain names, TTPs, and

[23] M. Lennon, 'Lockheed: Attackers Went Quiet After APT1 Report Exposed Chinese Hackers' (14 August 2014), *Security Week*: https://www.securityweek.com/lockheed-attackers-went-quiet-after-apt1-report-exposed-chinese-hackers

[24] Mandiant 2014 (n. 17), 28.

[25] id., 30.

[26] id., 27.

[27] e.g. Mimikatz and gsecdump: id., 34.

[28] DoJ Indictment (n. 8).

[29] id., 35.

[30] 'M-Trends 2014 Annual Threat Report: Beyond the Breach' (April 2014), *Mandiant*: https://www.fireeye.com/current-threats/annual-threat-report/mtrends/rpt-2014-mtrends.html, 20.

[31] Mandiant 2014 (n. 17), 20.

the identities of five members of APT1. All activity attributable to APT1 ceased thereafter.

Several plausible reasons offer to account for the efficacy of these disclosures. APT1's reliance on custom malware is the foremost factor. Capability development is a time-intensive and technically difficult task; and Mandiant reported that APT1 had utilized over forty malware families.[32] The process of retooling, as APT1 was observed doing, therefore reduced APT1's operational capacity. Additionally, Mandiant's report determined that a certain Mei Qiang belonged to 'a smaller group of highly capable developers within APT1'.[33] That observation is instructive.

Even if Qiang continued to work for APT1 post disclosure, new operational security efforts would likely disrupt the development workflow. Such a state of changing circumstances and protocols may have informed the 153-day disruption period after Mandiant's report. It fails, however, to explain APT1's disappearance after the Department of Justice indictment.

A second factor relates to the Mandiant report *qua* public disclosure. It was the first major revelation of its kind; and neither the international reaction nor the impact on APT1's operations were immediately clear. The report also represented a reckoning between two policies that had existed uncomfortably in parallel. On the one hand, the U.S. government had been reluctant to attribute cyber operations in public—far more than today. The Chinese government, on the other hand, had repeatedly insisted that it did not conduct cyber espionage for economic purposes. Mandiant's report revealed the mendacity of any such claims, and it did so definitively.

A third factor inheres within a qualitative difference between the report and the indictment. The Mandiant report was merely informative: it did not seek *per se* to effect a change. This does not, however, suggest that the authors of the report did not expect, or even desire, APT1 to change its behavior based on what was revealed – they most certainly did. Nevertheless, unlike the indictment, the report did not instantiate an act. On the other hand, the indictment constitutes an action against the five named PLA officers. That had real-world consequences for the ability of these individuals to travel. Furthermore, the indictment clearly communicates the intent (both present and future) of the U.S., insofar as the disclosure represented the U.S. government's willingness to confront China over its duplicity in cyber matters. CrowdStrike's Dmitri Alperovitch has noted the significant impact of indictments on PLA cyber operations, specifically citing the disruptions of APT1, APT3, and APT10.[34]

[32] Mandiant 2013 (n. 1), 5.

[33] id., 58.

[34] D. Alperovitch, 'Global Threat Brief' (26 February 2020), *RSA Conference*: https://www.rsaconference.com/usa/us-2020/agenda/hacking-exposed-global-threat-brief

As mentioned above, APT1's activity ceased after the indictment. Yet it seems highly unlikely that its talent and resources were simply disbanded. The group no longer exists in its formerly recognizable form, to be sure; however, a likelier explanation for its disappearance is that its personnel, capabilities, and targets were reassigned to another Chinese government cyber organization. Two factors make this plausible. First, the Chinese have several mature computer network operations units. This preexisting infrastructure facilitates the reconstitution of one APT into another to obfuscate its activities, while concurrently minimizing the disruption to the state's overall capacity to conduct cyber operations. Conversely, if a state lacks a large, mature cyber capability, its ability to disband and reassign an APT would be limited. Second, APT1 targeted a wide range of industries. The variety of its victims eases the process of transferring any sector-specific responsibilities to another group without attracting undue attention. Were APT1 highly specialized—only targeting, for example, SCADA systems—it would be more difficult to transfer their capability unawares.

Public disclosures significantly impacted APT1's operations and directly contributed to their disappearance.

Although it appears likely that the Mandiant report and the U.S. indictment significantly disrupted APT1, if not directly precipitating their disappearance, alternate explanations exist. In a 2016 report, FireEye noted that overall Chinese cyber activity had declined since mid-2014. They furnished two explanations: internal factors, such as President Xi's efforts to centralize and consolidate military cyber operations in the Strategic Support Force; and external factors, such as public disclosures and subsequent U.S. responses.[35] Further, Elsa Kania and John Costello note that PLA cyber units have focused more on military objectives since the Strategic Support Force was created. Political and economic targets, in turn, have shifted to the Ministry of State Security.[36] Still another internal factor is President Xi's campaign against corruption, which for cyber units often involved moonlighting to conduct financially motivated theft.[37] The 2015 agreement between Presidents Obama and Xi not to conduct espionage for commercial gain may also have contributed

[35] iSight Intelligence, 'Red Line Drawn: China Recalculates Its Use of Cyber Espionage' (21 June 2016), *FireEye*: https://www.fireeye.com/blog/threat-research/2016/06/red-line-drawn-china-espionage.html

[36] E. Kania & J. Costello, 'The Strategic Support Force and the Future of Chinese Information Operations', *The Cyber Defense Review* 3.1 (2018), 106 f.

[37] M. Hvistendahl, 'The Decline in Chinese Cyberattacks: The Story Behind the Numbers' (25 October 2016), *MIT Technology Review*: https://www.technologyreview.com/2016/10/25/156465/the-decline-in-chinese-cyberattacks-the-story-behind-the-numbers/

to the decline.[38] Nevertheless, whatever effect these other factors may have had on Chinese cyber operations (and they almost certainly did), the timeline of disruption following the two abovementioned disclosure events reveals that public disclosures significantly impacted APT1's operations and directly contributed to their disappearance.

[38] iSight Intelligence 2016 (n. 35).

APT10 (CHINA)

Timeline: Activity Levels and Disclosures

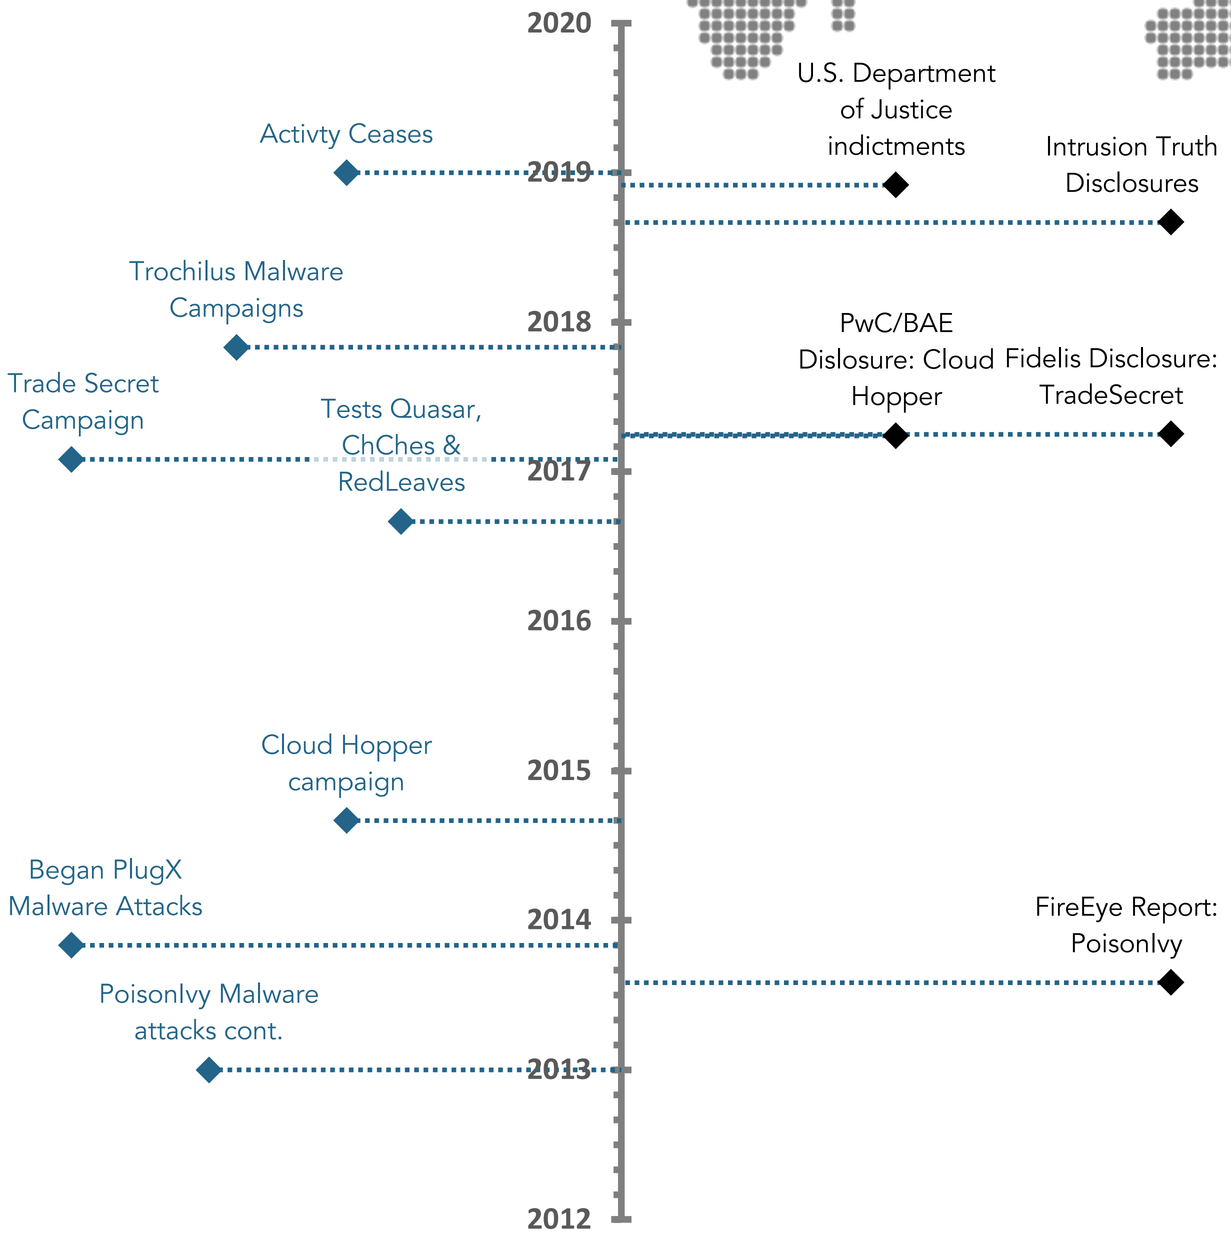

2020

Activity Ceases

U.S. Department of Justice indictments

2019

Intrusion Truth Disclosures

Trochilus Malware Campaigns

2018

PwC/BAE Dislosure: Cloud Hopper

Fidelis Disclosure: TradeSecret

Trade Secret Campaign

Tests Quasar, ChChes & RedLeaves

2017

2016

Cloud Hopper campaign

2015

Began PlugX Malware Attacks

2014

FireEye Report: PoisonIvy

PoisonIvy Malware attacks cont.

2013

2012

IMPACT OF DISCLOSURES

INTRODUCTION

APT10 is a Chinese threat actor believed to have ties to the Chinese Ministry of State Security (MSS). This highly prolific threat group has conducted large espionage campaigns against both intellectual property and Western defense information. APT10 has been the subject of several high-profile disclosures, including a 2017 report detailing the extent and nature of its campaign against managed service providers (known as Cloud Hopper) and an associated U.S. Department of Justice indictment in 2018 against several Chinese nationals believed to be associated with the group. Although most public reporting on APT10's activities has done little to impact its operations, the publication of the Cloud Hopper report disrupted APT10 activity for approximately seven weeks and the associated U.S. Department of Justice indictment appears to have halted APT10's campaigns completely.

THE GROUP

APT10 (MenuPass, Cloud Hopper, Red Apollo, CNVX, Stone Panda) is a Chinese threat actor known for its expansive espionage campaigns against sources of U.S. and Japanese defense and government information as well as sensitive trade secrets in an array of sectors. Several reports from the online group Intrusion Truth have tied APT10 activity to individuals associated with the Chinese MSS. Although it is believed to have begun operations around 2006, APT10 did not receive widespread attention until a 2013 FireEye report that associated it with PoisonIvy, a remote access tool (RAT) which had been used against U.S. and overseas defense contractors since 2009.[1]

APT10 is best known for an espionage campaign against managed services providers (MSP). This campaign was disclosed in a 2017 collaborative report from PwC UK and BAE Systems, *Operation Cloud Hopper.* [2] The campaign targeted MSPs and their clients from around the world—Canada, France, South Africa, Australia, Japan, India, Norway, the United States. [3] By targeting MSPs, APT10 could pivot to its desired target, a client network, more easily and more covertly than by attacking them directly. The campaign also marked a shift in APT10's targeting. Whereas previous efforts focused on Western defense contractors, the Cloud Hopper campaign saw the expansion of APT10's targets to include firms in the

[1] 'Poison Ivy: Assessing Damage and Extracting Intelligence' (August 2013), *FireEye*: https://www.fireeye.com/content/dam/fireeye-www/global/en/current-threats/pdfs/rpt-poison-ivy.pdf

[2] 'Operation Cloud Hopper' (April 2017), *PwC*: https://www.pwc.co.uk/cyber-security/pdf/cloud-hopper-report-final-v4.pdf

[3] 'APT10: A Brief Look into the Chinese Hacker Group's Targets and Operations' (24 November 2018), *Cyware*: https://cyware.com/news/apt10-a-brief-look-into-the-chinese-hacker-groups-targets-and-operations-e95ac4c3

finance, biotech, electronics, and telecommunication sectors.[4]

TIMELINE

APT10's espionage campaigns fall into two general periods. The first phase lasted from at least 2006 to 2013. During this time, APT10 targeted U.S. and overseas defense contractors with PoisonIvy. After the release of the 2013 FireEye report, which detailed the use of that malware by various actors, APT10 shelved PoisonIvy and began a retooling and replatforming effort. [5] The second phase spanned 2014 through the beginning of 2019. During this interval APT10 used a combination of bespoke and open source tools to conduct its operations. *Operation Cloud Hopper* suggests that APT10 undertook both retooling efforts and espionage campaigns simultaneously, allowing the group to stay ahead of defenders and mitigate the impact of public disclosures.[6] It is notable, therefore, that *Operation Could Hopper* alone appears to have succeeded in disrupting APT10 activity, albeit in combination with an aggressive multi-stakeholder response on client networks—and even then only for seven weeks.[7]

APT10 undertook both retooling efforts and espionage campaigns simultaneously, allowing the group to stay ahead of defenders and mitigate the impact of public disclosures.

Subsequent to this disruption, APT10 activity resumed with the targeting of Japanese government entities and MSPs.[8] This activity ceased after the U.S. Department of Justice indictment (17 December 2018) of two individuals who were believed to be associated with APT10. [9] Private cybersecurity firms corroborated the cessation of activity. This hiatus has apparently held: the TTPs, domains, and IPs associated with APT10 activity have not been publicly reported as of April 2020.

2006 First reports that U.S. and overseas defense contractors and agencies are targeted.[10]

[4] B. Barrett, 'How China's Elite Hackers Stole the World's Most Valuable Secrets' (20 December 2018), *Wired*: https://www.wired.com/story/doj-indictment-chinese-hackers-apt10/

[5] PwC 2017 (n. 2), 5.

[6] id., 16.

[7] Interview with a senior cybersecurity consultant.

[8] Japanese government: A. Matsuda & I. Muhammad, 'APT10 Targeting Japanese Corporations Using Updated TTPs' (13 September 2018), *FireEye*: https://www.fireeye.com/blog/threat-research/2018/09/apt10-targeting-japanese-corporations-using-updated-ttps.html. MSPs: Insikt Group, 'APT10 Targeted Norwegian MSP and U.S. Companies in Sustained Campaign' (6 February 2019), *Recorded Future*: https://www.recorded-future.com/apt10-cyberespionage-campaign/

[9] viz. Zhu Hua and Zhang Shilong.

[10] 'United States of America v. Zhu Hua and Zhang Shilong' (17 December 2018), *Untied States District Court, Southern District of New York*: https://www.justice.gov/opa/press-

2013	**Aug**: FireEye releases PoisonIvy report, documenting APT10 usage of malware.[11]
2014	Beginning of the Cloud Hopper campaign, targeting MSPs around the world.[12] **Late 2014**: APT10 targets European organizations.[13]
2016	**Sep–Nov**: APT10 campaign against Japanese academics and manufacturing organizations.[14] **Late 2016**: APT10 retools and tests Quasar, ChChes and RedLeaves malware families.[15]
2017	**Early 2017**: Beginning of campaign against Ministry of Foreign Affairs in Japan and the National Foreign Trade Council in the United States.[16] **Apr**: *Operation Cloud Hopper* released; APT10 activity against MSPs disrupted temporarily. **Nov–Sep 2018**: revamped campaigns against MSPs and other companies with sensitive intellectual property.[17]
2018	**Dec**: U.S. DoJ indictment released; APT10 activity ceases.

TYPOLOGY OF ATTACKS

APT10 gained access to victims' systems through two primary methods. The first is spear-phishing. Its spear-phishing campaigns displayed keen insight into its targets, typically aligning the malicious email's content with the target's own interests at the time. This implies that APT10 conducted substantial reconnaissance prior to its operations.[18] That procedural decision may also have influenced its preferred vector into a network—namely, decoy documents. [19] In certain cases, such as the campaigns against Japanese organizations, APT10 registered domains with names similar to legitimate entities as another means of deceit.[20] It also used web reconnaissance tools to conduct research on potential targets for later spear-phishing campaigns.[21]

The second infiltration method used by APT10 depended on the infrastructure linking MSPs and their clients. In short, it stole

release/file/1121706/download. Hereafter referenced as DoJ Indictment

[11] FireEye 2013 (n. 1), 25–31.

[12] DoJ Indictment (n. 10), *passim*.

[13] PwC 2017 (n. 2), 17.

[14] J. Miller-Osborn & J. Grunzweig, 'MenuPass Returns with New Malware and New Attacks Against Japanese Academics and Organizations' (16 February 2017), *Palo Alto Networks*: https://unit42.paloaltonetworks.com/unit42-menupass-returns-new-malware-new-attacks-japanese-academics-/organizations/

[15] PwC 2017 (n. 2), 16.

[16] Threat Research Team, 'Operation TradeSecret: Cyber Espionage at the Heart of Global Trade' (5 April 2017), *Fidelis Cybersecurity*: https://www.fidelissecurity.com/threatgeek/threat-intelligence/cyber-espionage-global-trade/

[17] Matsuda & Muhammad (n. 8).

[18] PwC 2017 (n. 2), 16.

[19] id., 11.

[20] id., 12.

[21] e.g. Scanbox: Fidelis Cybersecurity 2017 (n. 16).

or falsified network credentials from the MSP in order to pivot onto the target network.[22] This technique obviated the risk of detection that it would incur by spear-phishing its target directly. The Cloud Hopper campaign made significant use of this technique, generally after a spear-phishing email afforded it access to an MSP's network. Having entered said network, APT10 worked quickly to infect additional systems with malware in order to provide itself sustained access. During the network reconnaissance phase, APT10 sought out usernames and passwords that could directly access the target network.[23]

2006 to Mid-2013

APT10's initial operations targeting the U.S. defense industrial base relied heavily on spear-phishing. In or around 2009, APT10 began to disseminate the PoisonIvy malware in its phishing emails along with several less frequently used RATs. These spear-phishing emails were, again, well-tailored to intended targets and included malicious attachments that infected the victim system.[24]

Mid-2013 to 2014

APT10 underwent a retooling process after FireEye's report on PoisonIvy. The decrease in its operations coincided with the observation from cyber threat intelligence groups that PoisonIvy was no longer being used directly against its targets.[25]

2014 to 2016

In 2014 or shortly before, APT10 began using the PlugX malware in its spear-phishing attacks. Multiple variants of the PlugX software were deployed between 2014 and 2016, and with increasing sophistication. This campaign supported the larger Cloud Hopper operation and its concurrent efforts to spy on Japanese organizations.

Late 2016

APT10 underwent another retooling process, during which it developed and tested versions of the Quasar, ChChes, and RedLeaves malware families. These tools were incorporated slowly, sometimes used with PlugX and PoisonIvy, and were seen until the end of 2018.[26] Its campaign against Japanese academics from September to November 2016 offers a representative example of its operations at this time. The ChChes trojan (or another RAT) was imbedded in an attachment of a spear-phishing email and provided the group with initial access. After the victim opened the malicious attachment, ChChes would beacon back to C2 nodes with an MD5 hash representing the victim; from that point, the infected machine received additional malware through this C2 channel.[27]

2017 to 2018

APT10 continued to use the RedLeaves and ChChes malware in its spear-phishing campaigns, but now alongside a new backdoor,

[22] Cyware 2018 (n. 3).

[23] PwC 2017 (n. 2), 17.

[24] FireEye 2013 (n. 1), 25 f.

[25] PwC 2017 (n. 2), 16.

[26] ib.

[27] e.g. PoisonIvy or PlugX: Miller-Osborn & Grundzweig 2017 (n. 14).

Uppercut. It also employed variations of Quasar and the Trochilus malware family for network persistence and espionage. In 2018 it began to use stolen login credentials for remote desktop applications, such as Citrix, to gain initial access.

Although the publication of *Operation Cloud Hopper* disrupted the campaigns against MSPs and Japanese organizations were disrupted for approximately seven weeks, campaigns against certain Japanese organizations resumed in late 2017. Furthermore, from November 2017 onwards, APT10 reportedly resumed targeting MSPs. In one attack against Visma, a Norwegian MSP, the group used stolen remote desktop credentials to gain access, escalate network privileges, and infect systems with Trochilus malware. It then packaged and exfiltrated sensitive data through Dropbox. It used similar TTPs during this interval to attack an international apparel company, a U.S. law firm specializing in intellectual property law, [28] and companies in the Japanese media sector.[29]

DISCLOSURE EVENTS

Until the cessation of operations in December 2018, APT10 shifted its TTPs and infrastructure multiple times in order to stay ahead of defenders. Although many of the disclosures had little impact on its operations—especially those between April 2017 and September 2018—three disclosures did have a marked effect.

The first effective disclosure was FireEye's 2013 report on PoisonIvy. As mentioned above, usage of that tool decreased after the report was released, though it remains unclear whether APT10 was engaged in contemporaneous campaigns that used other malware. Thereafter, APT10 began a retooling and re-platforming initiative, lasting from 2014 to 2018, that saw the continuous development and deployment of new tools against its targets. This arms-race strategy allowed APT10 to conduct operations largely unimpeded despite a higher number of disclosures in 2017 and 2018.[30]

Operation Cloud Hopper was the next disclosure to have a palpable effect—a disruption in activity for approximately seven weeks. The positive effect of the report was due in part to PwC's coordination of its release with the victims and cybersecurity firms, who used the pre-publication time to catch APT10 unawares and implement plans to regain control of their networks.

The most effective disclosure against APT10 to date was the U.S. Department of Justice indictment of two of the group's members. All recognizable APT10 activity ceased following the indictments, which comports with the impact of similar U.S. indictments of Chinese threat groups associated with the government.

[28] Recorded Future 2019 (n. 8).

[29] Matsuda & Muhammad 2018 (n. 8).

[30] Cobalt Group adopted a similar approach in 2017.

FireEye PoisonIvy Report: August 2013

FireEye's study on observed cases of PoisonIvy was one of the first reports to disclose significant information regarding APT10's activity and TTPs. It detailed three different cyber-espionage campaigns that relied on the PoisonIvy malware, one of which was attributed to APT10. The report provided a technical analysis of PoisonIvy, including its customizable features, configurations, communication methods, and how to defend against it.[31] Further, it expounded APT10's usage of PoisonIvy, attack vectors, weaponization, targets, and passwords, along with the domains and IP addresses used by APT10.[32]

The most effective disclosure against APT10 to date was the U.S. Department of Justice indictment of two of the group's members.

Palo Alto Unit 42 Disclosure: Feb 2017

This report provided information on the APT10 campaign against Japanese academics from September to November 2016. The targeted individuals primarily worked in STEM fields such as Japanese pharmaceutical companies and the U.S.-based subsidiaries of Japanese manufacturing organizations. Unit 42 noted the use of PlugX and PoisonIvy in the campaign, and included a technical analysis of the newly-discovered ChChes malware.

Furthermore, the report identified the infrastructure overlap between APT10's previously reported C2 structure and the domains and IP addresses used in this campaign. It also included indicators of compromise for each malware family and domain names used in APT10's C2 operations.[33]

Operation Cloud Hopper Report: April 2017

PwC's *Operation Cloud Hopper* provided the definitive account of APT10's campaign against MSPs. Part narrative and part technical, the report detailed the pervasive nature of APT10's activities, including targets, apparent goals, and methodology. The various elements of a typical APT10 operation were documented from initial compromise and reconnaissance to exfiltration. A time-based analysis of APT10's activity tied the group to China, while a comparison of APT10's targets to other observed Chinese hacking operations provided further insight into its likely motivations. The detailed timeline and history of APT10's malware, how it was used, and how the group changed it over time to accommodate new goals were particularly noteworthy elements. [34] The technical annex provided a thorough examination of APT10's malware, tools, and its infrastructure from 2009 to mid-2016.[35]

Operation TradeSecret Report: April 2017

Fidelis Security's report *Operation TradeSecret* offered a different vantage into APT10's

[31] FireEye 2013 (n. 1), 9 f.

[32] id., 25–31.

[33] Miller-Osborn & Grundzweig 2017 (n. 14).

[34] PwC 2017 (n. 2).

[35] id., Technical Annex.

initial target reconnaissance activities, the research needed before launching spear-phishing campaigns. Rather than the usual suspect of MSPs and defense contractors, Fidelis Security investigated APT10's targeting of lobbyists involved in U.S. foreign trade policy. APT10 used ScanBox to compromise the webpages that were used to register for meetings of the National Foreign Trade Council. Similarly, it identified the Japanese Ministry of Foreign Affairs as a target. The report also included a more technical report on the TTPs used in these attacks, an analysis of ScanBox, and a list of the domains associated with the activity.[36]

FireEye Disclosure: April 2017

In this disclosure, FireEye documented what they dubbed 'APT10's resurgence' from 2016 to early 2017. They identified and summarized several pieces of malware and discussed attack vectors and methods similar to those identified in *Operation Cloud Hopper*.[37]

Accenture RedLeaves Report: April 2018

Accenture reported that APT10 had again been targeting Japan with the RedLeaves malware. They also analyzed the capabilities of RedLeaves from a sample acquired in January 2018, and they registered four domains, two IP addresses, and indicators of compromise associated with the malware.[38]

Intrusion Truth Disclosures: July to September 2018

Intrusion Truth released several reports during the summer of 2018 which identified four individuals and two companies that may have been associated with APT10 operations. They ultimately tied APT10 to the Tianjin bureau of the Chinese Ministry of State Security.[39]

FireEye Disclosure: September 2018

This report covered APT10's Uppercut/Anel malware campaign against the Japanese media sector. FireEye documented the phishing emails, malicious attachments used in the attack, and the available information on the Uppercut workflow, including APT10's alterations of the malware over time and concomitant changes to the indicators of compromise.[40]

U.S. DoJ Indictment: December 2018

The indictment identified two members of APT10, Zhang Shilong and Zhu Hua. It detailed their activities as employees of the Huaying Haitai Company on behalf of the Tianjin State Security Bureau of the Chinese Ministry of State Security. They were

[36] Fidelis Cybersecurity 2017 (n. 16).
[37] iSIGHT Intelligence, 'APT10 (MenuPass Group): New Tools, Global Campaign Latest Manifestation of Longstanding Threat' (6 April 2017), *FireEye*: https://www.fireeye.com/blog/threat-research/2017/04/apt10_menupass_grou.html
[38] J. Ray, 'Hogfish Alert' (23 April 2018), *Accenture*: https://www.accenture.com/us-en/blogs/blogs-hogfish-alert
[39] 'APT10 Was Managed by the Tianjun Bureau of the Chinese Minsitry of State Security' (15 August 2018), *Intrusion Truth*: https://intrusiontruth.wordpress.com/category/apt10/
[40] Matsuda & Muhammad 2018 (n. 8).

implicated in campaigns dating back to 2006 to steal information from U.S. government agencies and commercial and defense technology companies as well as APT10's global campaign to steal intellectual property through client-MSP infrastructure. The indictment identified attack methodology, TTPs, targets, and the role of the those indicted in these activities. Specifically, the indictment accused them of computer intrusion offensives, wire fraud, and aggravated identity theft.[41]

Recorded Future Disclosure: February 2019

This collaboration between the Insikt Group and Rapid7, published by Recorded Future, identified an APT10 campaign against Visma, a Norwegian MSP, and two other companies between November 2017 and September 2018. APT10 was seen using a new variant of the Trochilus malware alongside the Uppercut backdoor on which FireEye previously reported. APT10 used Citrix remote desktop credentials to access the three victim networks, and the report provides an attack timeline. More technically, it included an analysis of the Trochilus malware, an identification of the encryption used by APT10, and a registry of the domains and indicators of compromise that were associated with the attack.[42]

[41] DoJ Indictment (n. 10), 15–20.

[42] Recorded Future 2019 (n. 8).

COBALT (CRIMINAL GROUP)

Timeline: Activity Levels and Disclosures

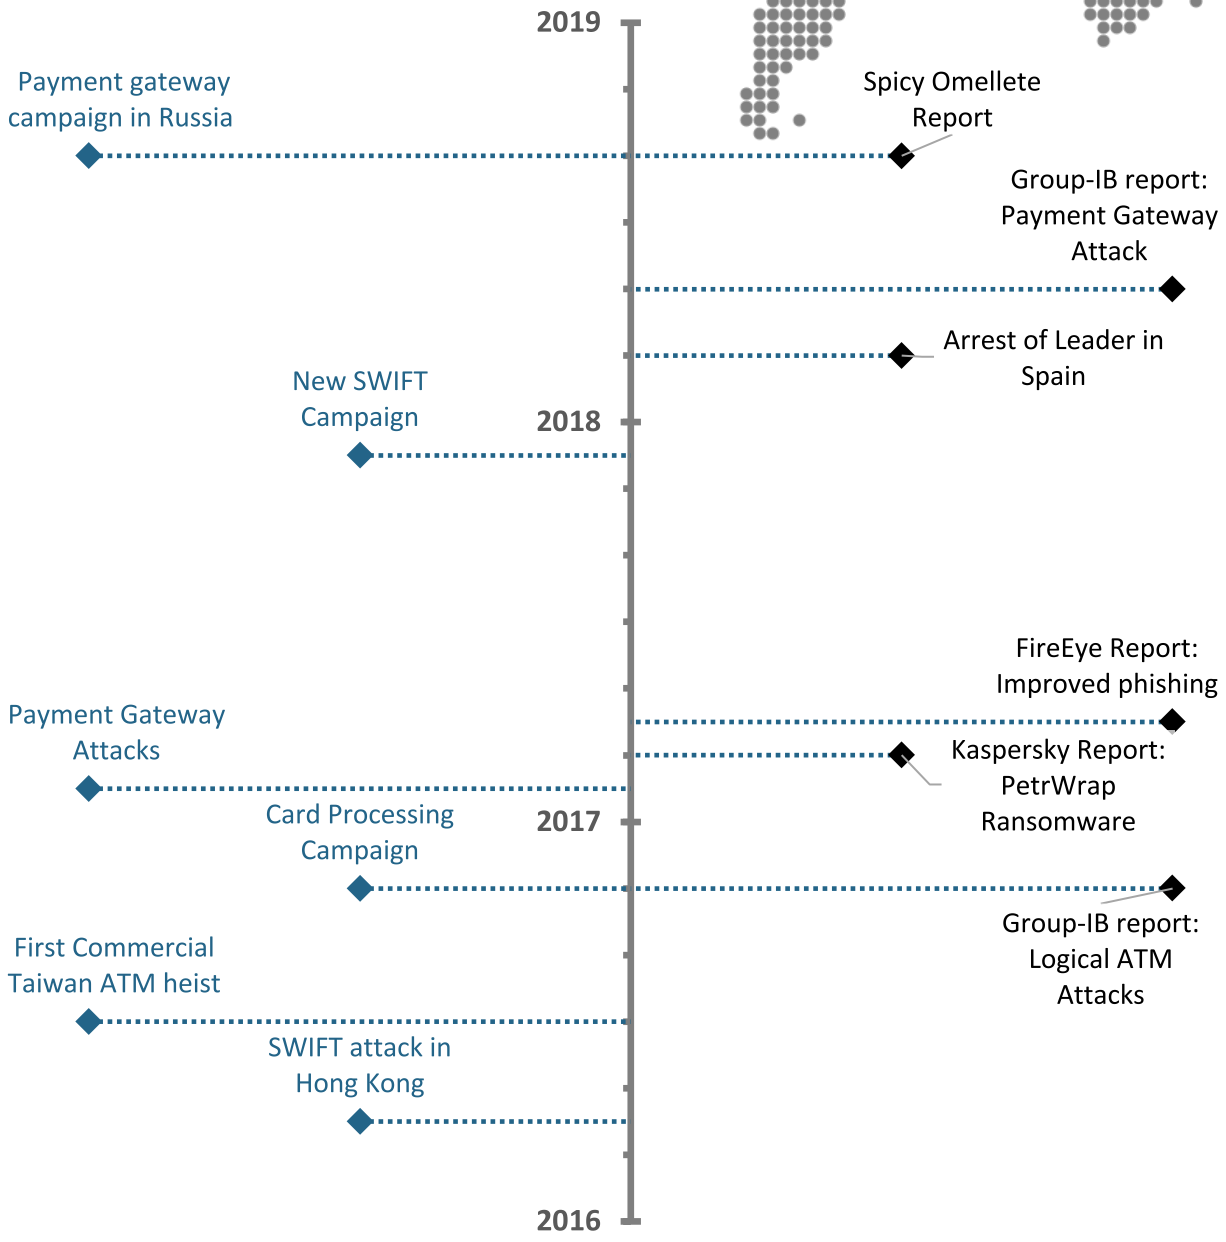

2019

Payment gateway campaign in Russia

Spicy Omellete Report

Group-IB report: Payment Gateway Attack

Arrest of Leader in Spain

New SWIFT Campaign

2018

FireEye Report: Improved phishing

Payment Gateway Attacks

Kaspersky Report: PetrWrap Ransomware

Card Processing Campaign

2017

Group-IB report: Logical ATM Attacks

First Commercial Taiwan ATM heist

SWIFT attack in Hong Kong

2016

IMPACT OF DISCLOSURES

INTRODUCTION

Disclosures have not affected Cobalt Group's attacks and behavioral patterns, nor do they seem likely to undermine its criminal enterprise. Cobalt's profit motive, combined with its development of one-off in-house malware, renders disclosures largely impotent. This motive also allows the group to pivot easily and attack any bank anywhere in the world. Furthermore, the success of Cobalt's operations is not predicated on any single vulnerability or exploit. The Group exploits macro-level structures within a bank's infrastructure: utilizing the legitimate operation of the system rather than hijacking the processes towards another end. Thus it collects cash from ATMs, withdraws cash from debit cards, and uses payment gateways to transfer money to itself. The fix to these hacks therefore depends on an individual bank's combination of software and hardware. Disclosures consequently have limited prophylactic use, if released in time to identify the process currently under attack. But by the time that each disclosure has occurred, Cobalt has moved on.

THE GROUP

Cobalt Group ranks among the most aggressive nonstate cybercriminal groups in the world.[1] It has been responsible for some of the most sophisticated and costly attacks on the financial sector, with Europol estimating that it has stolen over a billion Euros from over a hundred banks in over forty countries—ranging from Russia and former Soviet states to Western Europe and East Asia. The Central Bank of Russia has deemed Cobalt to be the foremost cyberthreat to the Russian financial sector.[2] It has stolen over $1 billion to date.[3]

Since appearing in 2016, the development of new tools and tactics has greatly facilitated Cobalt's success. The eponymous software CobaltStrike provides it a dependable and up-to-date suite of penetration testing and social engineering tools, but the diversity of its campaigns and the novelty of its malware from 2017 strongly suggest that its greatest resource is a suite of in-house

[1] They feature regularly in annual cyberthreat trend reports, e.g. Group-IB, 'Hi-Tech Crime Trends 2018' (https://www.group-ib.com/resources/threat-research/2018-report.html), 'High-Tech Crime Trends 2019/2020' (https://www.group-ib.com/resources /threat-research/2019-report.html); iDefense, '2019 Cyber Threatscape Report', *Accenture Security*: https://www.accenture.com/_acnmedia/pdf-107/accenture-security-cyber.pdf ; Positive Technologies, 'Cybersecurity Threatscape 2018' (18 March 2019), https://www.ptsecurity.com/ww-en/analytics/cybersecurity-threatscape-2018/

[2] J. Kirk, 'Cobalt Cybercrime Gang Reboots After Alleged Leader's Bust' (28 May 2018), *BankInfoSecurity*: https://www.bankinfosecurity.com/billion-euro-cybercrime-gang-reboots-after-arrest-a-11037.

[3] J. Kirk, 'Cobalt Cybercrime Gang Reboots After Alleged Leader's Bust' (28 May 2018), *BankInfoSecurity*: https://www.bankinfosecurity.com/billion-euro-cybercrime-gang-reboots-after-arrest-a-11037.

developers who design and write new malware.[4] The structure of the group also appears to be fluid and adaptable to changes in personnel: Cobalt's operations saw no immediate effect from the apprehension of its leader in Alicante, Spain[5] or from subsequent and related arrests.[6] Europol released news of the leader's arrest (28 March 2018) well after the fact,[7] but this did nothing to retard a new attack in southeast Asia that April.

Lastly, Cobalt possibly overlaps with other known cybercriminal organizations from Eastern Europe, specifically the Buhtrap[8] and Carbanak gangs.[9] The diversification and expansion of Cobalt's workforce offers another defense against disclosures, insofar as it dilutes the risk to the Group's own infrastructure; but perhaps more importantly, cooperation with other (highly technical and successful) cybercriminal organizations expands its capabilities and available vectors of attack. Connections to these groups are registered below.

Buhtrap	Carbanak
TTPs are identical or very similar to those used by Carbanak.[10]	Cobalt's 2017 attacks used a secure shell (SSH) backdoor (previously unique to Carbanak's 2014 campaign) and used a similar theft schema.[11]
Cobalt's first attack occurred three months after the disclosure of the source code for the Buhtrap malware. (Cobalt takes two	Carbanak's SSH backdoors and C2 addresses were located on the same subnets,[12] which shows that a single group likely controlled them. The same is true for the SSH backdoor and CobaltStrike C2 server used

[4] This expansion of capacity is often referred to as Cobalt's "evolution", for example: 'Cobalt Strikes Back: an Evolving Multinational Threat to Finance' (1 August 2017), *Positive Technologies*: https://www.ptsecurity.com/upload/corporate/ww-en/analytics/Cobalt-2017-eng.pdf; 'Cobalt: Evolution and Joint Operations' (May 2018), *Group-IB*: https://www.group-ib.com/resources/threat-research/cobalt-evolution.html

[5] J. Vijayan, '2018 Arrests Have Done Little to Stop Marauding Threat Group' (8 May 2018), *Dark Reading*: https://www.darkreading.com/attacks-breaches/2018-arrests-have-done-little-to-stop-marauding-threat-group/d/d-id/1334652; Kirk 2018 (n. 2).

[6] '**Кіберполіція викрила українського хакера у взламі комп'ютерів світових банків та готелів**' (26 March 2018), *Cyber Police of Ukraine* (**Кіберполіції України**): https://cyberpolice.gov.ua/news/kiberpolicziya-vykryla-ukrayinskogo-xakera-u-vzlami-kompyuteriv-svitovyx-bankiv-ta-goteliv-4470/

[7] Europol, 'Mastermind Behind €1 Billion Cyber Bank Robbery Arrested in Spain' (26 March 2018): https://www.europol.europa.eu/newsroom/news/mastermind-behind-eur-1-billion-cyber-bank-robbery-arrested-in-spain

[8] A portmanteau of *buhgalter* (**бухгалтер**, "accountant") and the English word "trap".

[9] See further: Group-IB 2018 (n. 4).

[10] J-I. Boutin, 'Operation Buhtrap, the Trap for Russian Accountants' (9 April 2015), *WeLiveSecurity*: https://www.welivesecurity.com/2015/04/09/operation-buhtrap/

[11] The backdoor had identical RSA and public keys installed on the two targeted Linux servers for data transfer to the C2 servers.

[12] 190.123.35.0/24, 190.123.36.0/24.

months on average for infiltration and reconnaissance.)	by Cobalt Group in 2017,[13] which again suggests single ownership of both elements.
The phishing emails used in the 2016 SWIFT attack employed techniques identical to those used by the Buhtrap group.	Cobalt's campaigns have become less technically complex over time: relying more on preexisting elements of bank's networks and requiring the compromise of less software. This streamlining more closely resembles Carbank's criminal infrastructure.

TIMELINE

Cobalt has undertaken four major campaigns to date. From the information available,[14] no clear pattern exists between its heists or its use of new TTPs and disclosure events. Cobalt appears to have a pattern of activity and rest, approximately every two months, which is possibly a deliberate aspect of its business model or strategy. The timeline cannot adequately show the extent to which Cobalt appears to shift towards targets in Asia following a disclosure or otherwise pivot away from Western Europe and the U.S. with its phishing emails. Despite the apparent shift, there are not enough instances of this to label it a true behavioral pattern.

2016

Mar: Last attack from Buhtrap until 2019;[15] activation and configuration of the servers used in the upcoming Hong Kong SWIFT attacks.

Apr: Theft through SWIFT in Hong Kong; theft through SWIFT in Ukraine.[16]

Jun: First attack in Russia as Cobalt.[17]

Jul: Successful attack against First Commercial Bank ATMs in Taiwan.[18]

[13] 89.37.226.0/24.

[14] The definitive source for which, to date of publication: Group-IB 2018 (n. 4).

[15] P. Paganini, 'Buhtrap Group Stole Tens of Millions of Dollars from Russian Banks' (18 March 2016), *Security Affairs*: https://securityaffairs.co/wordpress/45405/cyber-crime/buhtrap-group-attacks.html. Note, however, that erroneous reports of planned activity continued, e.g. E. Gerden, 'Russian Hacker Group Targeting Largest EU Banks' (11 April 2016), *SC Magazine UK*: https://www.scmagazine.com/home/security-news/russian-hacker-group-targeting-largest-eu-banks

[16] J. Kovensky, 'Hackers Reportedly Steal $10 Million from a Ukrainian Bank Through SWIFT Loophole' (25 June 2016), *Kyiv Post*: https://www.kyivpost.com/article/content/ukraine-politics/hackers-steal-10-million-from-a-ukrainian-bank-through-swift-loophole-417202.html

[17] V. Mateeva, 'Secrets of Cobalt' (15 August 2017), *Group-IB*: https://www.group-ib.com/blog/cobalt

[18] L. Chung, 'How Taiwanese Police Cracked NT$83 Million ATM Heist' (6 August 2016), *South China Morning Post*: https://www.scmp.com/news/china/money-wealth/article/1999019/how-taiwanese-police-cracked-nt83-million-atm-heist; C. Cimpanu, 'Events Behind July 2016 Taiwan ATM Heists Are Coming to Light' (27 January 2017), *Bleeping Computer*: https://www.bleepingcomputer.com/news/security/events-behind-july-2016-taiwan-atm-heists-are-coming-to-light/

Sept: Preparation for the upcoming campaign against card processors.

Nov: First successful attack on card processing in Kazakhstan; ATM jackpotting spree in Europe;[19] Group-IB publishes a report on its ATM hacks.[20]

2017 **Feb**: First attack using PetrWrap ransomware at a small Russian bank;[21] beginning of supply chain attacks on card processing service providers;[22] compromise of an IT integrator.[23]

Mar: Attacks on e-wallet and payment terminal companies; phishing emails see improved quality; banks penetrated through their supply chains.[24]

Apr: First attack on a payment gateway using a unique program.[25]

May: New JavaScript backdoor implemented; first use of decoy documents.[26]

Aug: First attack on Russian telecommunications companies; attack thwarted, goals unclear.[27]

[19] J. Finkle, 'Hackers Target ATMs Across Europe as Cyber Threat Grows' (21 November 2016), *Reuters*: https://in.reuters.com/article/cyber-banks-atms-idINKBN13G254; Z. Zorz, 'Cobalt Hackers Executed Massive, Synchronized ATM Heists Across Europe, Russia' (22 November 2016), *Help Net Security*: https://www.helpnetsecurity.com/2016/11/22/cobalt-hackers-synchronized-atm-heists/

[20] 'Cobalt: Logical Attacks on ATMs' (November 2016), *Group-IB*: https://www.infosecurityeurope.com/__novadocuments/459980?v=636576764177630000

[21] A. Ivanov & F. Sinitsyn, 'PetrWrap: The New Petya-Based Ransomware Used in Targeted Attacks' (14 March 2017), *Kaspersky*: https://securelist.com/petrwrap-the-new-petya-based-ransomware-used-in-targeted-attacks/77762/; L. Abrams, 'The Week in Ransomware - March 17th 2017 - Revenge, PetrWrap, and Captain Kirk' (18 March 2017), *Bleeping Computer*: https://www.bleepingcomputer.com/news/security/the-week-in-ransomware-march-17th-2017-revenge-petrwrap-and-captain-kirk/

[22] Positive Technologies 2017 (n. 4), 3; I. Arghire, 'Cobalt Hackers Now Using Supply Chain Attacks' (2 August 2017), *Security Week*: https://www.securityweek.com/cobalt-hackers-now-using-supply-chain-attacks

[23] 'Cobalt's Latest Attacks on Banks Confirm Connection to Anunak' (29 May 2018), *Group-IB*: https://www.group-ib.com/media/group-ib-cobalts-latest-attacks-on-banks-confirms-connection-to-anunak/

[24] ib.

[25] Group-IB 2018 (n. 4), 18–20.

[26] id., 30, 21 resp.; M. Gorelik, 'Cobalt Group 2.0' (8 October 2018), *Morphisec*: https://blog.morphisec.com/cobalt-gang-2.0

[27] L. Bermejo, R. Giagone, R. Wu, F. Yarochkin, 'Backdoor-Carrying Emails Set Sights on Russian-Speaking Businesses' (7 August 2017), *Trend Micro*: https://blog.trendmicro.com/trendlabs-security-intelligence/backdoor-carrying-emails-set-sights-on-russian-speaking-businesses/

Sept: Joint attack with Carbanak against payment gateways; use of new program, InfoStealer v. 0.2, which is wholly hosted in memory.[28]

Nov: Nov. 14, CVE-2017-11882 is published and patched by Microsoft; [29]Nov. 21, proof of concept is published for that exploit on GitHub;[30] Cobalt immediately begins a large phishing campaign against financial institutions with a malicious document that initially goes undetected by AV.[31]

Dec: Cobalt's spear phishing emails contain a link to a new Java applet (dropper).

2018 **Feb**: Cobalt stops using Java.

Mar: Cobalt's leader and chief malware author is arrested in Spain; activity reduced in former Soviet Union; Cobalt begins phishing campaign under the guise of U.S. companies IBM and Verifon and the international organization Spamhaus.[32]

Apr: Cobalt is detected in a Swedish company's network.[33]

May: Cobalt begins new phishing campaign against Russia and post-Soviet states under the guise of Kaspersky.[34]

Aug: SWIFT and card processing attacks in Russia.[35]

TYPOLOGY OF ATTACKS[36]

A key aspect of Cobalt's relative immunity to information disclosure events is the strategy behind its attacks. Aside from regularly changing its tactics and victims, Cobalt's attacks do not hinge on any single vulnerability

[28] Group-IB 2018 (n. 4), 34 f.

[29] See further: National Vulnerability Database, 'CVE-2017-11882 Detail' (14 November 2017), *NIST*: https://nvd.nist.gov/vuln/detail/CVE-2017-11882; T. Spring, 'Microsoft Patches 17-Year-Old Office Bug' (15 November 2017), *Threat Post*: https://threatpost.com/microsoft-patches-17-year-old-office-bug/128904/

[30] https://github.com/embedi/CVE-2017-11882

[31] J. Manual & J. Salvio, 'Cobalt Malware Strikes Using CVE-2017-11882 RTF Vulnerability' (27 November 2017), *Fortinet*: https://www.fortinet.com/blog/threat-research/cobalt-malware-strikes-using-cve-2017-11882-rtf-vulnerability.html. The bug was still a viable vector of attack two years later: T. Seals, 'Microsoft Warns of Email Attacks Executing Code Using an Old Bug' (10 June 2019), *Threat Post*: https://threatpost.com/microsoft-arbitrary-code-execution-old-bug/145527/

[32] Group-IB 2018 (n. 4), *passim*.

[33] id., 7.

[34] ClearSky Cyber Security, '2018 Cyber Events Summary Report "Year of the Dragon"' (2019), 25: https://www.clearskysec.com/wp-content/uploads/2019/02/ClearSky-End_of_Year_Report-2018.pdf

[35] ASERT Team, 'Double the Infection, Double the Fun' (30 August 2018), *NetScout Systems*: https://www.netscout.com/blog/asert/double-infection-double-fun

[36] The detailed information in this section is only available in the Group-IB report *Cobalt: Evolution and Joint Operations* from May 2018 (n. 4). That document is the definitive account of Cobalt's operations from inception through the date of publication and is the source for what follows.

in a network or brute force techniques—unlike, for instance, many hacks of cryptocurrency exchanges. Instead, Cobalt coopts the normal processes and workflow of banks' networks in small but meaningful ways to siphon out funds. Disclosures may curtail the phishing stage of a new campaign or signal that a cybersecurity firm has removed Cobalt from a network, but a Cobalt heist has yet to be caught *in medias res*. Nevertheless, whether a disclosure would deter Cobalt is not immediately clear because its attacks, as described below, target universal aspects of banks' infrastructure. The usual specifics included in a disclosure may not necessarily translate to another bank's situation in the same way that a report about a new trojan would.

Cobalt coopts the normal processes and workflow of banks' networks in small but meaningful ways to siphon out funds.

Hong Kong SWIFT

Cobalt Group's first recorded activity was a campaign to steal money from a bank in Hong Kong through the SWIFT system. This attack began on 20 March 2016, with the uploading of the Cobalt Strike payload to a server in Germany which subsequently attacked the HK bank. This was repeated two days later with a server in the US; however, the Metasploit framework was also found on that server, which may suggest that another group was involved in the attack. Cobalt was able to transfer money through the SWIFT system by compromising the credentials of SWIFT operators in the bank, implementing a unique JavaScript in authorization forms and unique malware to search for SWIFT payment confirmation messages.

The Group gained initial entry into the victim bank's network with a phishing email. With the malware in place on the workstation and in contact with the C2 server, it waited until the SWIFT operator logged into SWIFT portal. The malware then embedded a malicious JavaScript script in the original portal login page which logged the credentials of the operator and relayed them to the C2 server. Cobalt used the stolen credentials to access the SWIFT server and compromise it with a credential harvester and a log suppressor targeting fraudulent payment messages. It then used these SWIFT server credentials to initiate fund transfers. The theft itself occurred on 28 April. Whether the Group needed the entire month to prepare or if other considerations were involved is unclear.

Jackpotting ATMs

July 2016 saw the culmination of Cobalt Group's second major campaign. This time, it targeted the ATM network of First Commercial Bank in Taiwan. To have ATMs surrender their cash, Cobalt needed to find the segment of the Bank's internal network that controlled the ATMs. The Group again achieved initial compromise through phishing emails sent from two servers in Russia in June 2016 under the guise of the European Central Bank, Wincor Nixdorf (an ATM manufacturer), and local banks. If the exploit contained in the phishing email was successful at compromising the recipient's workstation,[37] the malware injected the Cobalt Strike payload "Beacon" into the memory, which allowed the Group to deliver further payloads to the workstation and ultimately control it.[38] Cobalt then expanded its control over workstations on the network—building, in effect, a network of infected machines within the Bank's own network. A single CobaltStrike console installed on a remote server could then control this subnetwork of infected workstations and use it to branch out to further hosts on the internal network. Researchers first documented the use of the Cobalt Strike Beacon against a financial institution the month before the attack in Taiwan, in June 2016 against a bank in Russia.

With control of the Bank's network and redundant access channels established, Cobalt sought out the workstations of employees who interfaced with the ATM servers and software. With access to these machines, it instructed individual ATMs to dispense cash from the different denomination cassettes within the physical machine. Money mules would appear at the ATMs, make a phone call, and then collect the cash in large duffle bags as it was ejected.

The methods and malware used to compromise First Commercial Bank in Taiwan are functionally identical to those used by the Buhtrap Group against Russian banks' ATMs during the interval of August 2015 to January 2016.[39]

Credit/Debit Card Processing

In September 2016, two months after the attack on First Commercial Bank's ATMs, Cobalt Group gained access to the network of a Kazakh bank. Over the next two months, the Group prepared to conduct a new manner of attack; namely, theft through a bank's card processing system. A major issue with Cobalt's previous operation in Taiwan (and subsequent attempts in Russia and Romania) was that local law enforcement tended to find and arrest the money mules—and with some rapidity.[40] A plausible element in the

[37] Using CVE-2015-1641.

[38] The initial injection into memory is not permanent, and it would be lost should the system restart. To gain permanent control over the workstation, Beacon scans for applications included in autorun and substitutes those legitimate files with identically named malicious executables. Hence normal services will automatically launch malicious applications after a restart or other even that wipes the memory clear. That does not, however, establish permanent access to the local network: the infected workstation can still be shut down or have its OS changed/reinstalled. The Group therefore seeks to escalate its privileges on the network to establish their access to the internal network independent of any single workstation.

[39] The last confirmed attack by the Buhtrap Group on a bank was March 2016. The criminal group used to launder the stolen money was arrested in May 2016; and June 2016 registers the first attack on a Russian bank with Cobalt Strike.

[40] J. Pan, 'East European Trio Indicted Over First Commercial Heist' (14 September 2016), *Taipei Times*: https://www.taipeitimes.com/News/front/archives/2016/09/14/2003655108; T. Ferry, 'Looking Back at the First Bank's ATM Heist' (15 February

mules' quick apprehension was the coincidence of the bank's location and the location of the hacked ATMs, which is to say that First Bank only had ATMs in Taiwan. Cobalt now sought to cash out from ATMs independent of the bank to which the account belonged or not in the same country as the bank's headquarters. This new strategy also simplified the collection of the stolen cash: the Group's hacker wing no longer needed to coordinate a cyberattack with the mules' physical presence at the correct ATM. Moreover, withdrawing cash in another county puts distance between the mules, local law enforcement, and the bank's security team.

The operation itself was markedly simpler than prior operations. The first phase was identical to what was seen at First Commercial Bank, beginning with phishing emails. The malicious attachments in these emails contained the CobaltStrike payload, which the Group then used to establish an infected sub-network on the bank's internal network. Over the next two months (9 September–10 November), it used CobaltStrike to collect data on domains and local user accounts. After Cobalt established persistence in the network, it spent three weeks conducting reconnaissance on the card processing systems while members opened accounts at the bank. A brief hiatus ensued as the organization waited for the debit cards to be delivered. On 18 December 2016 (Sunday), the Group effected the necessary changes to its accounts; and in Russia, Latvia, Estonia, France, Austria, Germany, the Netherlands, and Belgium, mules began to withdraw cash. Bank officials quickly noticed the changes on Monday and had all illicit accounts deactivated before noon.

Payment Gateway

In late March 2017 Cobalt began a spear-phishing campaign against electronic wallet and payment terminal companies in Russia and Ukraine. During the reconnaissance phase, it discovered several payment gateway servers that process requests to transfer money. These gateways are most often used to convey small sums from one account to another: Cobalt had to automate transactions to steal a large amount. To do so, it created a new program that sent fraudulent transfer requests and the criminal recipient's account information.

Here again, Cobalt had the technical ability to steal money through various parts of a bank's network but lacked a safe means to cash out. Nevertheless, targeting a small-transactions gateway, which processes thousands of requests each day, afforded it some degree of cover via the volume of transaction traffic and the paucity of the sums, which it used to evade fraud detection and countermeasures.

DISCLOSURE EVENTS

As mentioned above, disclosures of Cobalt's TTPs and activity have only occurred *post facto*. Banks, both victims and near-victims,

2017), *Taiwan Business Topics*: https://topics.amcham.com.tw/2017/02/looking-back-at-the-first-banks-atm-heist/

have largely remained silent on any relevant incidents. But their reticence sounds a deafening undertone for the broader threatscape, which is itself starved for reports that are contemporary, detailed, and public. With over $1.2 billion stolen to date, one would be justifiably incredulous of the idea that banks, cybersecurity firms, or law enforcement lack a reliable means to preclude Cobalt from bank robberies or otherwise obviate its four known techniques. And, yet, such seems to be the case. Most of the relevant disclosures from cybersecurity firms are phishing detection notices, reports on specific malware, or analyses of completed campaigns. Only Group-IB's (non-public) report *Cobalt Evolution* detailed how Cobalt steals money. Indeed, perhaps the largest setback to one of Cobalt's campaigns was self-inflicted: it forgot to blind carbon copy the 1,880 targets in a Kazakh bank spear-phishing campaign during March 2017, and instead put them all in the "to" field—a mistake it made again, in November of that year, for a campaign against Russian and Turkish banks.[41] Even with such forewarning, Cobalt had several successful attacks against Russian banks later that year. That is perhaps the most damning evidence against the utility of public disclosures.

Perhaps the largest setback to one of Cobalt's campaigns was self-inflicted

What follows is not a comprehensive list, but a selection of the most important and most salient reports. For reasons of time and manpower, it lies beyond the remit of this project to sift through the vast number of warnings and advisories from cybersecurity firms and blogs that occur on any given Cobalt development. That is not to say, however, that they lack importance, merely that there is no open-source method to determine their effect while staving away conjecture.

PetrWrap Ransomware Report: May 2016

Kaspersky revealed the existence of a new Petya variant and provided a full technical analysis.[42] A little less than a year later, in February 2017, Cobalt began using the ransomware.

Logical ATM Attacks: November 2016

During the spate of ATM attacks in Europe, Group-IB published a report on the First Bank ATM heist from July that year.[43] The ATM attacks resumed in Russia approximately one year later.

Positive Technologies: After May 2017

Sometime between May 2017 and the new year, Positive Technologies published a review of Cobalt's phishing emails from earlier that year and provided a general explanation of how the malware worked.[44] This

[41] Y. Klijnsma, 'Gaffe Reveals Full List of Targets in Spear Phishing Attack Using Cobalt Strike Against Financial Institutions' (28 November 2017), *RiskIQ*: https://www.riskiq.com/blog/labs/cobalt-strike/

[42] F. Sinitsyn, 'Petya: The Two-In-One Trojan' (4 May 2016), *Kaspersky*: https://securelist.com/petya-the-two-in-one-trojan/74609/

[43] Group-IB 2016 (n. 20).

[44] Positive Technologies 2017 (n. 4).

report was non-technical and so wanted for the usual IPs, TTPs, and hashes that one might expect in a report of its length.

Arrest of Leader in Spain: March 2018

Europol arrested the supposed leader of Cobalt Group,[45] with subsequent arrests in Ukraine.[46] Operations continued without a perceptible pause.

Group-IB 'Cobalt: Evolution' Report: May 2018

Group-IB published a comprehensive report on Cobalt Group's to-date activity, but only for paying customers.[47] It provided a detailed explanation for each of Cobalt's four known theft techniques in addition to the full suite of technical indicators for its known malware.

Bitdefender's APT Blueprint: After May 2018

Sometime after May 2018 (publication date is unlisted), Bitdefender published a white-paper that detailed every step in Cobalt's mid-2018 spear-phishing campaign, which had notably pivoted away from East Europe and the former Soviet Union. The report was highly detailed, including all the usual technical indicators and a complete timeline of its on-network activity that showed its activity on a minute-by-minute basis.[48] Cobalt merely refocused on Russia and the former Soviet states during the second half of the year, and with its usual success.

ASERT Team: August 2018

Shortly before Cobalt could count any success in its new Russian bank campaign, ASERT put out a technical analysis of the malware which it used in that phishing campaign. Slightly more than half a month separated initial detection and publication.[49]

[45] Europol 2018 (n. 7).

[46] Ukrainian Cyber Police 2018 (n. 6).

[47] Group-IB 2018 (n. 4).

[48] 'An APT Blueprint: Gaining New Visibility into Financial Threats', *Bitdefender*: https://www.bitdefender.com/files/News/CaseStudies/study/262/Bitdefender-WhitePaper-An-APT-Blueprint-Gaining-New-Visibility-into-Financial-Threats-interactive.pdf

[49] ASERT Team 2018 (n. 35).

APT33 (IRAN)

Timeline: Activity Levels and Disclosures

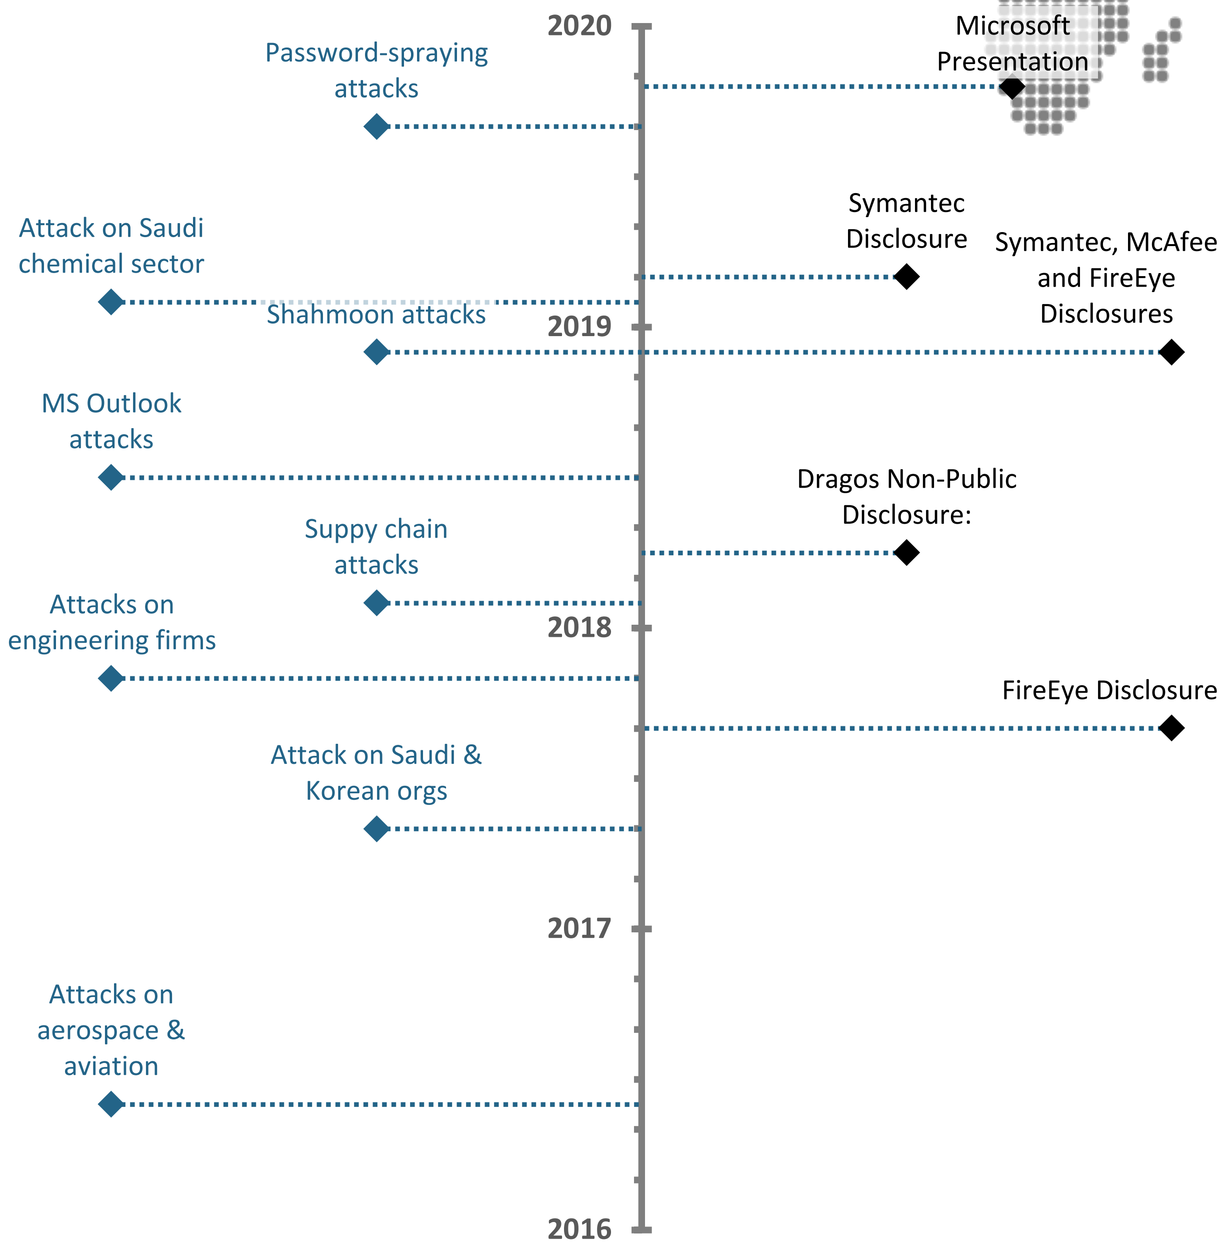

2020

Microsoft Presentation

Password-spraying attacks

Symantec Disclosure

Symantec, McAfee and FireEye Disclosures

Attack on Saudi chemical sector

Shahmoon attacks

2019

MS Outlook attacks

Dragos Non-Public Disclosure:

Suppy chain attacks

Attacks on engineering firms

2018

FireEye Disclosure

Attack on Saudi & Korean orgs

2017

Attacks on aerospace & aviation

2016

IMPACT OF DISCLOSURES

INTRODUCTION

APT33 is an Iranian group whose actions align with the strategic and political objectives of the Iranian government, primarily its national defense and economic priorities—though it remains unclear how closely the group is linked to the regime in Tehran. As a cyber actor, it has been resilient against information disclosures: various cyber threat intelligence groups have disclosed TTPs and network artifacts with little perceivable effect (deterrent or otherwise) on the frequency of APT33's operations or capabilities.[1] Instead, APT33 has continued to expand its list of targets and increase the destructiveness of its attacks, which have shifted from espionage to sabotage through wiper attacks.

THE GROUP

APT33 (Elfin, Magnallium, Holmium)[2] is generally accepted to be an Iranian cyberthreat actor. The extent of its ties to the Iranian government are unclear, though at least one of its members has been linked to the Nasr Institute—a group controlled directly by the Iranian government. Yet signs of this connection exist elsewhere. It employs hacking tools that are commonly used by Iranian hackers, as well as DNS servers suspected to be used by other Iranian cyber groups; and, tellingly, its hours of operation coincide with the Iranian workday and workweek.

Cyber threat intelligence firms have tracked APT33 since at least 2013.[3] The group did not, however, become a major threat actor until 2017, when it and other Iranian-linked groups became increasingly active and increased their presence, persistence, and sophistication. [4] Throughout this interval, APT33 primarily targeted Middle Eastern countries and companies that Tehran views as regional rivals—Saudi Arabia, the UAE, Jordan, and Morocco.[5] But its scope is not limited to that region. APT33 has also attacked the United States, the United Kingdom, Italy, Germany, Belgium, the Czech

[1] Finding confirmed in an interview with FireEye analyst Ben Read.

[2] Respectively: 'APT33' (28 June 2019), *Mitre*: https://attack.mitre.org/groups/G0064/; 'Magnallium', *Dragos*: https://dragos.com/resource/magnallium/; AFP, 'Iranian Hackers Caused Losses in Hundreds of Millions: Microsoft Researchers' (7 March 2019), *Radio Farda*: https://en.radiofarda.com/a/iranian-hackers-caused-losses-in-hundreds-of-millions-microsoft-researchers/29808137.html

[3] J. O'Leary, J. Kimble, K. Vanderlee, N. Fraser, 'Insights into Iranian Cyber Espionage: APT33 Targets Aerospace and Energy Sectors and Has Ties to Destructive Malware' (20 September 2017), *FireEye*: https://www.fireeye.com/blog/threat-research/2017/09/apt33-insights-into-iranian-cyber-espionage.html

[4] '2017 in Cyber Perspective: the Rise of Iranian Threat Actors, Repeat Attacks and Other Trends' (5 April 2018), *Cyware*: https://cyware.com/news/2017-in-cyber-perspective-the-rise-of-iranian-threat-actors-repeat-attacks-and-other-trends-76013e05

[5] ib.

Republic, China, South Korea, and India.[6] Many of its operations outside the Middle East have focused on companies with ties to the Middle East through holdings and business but which are based outside the region.[7]

As APT33 has targeted more countries over time, the number of industries affected has likewise expanded. This trend began in 2017, before which it was primarily targeting military and commercial aviation companies and the energy sector (specifically petrochemical production).[8] Since then, APT33 has moved on to the aerospace, oil and gas, electric, government, research, chemical, engineering, manufacturing, consulting, finance, and telecommunications sectors, as well as the manufacturers, suppliers, and maintainers of industrial control systems equipment.[9]

TIMELINE

Since its formation in 2013, APT33 has carried out numerous campaigns and attacks. Although it seems to alternate between periods of attack and rest, operations have appeared continuously since redoubling its efforts in 2017; and in 2018 it expanded its focus to include Europe and North America, beyond its 2017 remit of the Middle East.[10] Its activity has been particularly visible since late 2018, as its TTPs began to change more regularly and its attacks took on a higher degree of frequency and destructivity. Password-spraying attacks, for instance, were conducted throughout most of 2019.[11] Therefore, Intervals of low reported activity cannot be assumed to represent periods of dormancy: such a hiatus more likely represents a preparatory phase for the next wave of attacks.

[6] The reporting on these sundry attacks is predictably voluminous. The following reports illustrate the widening scope of APT33's targets: O'Leary *et al.* 2017 (n. 3); Critical Attack Discovery and Intelligence Team, 'Shamoon: Destructive Threat Re-Emerges with New Sting in Its Tail' (14 December 2018), *Symantec*: https://www.symantec.com/blogs/threat-intelligence/shamoon-destructive-threat-re-emerges-new-sting-its-tail; AFP, 'Iranian Hackers Caused Losses in Hundreds of Millions: Microsoft Researchers' (7 March 2019), *Radio Farda*: https://en.radiofarda.com/a/iranian-hackers-caused-losses-in-hundreds-of-millions-microsoft-researchers/29808137.html; Critical Attack Discovery and Intelligence Team, 'Elfin: Relentless Espionage Group Targets Multiple Organizations in Saudi Arabia and U.S.' (27 March 2019), *Symantec*: https://www.symantec.com/blogs/threat-intelligence/elfin-apt33-espionage

[7] 'Magnallium', *Dragos*: https://dragos.com/resource/magnallium/

[8] O'Leary *et al.* 2017 (n. 3).

[9] Some illustrative examples: Dragos (n. 2); Critical Attack Discovery and Intelligence Team 2018 (n. 6); Critical Attack Discovery and Intelligence Team 2019 (n. 6); A. Greenberg, 'A Notorious Iranian Hacking Crew Is Targeting Industrial Control Systems' (11 November 2019), *Wired*: https://www.wired.com/story/iran-apt33-industrial-control-systems/

[10] Dragos (n. 2).

[11] Greenberg 2019 (n. 9).

2013

APT33 forms.[12]

2016

Mid-2016 to early 2017: APT33 compromises aerospace and aviation companies in Saudi Arabia and the U.S.; compromises oil and petrochemical company in South Korea.[13]

2017

May: APT33 targets a Saudi organization and a South Korean business conglomerate.[14]

Sep: FireEye releases first major disclosure on APT33 activities.[15]

Nov: APT33 uses stolen credentials and publicly available tools, targeting within the engineering sector to escalate privileges and steal more credentials.[16]

2018

Dragos releases a non-public disclosure on APT33.[17]

Feb: APT33 compromises a U.S. company to steal information.[18]

July-Aug: APT33 uses stolen credentials and publicly available tools, again targeting within the engineering sector to modify users' MS Outlook client homepages to enable code execution and persistence.[19]

Dec: APT33 uses Shamoon and a new wiper variant against numerous targets in the oil and gas sector;[20]Symantec, McAfee, and FireEye report on the campaign.

2019

Feb: APT33 conducts remote access and privilege exploitation against a Saudi chemical sector target.[21]

Mar: Symantec reports on the February attack.[22]

Sep-Nov: APT33 conducts narrowed password-spraying attacks against smaller number of organizations, but with increased number of attacks against those organizations; focus on gaining entry to ICS.[23]

Nov: Microsoft reports on the password-spraying campaign at CYEBRWWARCON.[24]

[12] Mitre 2019 (n. 2).

[13] O'Leary *et al.* 2017 (n. 3).

[14] ib.

[15] ib.

[16] G. Ackerman, R. Cole, A. Thompson, A. Orleans, N. Carr, 'OVERRULED: Containing a Potentially Destructive Adversary' (21 December 2018), *FireEye*: https://www.fireeye.com/blog/threat-research/2018/12/overruled-containing-a-potentially-destructive-adversary.html

[17] Dragos (n. 2).

[18] Critical Attack Discovery and Intelligence Team 2019 (n. 6).

[19] Ackerman *et al.* 2018 (n. 16).

[20] Critical Attack Discovery and Intelligence Team 2018 (n. 6).

[21] Critical Attack Discovery and Intelligence Team 2019 (n. 6).

[22] ib.

[23] Greenberg 2019 (n. 9).

[24] ib.

TYPOLOGY OF ATTACKS

APT33 regularly uses spear-phishing emails containing fraudulent job opportunities and offers to gain initial access to targets. But as its capabilities have grown, so too have its means of entry. It has been seen utilizing pre-stolen credentials and exploiting publicly available vulnerabilities, which culminated in the mass password-spraying attacks of 2019.

Public disclosures are unlikely to affect APT33 in any meaningful manner, and for various reasons. First, it uses both commodity and bespoke tools, which allows it to switch between tools with relative ease. Second, although targeted organizations monitor their networks for the presence of known malware signatures, they cannot consistently protect their employees from well-made spear phishing emails. So long as employees continue to click on links and open malicious attachments, APT33 will be able to gain a foothold in the network. Finally, APT33 serves as an Iranian proxy group rather than an explicit part of the Iranian government, thereby granting Tehran plausible deniability for APT33's actions. While APT33 does not particularly care about its tools and its targets being disclosed, it does care about being directly linked to the Iranian government and thus takes great pains to carefully manage its infrastructure and activities to prevent researchers from connecting the two.[25] As a result, there is a dearth of incentive for Tehran to pressure the group into inactivity.

APT33 does not particularly care about its tools and its targets being disclosed.

Mid-2016 through Mid-2017

Prior to this period, APT33 relied heavily on spear-phishing to gain access into networks, primarily those of Saudi companies. In 2016 it began domain masquerading to improve the quality of its phishing emails,[26] which themselves contained fallacious job opportunities and offers. These attacks saw the use of multiple custom backdoors in conjunction with publicly available tools, likely to offset or minimize the research and development costs of the custom tools. Some of these attacks also used a dropper similar to a Shamoon variant.[27]

Late-2017 through Mid-2018

APT33 conducted a series of attacks against engineering firms to gain access, maintain presence, and farm credentials. Using already-stolen credentials, the group was able to map target networks, modify MS Outlook client homepages, and use commodity tools for privilege escalation and credential theft. Multiple variations of this homepage exploit also debuted during this span. If it were

[25] Interview with a senior cybersecurity consultant.

[26] i.e. registering domains that appear to be owned and operated by Saudi companies and their Western partners, in order to have the sender's address in the phishing email appear plausibly valid.

[27] O'Leary *et al.* 2017 (n. 3).

caught and its presence contained on a network, it would reestablish access and maintain persistence through password spraying.

December 2018

APT33 conducted a supply chain attack against several Middle Eastern companies via their European suppliers, again using the job offer template for the initial spear-phishing effort. It used the Shamoon v.3 wiper as part of a toolkit alongside several other modules, including the Filerase wiper, that made the toolkit extremely destructive. Shamoon v.3 is itself very similar to v.2, which suggests that APT33's goals included testing new wiper modules.[28] Working in tandem, Shamoon v.3 erased the master boot records of an infected computer while Filerase deleted and overwrote files.[29]

February 2019

APT33 targeted a Saudi company in the chemical sector with a spear-phishing attack that again alleged to link to a job description. Once inside the network, APT33 used both custom and commodity malware to exploit a known vulnerability that yielded the ability to archive, compress, and extract files. [30] Symantec ultimately prevented it from doing so.

September 2019 through November 2019

Using its own botnets and private VPNs, APT33 spent several months on a password-spraying campaign that tested common passwords against user accounts at tens of thousands of organizations.[31] From September to November, its focus narrowed to a select number of organizations per month while simultaneously increasing the number of accounts targeted at those organizations. Motivations for the attack are ostensibly related to gaining access to industrial control systems, but researchers believe that these companies are upstream in the supply chain from the actual targets. APT33 could be targeting critical infrastructure organizations that use these control systems.[32]

DISCLOSURE EVENTS

Despite numerous cyber threat intelligence groups reporting on APT33 since 2017, the group has remained undeterred from action. Its TTPs have, however, shifted over time; and with that, the network and host artifacts have also changed. Two likely explanations for this shift are directly opposed. On the one hand, these changes could reflect the overall escalation in the intensity and severity of the threat posed by APT33. That is to

[28] T. Roccia, J. Saavedra-Morales, & C. Beek, 'Shamoon Attackers Employ New Tool Kit to Wipe Infected Systems' (19 December 2018), *McAfee*: https://www.mcafee.com/blogs/other-blogs/mcafee-labs/shamoon-attackers-employ-new-tool-kit-to-wipe-infected-systems/

[29] Critical Attack Discovery and Intelligence Team 2018 (n. 6).

[30] Critical Attack Discovery and Intelligence Team 2019 (n. 6).

[31] T. Seals, 'APT33 Mounts Focused, Highly Targeted Botnet Attacks Against U.S. Victims' (14 November 2019), *Threatpost*: https://threatpost.com/apt33-mounts-targeted-botnet-attacks-us/150248/

[32] Greenberg 2019 (n. 9).

say, the shift could ensue from a natural expansion of its ambitions and the desire to enhance the destructiveness of its operations, suggesting that disclosures fail to alter its behavior. On the other hand, the increasingly aggressive nature of its operations could be a result of the forced obsolescence of TTPs revealed in disclosures. Yet intelligence rarely tolerates such intellectual purity. The reality of the situation is no less plausibly a mix of both theories and other factors heretofore unknown or indiscernible. It should also be noted that the U.S. Department of Justice has not indicted any members of APT33.

The increasingly aggressive nature of its operations could be a result of the forced obsolescence of TTPs revealed in disclosures.

FireEye Disclosure: September 2017

FireEye's September 2017 report on APT33 was the first large disclosure made on the group. The report offered a general timeline of APT33's actions up to the publication date; a list of TTPs, including a description of the fake-job spear-phishing emails; and a description of how its tools were extremely similar to the dropper used in attacks containing Shamoon variants. The disclosure detailed the group's links to Iran and leaked the online handle of one of its members, whom FireEye linked to Iran's Nasr Institute. The report also included an appendix of malware families used by APT33, domain names, and indicators of compromise.[33]

Dragos Non-Public Disclosure: 2018

In 2018 Dragos published a report on APT33 for its paying customers. The report itself was not made public and therefore does not count as a public disclosure, but Dragos did release a summary that contained limited information on the group. The primary value of this précis to our report is its consistency with other reports and disclosures on APT33 from other firms. It explained how APT33's focus started local, mainly focused on Saudi energy and aviation firms, and expanded over time. It also covered APT33's use of spear-phishing as the primary means of gaining initial access, and touched on its seeming interest in industrial control systems—an observation that would be confirmed by the group's password-spraying attacks in 2019.[34]

Symantec Disclosure: December 2018

Symantec was the first firm to report on the December 2018 attack that used the new Shamoon v.3 and the Filerase wiper, and it linked the attack back to APT33. The disclosure did not report on TTPs but did describe tools used, protective measures to take, and indicators of compromise.[35]

McAfee Disclosure: December 2018

The McAfee disclosure on the December 2018 attack came a few days after the Symantec report, and it confirmed the use of

[33] O'Leary *et al.* 2017 (n. 3).

[34] Dragos (n. 2).

[35] Critical Attack Discovery and Intelligence Team 2018 (n. 6).

Shamoon v.3 and the Filerase wiper. The report included information on TTPs, including the fake-job spear-phishing emails, a Quran verse embedded in the code, tools, network and host artifacts, domain names, and indicators of compromise.[36]

FireEye Disclosure: December 2018

FireEye published its own report two days after the McAfee disclosure, one week after the original Symantec disclosure. FireEye was initially less confident that APT33 orchestrated the attack, though an update posted in May 2019 affirmed that APT33 was indeed responsible. The original report registered TTPs, network and host artifacts, domain names, and indicators of compromise—including those of bespoke malware. It also showed how APT33 escalated privileges once on a network and described how defenders could protect their networks.[37]

Symantec Disclosure: March 2019

Symantec published a report on how APT33's February 2019 attack unfolded. The disclosure reviewed its TTPs, such as the spear-phishing emails, dropper codes, and both the commodity and bespoke toolsets used in the attack. Symantec included their own illustration of how APT33 attacks unfold, albeit the old attack from February 2018 as the example. That example included C2 servers, IP addresses, TTPs, and indicators of compromise.[38]

Microsoft Presentation at CYBERWARCON: November 2019

In a presentation at the 2019 CYBERWARCON, an annual conference near Washington, D.C., a researcher from Microsoft described APT33's 2019 password-spraying campaign. The presentation outlined how the attack's focus narrowed throughout the year; and the speaker indicated that it appeared to be targeting industrial control systems. Furthermore, the presenter speculated that the password-spraying attack could be a supply chain attack, with APT33 using it to gain access through lateral movement to other parties linked via networks to the organizations presently under attack. The presenter also suggested that APT33 could be using the attack to gain a position to target critical infrastructure components linked to industrial control systems.[39]

[36] Roccia *et al.* 2018 (n. 28).

[37] Ackerman *et al.* 2018 (n. 16).

[38] Critical Attack Discovery and Intelligence Team 2019 (n. 6).

[39] Greenberg 2019 (n. 9).

APT34 (IRAN)

Timeline: Activity Levels and Disclosures

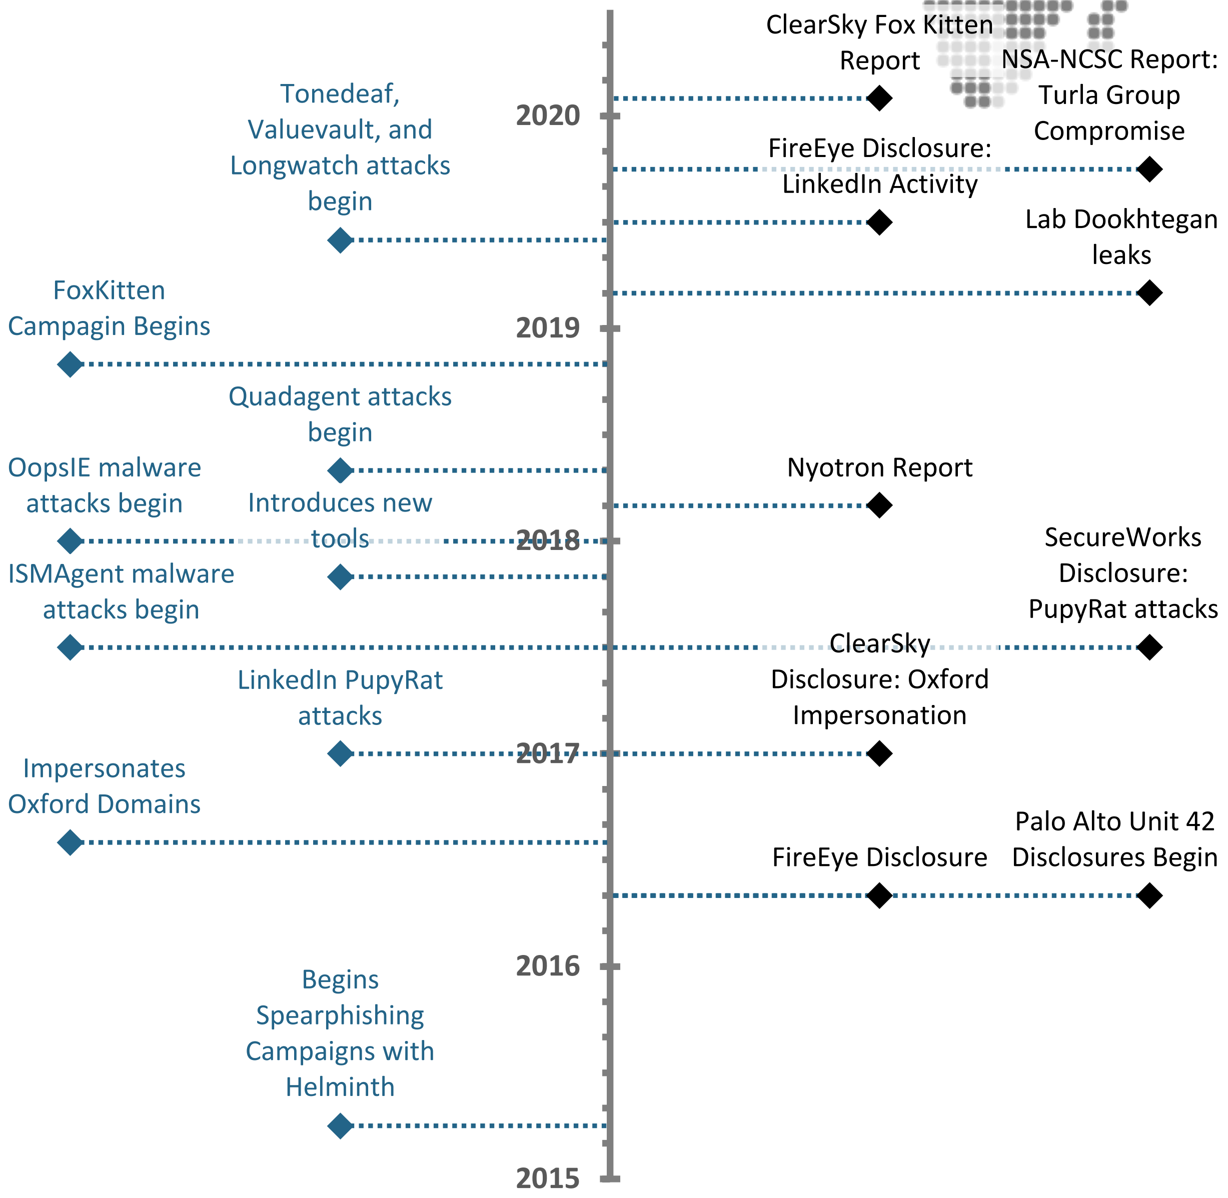

ClearSky Fox Kitten Report

NSA-NCSC Report: Turla Group Compromise

Tonedeaf, Valuevault, and Longwatch attacks begin

FireEye Disclosure: LinkedIn Activity

Lab Dookhtegan leaks

FoxKitten Campagin Begins

Quadagent attacks begin

OopsIE malware attacks begin

Nyotron Report

Introduces new tools

ISMAgent malware attacks begin

SecureWorks Disclosure: PupyRat attacks

ClearSky Disclosure: Oxford Impersonation

LinkedIn PupyRat attacks

Impersonates Oxford Domains

FireEye Disclosure

Palo Alto Unit 42 Disclosures Begin

Begins Spearphishing Campaigns with Helminth

IMPACT OF DISCLOSURES

INTRODUCTION

APT34 is an Iranian cyber threat group known for sustained cyber-espionage activities, which it is believed to do on behalf of the Iranian government. ATP34 consistently targets Middle Eastern organizations that hold substantial interest to Iranian state goals.[1] Despite (or perhaps because of) the many disclosures of its activity, APT34 regularly updates and upgrades its TTPs, and it continues to expand the scope of its operations.

THE GROUP

APT34 (OilRig, Helix Kitten) is believed to have close ties to the Iranian government, specifically to the Iranian Ministry of Intelligence. Concrete evidence of its existence dates no earlier than 2016; however, researchers believe that it has been active since 2014 and its first campaign likely started in late 2015. Since its initial detection, APT34 has remained highly active and invested in updating its toolset.[2] Its operations are classifiable as espionage and typically involve extracting data on the users and processes of target computer systems.[3] It gains access to networks through spear-phishing emails and other types of social engineering campaigns, both of which demonstrate a certain level of research on its targets. In 2019, it ostensibly began to experiment with new delivery vectors for its malware, including the impersonation of IT vendors and the exploitation of unpatched VPN and RDP client vulnerabilities.

APT34 primarily targets organizations in the Middle East, both private businesses and government agencies, that tend to be of certain interest with respect to Iranian state objectives.[4] Consequently, the target countries are often Iran's regional adversaries or competitors—Saudi Arabia, the UAE, Israel, Kuwait, Lebanon, Bahrain, Turkey, Qatar. More recently, APT34 expanded its sights to include infrastructure and companies based in the U.S., Europe, and Australia. Its targets have thus spanned a wide range of sectors, from government, energy, and telecommunications to higher education, hospitality, finance, and aerospace.[5]

[1] M. Sardiwal, V. Cannon, N. Fraser, Y. Longhe, N. Richard, J. O'Leary, 'New Targeted Attack in the Middle East by APT34, a Suspected Iranian Threat Group, Using CVE-2017-11882 Exploit' (7 December 2017), *FireEye*: https://www.fireeye.com/blog/threat-research/2017/12/targeted-attack-in-middle-east-by-apt34.html

[2] B. Lee & R. Falcone, 'OopsIE! Oilrig Uses ThreeDollars to Deliver New Trojan' (23 February 2018), *Palo Alto Networks, Unit42*: https://unit42.paloaltonetworks.com/unit42-oopsie-oilrig-uses-threedollars-deliver-new-trojan/

[3] Sardiwal *et al.* 2017 (n. 1).

[4] ib.

[5] T. Seals, 'OilRig APT Continues its Ongoing Malware Evolution' (13 September 2018), *Threatpost*: https://threatpost.com/oilrig-apt-continues-its-ongoing-malware-evolution/137444/; A. Meyers, 'Meet CrowdStrike's Adversary of the Month for November: HELIX KITTEN' (27 November 2018), *CrowdStrike*:

TIMELINE

The first attestation of APT34 occurred on 22 May 2016. On that day, FireEye issued a warning about an unknown threat actor which had sent emails with malicious attachments to Middle Eastern banks. Since then, cyber threat analysts have come to understand APT34 much better, assisted in no small part by APT34's high level of activity and reuse of assets.[6] It often uses its tooling and infrastructure, for instance, in multiple concurrent campaigns—a telling contrast to other groups, such as APT28 and APT29. Indeed, its consistent use of several tools and the same infrastructure significantly helped identify it as a single group.[7]

2015	First observed conducting spear-phishing against the Saudi defense sector.[8]
2016	**May–Oct**: APT34 undertakes a series of spear-phishing campaigns against Middle East organizations in various sectors using Helminth malware.[9]
2017	**Jan–Feb**: APT34 uses socially engineered LinkedIn accounts to deliver PupyRAT malware against Saudi technology and energy companies.[10] **Jul–Aug**: ISMAgent malware debuts; improved anti-analysis techniques; malware sent to a UAE government organization. **Nov:** APT34 introduces several new tools (e.g. ALMA Communicator, Powruner, Bondupdater, Twoface, RGDoor) to conceal C2 infrastructure better, beat anti-malware tools, and provide alternative backdoors; continues use of spear-phishing, fake job websites, and credential harvesting for access.
2018	**Jan**: APT34 introduces OopsIE malware and a new delivery document.

https://www.crowdstrike.com/blog/meet-crowdstrikes-adversary-of-the-month-for-november-helix-kitten/

[6] Two years after their detection, in February 2018, Palo Alto Networks' Unit 42 was able to assert that 'the OilRig group remains highly active in their attack campaigns while they continue to evolve their toolset' (Lee & Falcone 2018 [n. 2]).

[7] ib.

[8] R. Falcone & B. Lee, 'The OilRig Campaign: Attacks on Saudi Arabian Organizations Deliver Helminth Backdoor' (26 May 2016), *Palo Alto Networks Unit42*: https://unit42.paloaltonetworks.com/the-oilrig-campaign-attacks-on-saudi-arabian-organizations-deliver-helminth-backdoor/

[9] J. Grunzweig & R. Falcone, 'OilRig Malware Campaign Updates Toolset and Expands Targets' (4 October 2016), *Palo Alto Networks Unit42*: https://unit42.paloaltonetworks.com/unit42-oilrig-malware-campaign-updates-toolset-and-expands-targets/

[10] Counter Threat Unit Research Team, 'The Curious Case of Mia Ash: Fake Persona Lures Middle Eastern Targets' (27 July 2017), *Secureworks*: https://www.secureworks.com/research/the-curious-case-of-mia-ash

May–Aug: APT34 adds further anti-analysis features to OopsIE and Bondupdater; Quadagent malware debuts.

Nov: Attacks pivot to the telecommunications sector.

2019

Mar: Lab Dookhtegan leaks multiple tools, infrastructure details, TTPs, and individual personalities associated with APT34 via Telegram.

Jun: APT34 introduces Tonedeaf, Valuevault, and Longwatch; conducts spear-phishing attacks through fake accounts on LinkedIn.

Oct: NSA and NCSC release details on the Turla group and compromise APT34 infrastructure.

2020

Early 2020: APT34 continues use of Tonedeaf and Valuevault against a U.S. company; VPN and RDP vulnerability attacks, possibly dating back to 2017, are exposed.

TYPOLOGY OF ATTACKS

To achieve its cyberespionage objectives, APT34 relies primarily on in-house tools and techniques that it tailors for a specific use or operation. As is customary for these organizations, APT34 makes significant use of spear-phishing and social engineering to deliver malicious attachments with embedded executables.[11] It regularly updates its TTPs and tooling, and further increases their sophistication—particularly in regards to C2 obfuscation.[12] The bespoke nature of these attacks indicates robust target reconnaissance and research. For example, APT34's spear-phishing campaigns have used emails related to the target's internal IT system, including what appeared to be a conversation between two actual employees at the organization.[13] Another attack used a subject line that referred to a topic recently covered in an internal publication from the target.[14] To further its credibility with a target, APT34 has also sent emails from compromised accounts at organizations related to the target or that do business with them. These messages are often related to IT and corporate infrastructure.[15]

APT34 has used other intricate ploys to gain the trust of targets, including socially engineered attacks on social media. It uses these attacks after its more usual phishing efforts have failed: it uses fake social media profiles, ranging from amateur models to university researchers, to lure targets into clicking on a

[11] Meyers 2018 (n. 5).

[12] Sardiwal *et al.* 2017 (n. 1).

[13] S. Singh & Y. H. Chang, 'Targeted Attacks Against Banks in the Middle East' (22 May 2016), *FireEye*: https://www.fireeye.com/blog/threat-research/2016/05/targeted_attacksaga.html

[14] T. Seals, 'Oilrig Sends an OopsIE to Mideast Government Targets' (5 September 2018), *Threatpost*: https://threatpost.com/oilrig-sends-an-oopsie-to-mideast-government-targets/137220/

[15] Meyers 2018 (n. 5).

malicious attachment.[16] The group has also shown increased technical sophistication in its attempts to entice targets, including VPN portal spoofing, IT vendor impersonation, and security certificate theft.[17] In that vein, ClearSky Cyber Security recently identified an extensive campaign to exploit unpatched VPN and RDP services, purportedly extending as far back as 2017 and perhaps in conjunction with other Iranian cyberthreat groups. In this campaign, APT34 gained access to targets through a one-day exploit before companies had patched or by using stolen or faked credentials and certificates. After securing access, it implanted backdoors to maintain network presence, and then moved laterally to find and exfiltrate sensitive information.

The bespoke nature of these attacks indicates robust target reconnaissance and research.

May to October 2016

APT34 conducted a series of targeted spear-phishing campaigns against several financial and tech organizations in Saudi Arabia, a company in Qatar, and government organizations in Turkey, Israel, and the United States. The phishing emails were, however, detected during the initial stages of reconnaissance against these organizations, thereby giving cyber threat firms a significant period of observation. [18] The emails were sent from spoofed accounts and contained an attachment that delivered the Helminth backdoor onto the victim's machine. Throughout this period, APT34 made changes to both the attached document and the malware that lessened the chances that its actions would be detected.[19] During autumn 2016, APT34 attempted to deliver Helminth through a fake VPN portal, to which targets were directed through emails sent from the accounts of compromised IT vendors.[20] Meanwhile, in another effort, it used stolen certificates to impersonate the University of Oxford's domain, on which it hosted a 'job symposium' registration portal that hid malware.[21]

[16] Respectively: T. Spring, 'APT Group Uses Catfish Technique to Ensnare Victims' (27 July 2017), *Threatpost*: https://threatpost.com/apt-group-uses-catfish-technique-to-ensnare-victims/127028/[;] M. Bromiley, N. Klapprodt, N. Fraser, Y. Longhe, N. Schroeder, J. Rocchio, 'Hard Pass: Declining APT34's Invite to Join Their Professional Network' (18 July 2019), *FireEye*: https://www.fireeye.com/blog/threat-research/2019/07/hard-pass-declining-apt34-invite-to-join-their-professional-network.html

[17] ClearSky Research Team, 'Iranian Threat Agent OilRig Delivers Digitally Signed Malware, Impersonates University of Oxford' (5 January 2017), *ClearSky Cyber Security*: https://www.clearskysec.com/oilrig/

[18] Singh & Chang 2016 (n. 13).

[19] ib.; Grunzweig & Falcone 2016 (n. 9).

[20] ClearSky Research Team 2017 (n. 17).

[21] ib. The impersonation of the University of Oxford was particularly sloppy, neither spoofing nor trying to approximate a believable "ox.ac.uk" domain, nor an "@ox.ac.uk" email, nor the University's preferred san-serif font, nor the correct Oxonian blue (#002147). They also failed to realize that the University is comprised of constituent colleges and independent

January to February 2017

APT34 attempted to use social media profiles on LinkedIn to lure specific targets after email phishing attempts had failed. It curated two fake profile, possibly over the course of years, of a young London-based photographer and an amateur model named Mia Ash. These personae were used to target tech-focused employees at large Saudi energy and technology companies, to whom they sent attachments containing the PupyRAT malware.[22]

July to August 2017

APT34 targeted government organizations in the UAE with a Rich Text Format file that contained a new PowerShell-based backdoor—built on the recently released CVE-2017-0199. This document contained a variant of its Clayside script that pushed the new backdoor, called ISMAgent (itself a variant of the ISMDoor trojan), onto the victim's machine. This new backdoor granted additional C2 functionality with a domain generation algorithm called BondUpdater. [23] APT34 later delivered ISMAgent through a new injector, creatively dubbed ISMInjector, which further decreased the conspicuity of the installation and provided additional hedges against anti-virus detection.[24]

November 2017

APT34 introduced several new tools and updated versions of older ones that had been overserved in the wild for some time. Researchers observed the group using Mimikatz in coordination with a new DNS tunneling trojan, the ALMA Communicator, at a public utilities company in the Middle East.[25] This attack marked an enhancement in the accessibility and stealth of its C2 infrastructure, with the novel incorporation of the TwoFace webshell and the RGDoor backdoor.[26]

faculties, and therefore lacks a central hiring authority. Not to mention their dismal command of the English language, which one would imagine to be prerequisite before trying to impersonate the world's foremost English-speaking university. The University would never host a 'job symposium' because (a) that is not the correct use of "symposium", let alone in so formal an academic setting, and (b) the use of the word "job" in this context (as opposed to "career") is distinctly American.

[22] Counter Threat Unit Research Team 2017 (n. 10).

[23] R. Falcone & B Lee, 'OilRig Uses ISMDoor Variant; Possibly Linked to Greenbug Threat Group' (27 July 2017), *Palo Alto Networks Unit42*: https://unit42.paloaltonetworks.com/unit42-oilrig-uses-ismdoor-variant-possibly-linked-greenbug-threat-group/

[24] R. Falcone & B. Lee, 'OilRig Group Steps Up Attacks with New Delivery Documents and New Injector Trojan' (9 October 2017), *Palo Alto Networks Unit42*: https://unit42.paloaltonetworks.com/unit42-oilrig-group-steps-attacks-new-delivery-documents-new-injector-trojan/

[25] R. Falcone, 'OilRig Deploys "ALMA Communicator" – DNS Tunneling Trojan' (8 November 2017), *Palo Alto Networks Unit42*: https://unit42.paloaltonetworks.com/unit42-oilrig-deploys-alma-communicator-dns-tunneling-trojan/

[26] R. Falcone, 'OilRig uses RGDoor IIS Backdoor on Targets in the Middle East' (25 January 2018), *Palo Alto Networks Unit42*:

In late November, APT34 debuted DNSpionage, which it used to deliver Microsoft Office document with malicious macros.[27]

January 2018

APT34 attacked Middle Eastern insurance and financial companies with OopsIE and ThreeDollars, respectively a new malware and a new delivery document.[28]

May to August 2018:

APT34 premiered the Quadagent backdoor. Using stolen credentials from trusted organizations within a specific nation-state, it conducted spear-phishing attacks against a government entity and a tech service provider in the same country.[29] APT34 continued to add anti-analysis features to its OopsIE and Bondupdater tools.[30,31]

November 2018

In early November, CrowdStrike observed that APT34 was targeting a specific customer in the telecommunication vertical. While this represented a shift in targeting, APT34 utilized the same tools and TTPs.[32]

June 2019

APT34 impersonated researchers at the University of Cambridge on LinkedIn and offered fake job opportunities to certain individuals. It sent malicious messages containing the new Tonedeaf, Valuevault, and Longwatch malware.[33]

January to February 2020

APT34 used an updated version of Tonedeaf and Valuevault against the U.S. company Westat.[34] Together with APT33 and APT39, it launched a campaign to exploit unpatched VPN and RDP vulnerabilities in order to gain

https://unit42.paloaltonetworks.com/unit42-oilrig-uses-rgdoor-iis-backdoor-targets-middle-east/

[27] W. Mercer & P. Rascagneres, 'DNSpionage Campaign Targets Middle East' (27 November 2018), *Talos Intelligence*: https://blog.talosintelligence.com/2018/11/dnspionage-campaign-targets-middle-east.html

[28] Lee & Falcone 2018 (n. 2).

[29] B. Lee & R. Falcone, 'OilRig Targets Technology Service Provider and Government Agency with QUADAGENT' (25 July 2018), *Palo Alto Networks Unit42*: https://unit42.paloaltonetworks.com/unit42-oilrig-targets-technology-service-provider-government-agency-quadagent/

[30] R. Falcone, B. Lee, R. Porter, 'OilRig Targets a Middle Eastern Government and Adds Evasion Techniques to OopsIE' (4 September 2018), *Palo Alto Networks Unit42*: https://unit42.paloaltonetworks.com/unit42-oilrig-targets-middle-eastern-government-adds-evasion-techniques-oopsie/

[31] K. Wilhoit & R. Falcone, 'OilRig Uses Updated BONDUPDATER to Target Middle Eastern Government' (12 September 2018), *Palo Alto Networks Unit42*: https://unit42.paloaltonetworks.com/unit42-oilrig-uses-updated-bondupdater-target-middle-eastern-government/

[32] Meyers 2018 (n. 5).

[33] Bromiley *et al*. 2019 (n. 16).

[34] P. Litvak & M. Kajiloti, 'New Iranian Campaign Tailored to U.S. Companies Uses an Updated Toolset' (30 January 2020), *Intezer*: https://intezer.com/blog/apt/new-iranian-campaign-tailored-to-us-companies-uses-updated-toolset/

access to target networks in Israel and some Gulf states. This activity potentially extended back to 2017, and included a few previously unreported TTPs and malware that affected dozens of companies around the world.

DISCLOSURE EVENTS

Since 2016 there have been no less than twenty-five public disclosures, leaks, or reports on APT34's TTPs, tools, and infrastructure. Information disclosures appear to regulate APT34's behavior, insofar as it frequently updates its malware and TTPs and enhances the sophistication of its C2 infrastructure and anti-virus detection capabilities. Nevertheless, available evidence fails to suggest that disclosures disrupt its activity. Its TTPs often remain unchanged from one attack to the next, regardless of whether a disclosure occurred. One could therefore presume that, in the eyes of APT34, disclosures do not particularly jeopardize the efficacy of its tools and tactics. Hence the lack of correlation between the dates of disclosures and the longevity of its operations.

In the eyes of APT34, disclosures do not particularly jeopardize the efficacy of its tools and tactics.

Several reasons explain this lack of impact. For one, APT34 is increasing its vectors of attack. Although phishing emails still appear *de rigueur* in the opening phases of a campaign, APT34 has demonstrated the ability to capitalize on one-day exploits, and its credential-harvesting operations now obtain throughout the Middle East. Moreover, APT34 has increasingly broadened the scope of its targets in terms of both region and sector: techniques well known in one area may be entirely novel in another. Perhaps most important is its proven capacity for updating and upgrading its malware during the middle of an operation. A disclosure poses no threat if its information can be quickly rendered outdated or otherwise obviated through planned patching. Accordingly, the only disclosure to disrupt its activity was the Lab Dookhtegan leak, which both divulged an extensive amount of information and directly compromised individual members of APT34.[35]

FireEye Disclosure: May 2016

FireEye's report on a series of spear-phishing attacks against certain Middle Eastern banks marked the first disclosure of APT34's activity. Although the culprit remained unknown at the time, FireEye later attributed the activity to APT34 in December 2017. The initial report provided the basis for most subsequent reporting on its activity. The report identified the highly targeted nature of its spear-phishing emails and the primary delivery method as macro-enabled XLS files. The report also detailed the contents of the emails and XLS files, and analyzed what

[35] Thus FireEye's Ben Read, in an interview on 15 April 2020.

would become known as the Helminth malware.[36]

Palo Alto Unit 42 Disclosures: May and October 2016

Palo Alto Networks' Unit 42 issued two reports on activity similar to what was mentioned in the FireEye report in May 2016. The reports named the campaign OilRig, the malware Helminth, and the document Clayside. Unit 42 connected this activity to incidents involving the Saudi defense industry in 2015. Across the two reports, Unit 42 canvassed APT34's TTPs, analyzed the malware and its variants, the indicators of compromise, and Helminth's C2 infrastructure. Of note was Unit 42's discovery that variants of Helminth dropped files named fireeye.vbs and fireeye.ps1, suggesting that the authors of these updates had read FireEye's report.[37]

Variants of Helminth dropped files named fireeye.vbs and fireeye.ps1, suggesting that the authors of these updates had read FireEye's report.

ClearSky Cyber Security Disclosure: January 2017

ClearSky found and published the first examples of APT34 using fake VPN portals, fake websites, and stolen security certificates. ClearSky's report detailed how these websites downloaded the Helminth malware onto the target computers through both the VPN software and the malicious webpages. The downloads were signed with a valid but stolen security certificate from a legitimate software company. The report also discussed the four domains that APT34 used to impersonate the University of Oxford.[38]

SecureWorks Counter Threat Unit Disclosure: July 2017

This report highlighted APT34's use of fake social media accounts to conduct spear phishing. SecureWorks identified the TTPs employed by APT34 to lure targets with the account and deliver the PupyRAT malware. The report described the primary social media account of a fake persona known as Mia Ash and attributed the profile to APT34.

Nyotron Attack Response Center Report: March 2018

The Nyotron report was the first systematic catalogue of APT34's activity up to that point. It registered the TTPs and tools used to bypass defenses, establish persistence, escalate privileges, conduct internal reconnaissance, and move laterally. Notably, the report provided technical details of three previously undocumented C2 and data exfiltration capabilities.[39]

[36] Singh & Chang 2016 (n. 13).

[37] Falcone & Lee 2016 (n. 8).

[38] ClearSky Research Team 2017 (n. 17).

[39] Nyotron Attack Response Center, 'OilRig is Back with Next-Generation Tools and Techniques' (March 2018), *Nyotron*: https://nyotron.com/wpcontent/uploads/2018/03/Nyotron-OilRig-Malware-Report-March-2018b.pdf

Lab Dookhtegan Leak: March 2019

An anonymous entity dumped a trove of data that allegedly originated from APT34's operations. In sum, the disclosure included a list of 13,000 stolen credentials and compromised systems, one-hundred different web shell-launching URLs, various backdoors and DNS hijacking tools, screenshots of their operational platforms, sundry details on their C2 infrastructure, and details about specific individuals associated with the Iranian ministry of Intelligence. The leaked information suggested that APT34 had conducted operations against ninety-seven organizations across twenty-seven countries.[40]

FireEye Disclosure: July 2019

FireEye identified activity similar to what was found in the SecureWorks report on APT34's LinkedIn operations. FireEye's report covered APT34's impersonation of research staff at the University of Cambridge and its three new malware tools. The reported activity occurred in June, only a few months after the Lab Dookhtegan Leaks. FireEye also provided the expected analysis of the new malware and new indicators of compromise identified activity similar to what was found in the SecureWorks report on APT34's LinkedIn operations. FireEye's report covered APT34's impersonation of research staff at the University of Cambridge and its three new malware tools. The reported activity occurred in June, only a few months after the Lab Dookhtegan Leaks. FireEye also provided the expected analysis of the new malware and new indicators of compromise.[41]

NSA-NCSC Report: October 2019

This joint report between NSA and NCSC detailed the Turla Group's apparent compromise of APT34's C2 infrastructure. It provided two indicators of compromise, although whether the indicators signify activity from APT34 or Turla remains unclear.[42]

Intezer Disclosure: January 2020

Intezer reported on APT34's phishing campaign against the U.S. company Westat. The campaign used a bogus employee satisfaction survey embedded with malicious code, which in turn downloaded the version of the ToneDeaf malware treated by FireEye's report from July 2019. Intezer's document identified TTPs, C2 infrastructure, specific updates to the malware, and indicators of compromise. [43]

ClearSky Fox Kitten Report: February 2020

Utilizing undisclosed information from the cybersecurity firm Dragos, ClearSky issued a report on APT34's extensive use of

[40] B. Lee & R. Falcone, 'Behind the Scenes with Oilrig' (30 April 2019), *Palo Alto Networks Unit42*: https://unit42.paloaltonetworks.com/behind-the-scenes-with-oilrig/

[41] Bromiley *et al.* 2019 (n. 16).

[42] U.S. National Security Agency, U.K. National Cybersecurity Centre, 'Cybersecurity Advisory: Turla Group Exploits Iranian APT to Expand Coverage of Victims' (21 October 2019), https://media.defense.gov/2019/Oct/18/2002197242/-1/-1/0/NSA_CSA_TURLA_20191021%20VER%203%20-%20COPY.PDF

[43] Litvak & Kajiloti 2020 (n. 34).

unpatched VPN and RDP software to infiltrate target networks. They analyzed the tools and TTPs used in the attacks, with specific attention to how APT34 gained access, escalated privileges, moved laterally, and exfiltrated data. The report attributed the activity to APT34 but remained open to the possibility that APT33 or APT39 was responsible. It included indicators of compromise, hashes, and associated IP addresses.[44]

[44] ClearSky Research Team, 'Fox Kitten – Widespread Iranian Espionage-Offensive Campaign' (16 February 2020), *ClearSky Cyber Security*: https://www.clearskysec.com/fox-kitten/

APT38 (NORTH KOREA)

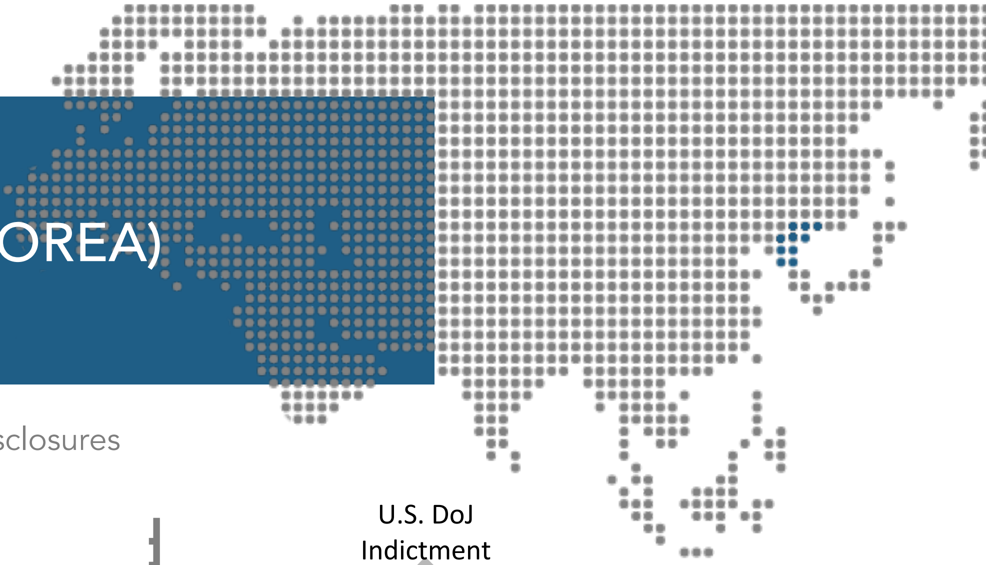

Timeline: Activity Levels and Disclosures

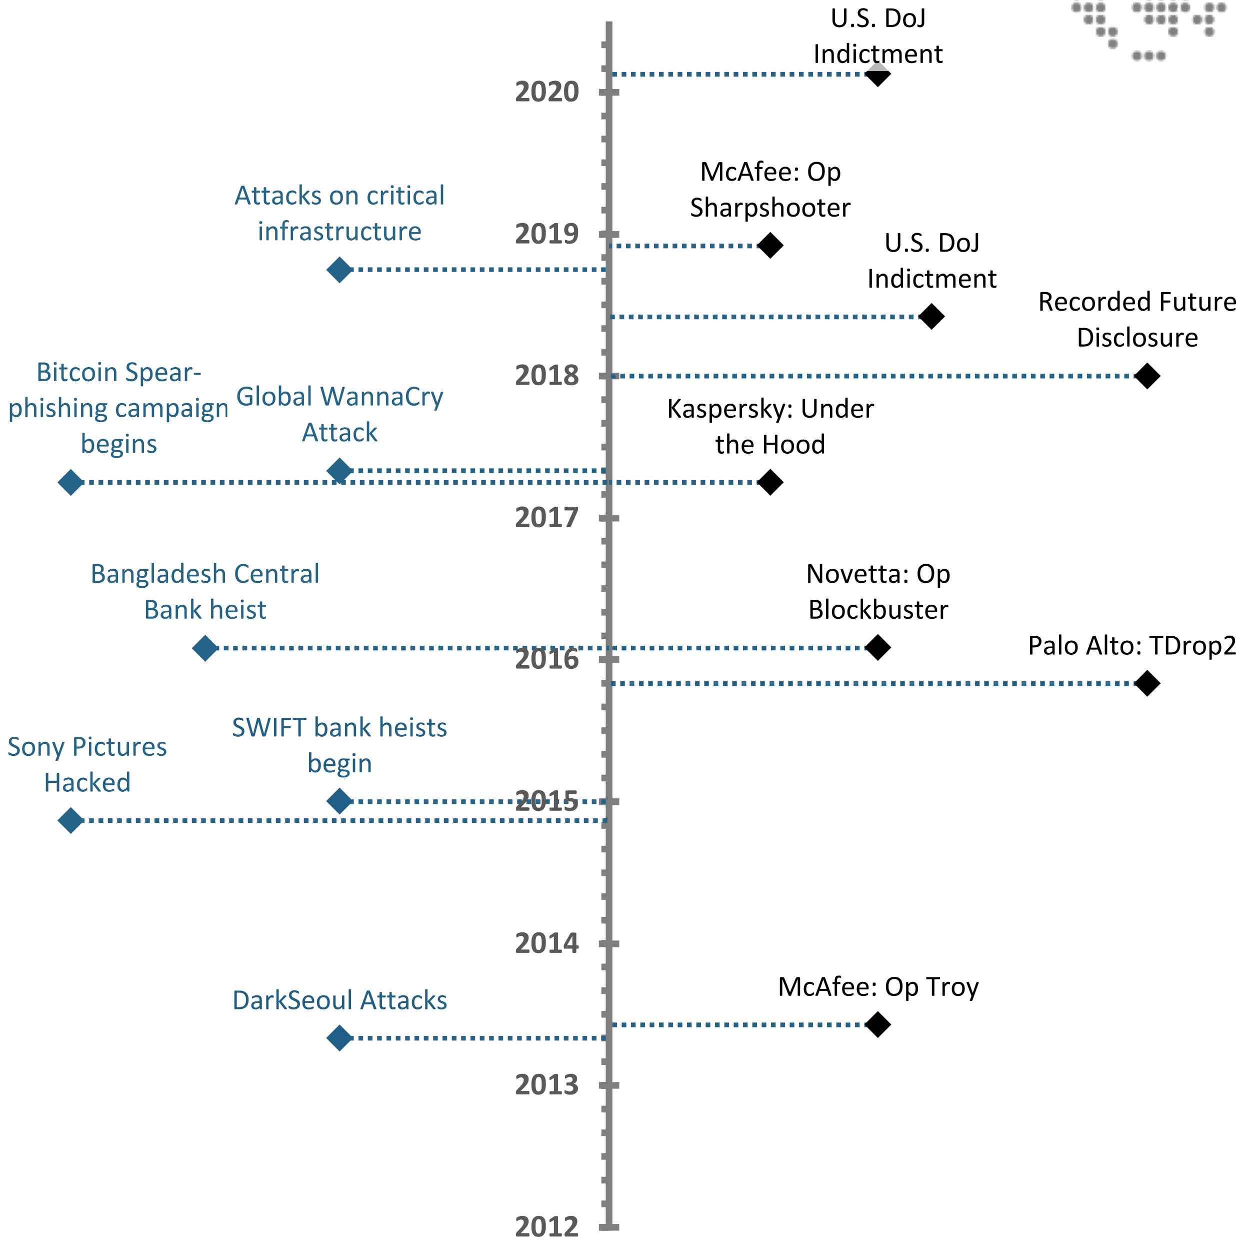

IMPACT OF DISCLOSURES

INTRODUCTION

Lazarus Group's imperviousness to disclosures has several likely sources. It's concealed by the secretiveness of North Korean society and, as part of the military structure, is unconcerned with foreign criminal investigations. It has also been primarily financially motivated since the 2013 United Nations sanctions on North Korea and its 2014 Sony Pictures hack. This new motivation has allowed Lazarus to diversify its target base: first attacking national banks then pivoting to cryptocurrency exchanges in 2017 when Bitcoin prices rose. As cryptocurrency prices settled, it began conducting FASTCash operations in a variety of locations from South America to Africa. In April 2020, the U.S. government cautioned that the group had started providing its cybercrime services to international clients—a move which could confound the attribution process further. Disclosures have been unable to change the decision-calculus of a regime with no alternative streams for revenue and little to lose.

THE GROUP

APT38 (Hidden Cobra, Zinc, Stardust Chollima, Guardians of Peace, WhoIs Group),[1] better known as the Lazarus Group, is a threat actor run by the North Korean state intelligence agency, the Reconnaissance General Bureau.[2] Lazarus' first campaign launched in 2009, first against the United States and then against South Korea.[3] Since then, it has increased its international notoriety with several significant incidents, such as its 2014 breach of Sony Pictures' networks;[4] its 2016 heist at Bangladesh's Central Bank through the SWIFT system;[5] and its 2017 release of WannaCry 2.0 into the world.[6] Although Lazarus' activities have drifted from espionage to disruption, its primary activity appears to be financial cybercrime—no doubt to help the North Korean regime weather an array of United Nations and U.S. sanctions. South Korean law enforcement has estimated that North Korea had a force of at least ten-thousand trained hackers as far back as 2011.[7]

[1] 'Lazarus Group', *Malpedia:* https://malpedia.caad.fkie.fraunhofer.de/actor/lazarus_group ; R. Sherstobitoff, I. Liba, & J. Walter, 'Dissecting Operation Troy: Cyberespionage in South Korea' (June 2013), *McAfee:* https://www.mcafee.com/enterprise/en-us/assets/white-papers/wp-dissecting-operation-troy.pdf

[2] 'United States of America v. Park Jin Hyok' (8 June 2018), *United States District Court, Central District of California:* https://www.justice.gov/opa/press-release/file/1092091/download. Hereafter referred to as DoJ Indictment.

[3] A. L. Johnson, 'Born on the 4th of July' (9 July 2009), *Symantec*: From https://community.broadcom.com/symantecenterprise/communities/community-home/librarydocuments/viewdocument?DocumentKey=d5fc6afb-02e8-423f-8feb-f77c68ec7c8a&CommunityKey=1ecf5f55-9545-44d6-b0f44e4a7f5f5e68&tab=librarydocuments

[4] Novetta 2016 (n. 1), 6.

[5] DoJ Indictment (n. 2), 56.

[6] id., 106.

[7] D. Gewirtz, 'Inside the Early Days of North Korea's Cyberwar Factory' (12 June 2018),

Lazarus Group hides behind the isolation of North Korea to present itself in different ways. It conducted Operation Troy from 2009–2013 as the WhoIs Group and NewRomanic Cyber Army Team. For its breach of Sony Pictures, it self-identified as a hacktivist group called the Guardians of Peace. In 2016, the data analytic company Novetta labelled it 'Lazarus Group', a term which has come to encompass a sizable amount of North Korean cybercrime operations. An indictment from the U.S. Department of Justice in 2018 identified an individual and several companies that Lazarus used as fronts for its operations, again suggesting that the relative obscurity of North Korea plays to its advantage outside the Korean Peninsula.[8]

There is little that the U.S. can do to deter the Group because the North Korean regime has few alternative streams for revenue.[9] Indictments from the U.S. government and its allies may send a signal that they consider such behavior a threat; however, these actions will fail to have an impact if the regime lacks something meaningful (or anything) to lose, or if the intended source of coercion does not impose a significant enough cost. The group operates akin to a criminal group, and immune to the fear of extradition or sanctions on individuals, there is little reason that public disclosures or even indictments will impact its operations.[10]

Immune to the fear of extradition or sanctions on individuals, there is little reason that public disclosures or even indictments will impact its operations.

Over the past decade, Lazarus' operations have expanded to such an extent that it now touts two highly successful subgroups, Andariel and Bluenoroff.[11] The former refers to the outfit that targets the South Korean government and its associated organizations: the latter focuses on global espionage and the monetization of cyber capabilities for the North Korean regime. This recategorization of its threat activity supervened on the identification of similar code in the malware used during these divergent operations. According to a senior security researcher with Kaspersky, '[Lazarus Group] adapt their toolkit to match, and then dispose and keep going'.[12]

ZDNet: https://www.zdnet.com/article/inside-the-early-days-of-north-koreas-cyberwar-factory/ (originally published in *Counterterrorism Magazine* 2012).

[8] DoJ Indictment (n. 2).

[9] Thus Jenny Jun, Non-Resident Fellow at the Atlantic Council Cyber Statecraft Initiative, and lead author of the 2018 CSIS report *North Korea's Cyber Operations*.

[10] Dmitri Alperovitch Interview.

[11] 'A Look into the Lazarus Group's Operations' (24 January 2018), *Trend Micro*: https://www.trendmicro.com/vinfo/us/security/news/cyber-attacks/a-look-into-the-lazarus-groups-operations

[12] K. Zetter, 'The Sony Hackers Were Causing Mayhem Years Before They Hit the Company' (24 February 2016), *WIRED Magazine*: https://www.wired.com/2016/02/sony-hackers-causing-mayhem-years-hit-company/

Most recently, in April 2020, the U.S. government suggested that Lazarus had begun to offer hacker-for-hire services to nation-state and criminal cyber groups.[13] In a New York Times article discussing the alert, FireEye's Director of Intelligence Analysis John Hultquist stated, 'we never knew that, and what it shows is the level to which North Korean hackers are maximizing their cyber capabilities'.[14]

TIMELINE

Researchers first observed Lazarus Group in 2009 but believe it formed two years prior.[15] Its first known campaign began on 4 July 2009 with a series of DDoS attacks, first against the United States and then South Korea later that week[16]. A hiatus ensued immediately thereafter and lasted through 2010.[17] Lazarus reappeared in 2011 with a DDoS and espionage campaign—dubbed the "Ten Days of Rain"—which occurred in March and May.[18] The group eluded serious public scrutiny until 2013, when McAfee's report *Dissecting Operation Troy* examined the DarkSeoul malware that ravaged systems at several South Korean targets in 2013.[19] The report expressed the belief that Lazarus had been present in those systems as far back as October 2009.[20]

Lazarus remained relatively enigmatic in this period before the Sony Pictures hack. It donned the mantle of hacktivists to evade, confuse, or otherwise misdirect its victims;[21] that tactic would again feature prominently in their November 2014 campaign against Sony Pictures, where it took responsibility for the defacement and exfiltration of confidential data under the guise of the "Guardians of Peace". An FBI investigation into the incident quickly identified North Korea as the culprit,[22] but several years and several more high-profile incidents passed before the U.S. government issued an indictment.

[13] U.S. CERT 2020 (n. 1).

[14] D. E. Sanger & N. Perlroth, 'U.S. Accuses North Korea of Cyberattacks, a Sign That Deterrence Is Failing' (15 April 2020), *The New York Times:* https://www.nytimes.com/2020/04/15/world/asia/north-korea-cyber.html

[15] Novetta 2016 (n. 1), 21.

[16] Johnson 2009 (n. 3).

[17] J. A. Guerrero-Saade & C. Raiu, 'Operation Blockbuster Revealed' (24 February 2014), *Kaspersky:* https://securelist.com/operation-blockbuster-revealed/73914/

[18] 'Ten Days of Rain' (July 2011), *McAfee*: https://www.mcafee.com/wp-content/uploads/2011/07/McAfee-Labs-10-Days-of-Rain-July-2011.pdf

[19] R. Sherstobitoff, I. Liba, J, Walter, 'Dissecting Operation Troy: Cyberespionage in South Korea' (June 2013), *McAfee*: https://www.mcafee.com/enterprise/en-us/assets/white-papers/wp-dissecting-operation-troy.pdf

[20] ib., 17.

[21] 'A Look into the Lazarus Group's Operations' (24 January 2018), *Trend Micro*: https://www.trendmicro.com/vinfo/us/security/news/cyber-attacks/a-look-into-the-lazarus-groups-operations

[22] 'Update on Sony Investigation' (19 December 2014), *Federal Bureau of Investigation*: https://www.fbi.gov/news/pressrel/press-releases/update-on-sony-investigation

After 2014, Lazarus pivoted from political campaigns to cybercrime. From January 2015 through October 2017, the group conducted successful bank heists in Ecuador,[23] the Philippines, Taiwan, and most memorably Bangladesh—where SWIFT form spelling errors reduced the intended haul of $1 billion $81 million.[24] Two exceptions came when the group failed to rob TPBank in Vietnam in late 2015 and an anonymous bank in South Asia in mid-2016.[25]

Despite the revelations from Operation Blockbuster *and the publicity of the Bangladesh Central Bank heist, Lazarus' operations against banks in South East Asia and Europe persisted and even grew in scope.*

Lazarus received attention in February 2016, the very month that the Bangladesh Central Bank was hit, when Novetta published *Operation Blockbuster*. This report pieced together the clues from Operation Troy and the Sony Pictures breach to expose and name the threat group as the Lazarus Group.[26] The disclosure revealed details on malware used, TTPs employed, and past incidents that could be attributed to the group.

Despite the revelations from *Operation Blockbuster* and the publicity of the Bangladesh Central Bank heist, Lazarus' operations against banks in South East Asia and Europe persisted and even grew in scope to include cryptocurrency exchanges. [27] In April 2017, Kaspersky released *Lazarus Under the Hood*, [28] which revealed more information on a finance-focused Lazarus subgroup that it called Bluenoroff. [29] That report revealed Bluenoroff's TTPs and tools, and by extension those of Lazarus.

Until this point, Lazarus' activities had been target-specific. That changed with WannaCry 2.0 in May 2017.[30] This ransomware swept across the world and only

[23] D. Barrett, & K. Burne, 'Now It's Three: Ecuador Bank Hacked via Swift' (19 May 2016), *Wall Street Journal*: https://www.wsj.com/articles/lawsuit-claims-another-global-banking-hack-1463695820

[24] J. Pagliery, 'Global Banking System Under Attack - What You need to know' (28 May 2016), *CNN*: https://money.cnn.com/2016/05/27/technology/
swift-bank-hack/index.html

[25] See, respectively: 'Vietnamese Bank Foils $1m Cyber Heist' (15 May 2016), *The Guardian*: https://www.theguardian.com/technology/2016/may/16/vietnamese-bank-foils-1m-cyber-heist; 'Lazarus Under The Hood' (3 April 2018), *Kaspersky*, 4–13: https://media.kasperskycontenthub.com/wp-content/uploads/sites/43/2018/03/07180244/Lazarus_Under_The_Hood_PDF_final.pdf

[26] Novetta 2016 (n. 1), 13.

[27] J. A. Guerrero-Saade, & P. Moriuchi, 'North Korea Targeted South Korean Cryptocurrency Users and Exchange in Late 2017 Campaign' (16 January 2018), *Recorded Future*: https://go.recordedfuture.com/
hubfs/reports/cta-2018-0116.pdf

[28] Kaspersky 2018 (n. 25).

[29] id., 3.

[30] A. L. Johnson, 'WannaCry: Ransomware attacks show strong links to Lazarus group' (22 May 2017), *Symantec*: https://community.broadcom.com/symantecenterprise/communities/community-

stopped when researchers found and activated a kill switch.

In June 2018, more than a year after WannaCry, the U.S. Department of Justice indicted a North Korean individual, Park Jin Hyok.[31] The Department of Justice identified Hyok by his travel to China,[32] his background in computer programming,[33] and his connection to Korea Expo Joint Venture—an alleged front for the North Korean hacking cell Lab 110.[34] The indictment confirmed the link between Lazarus and the Reconnaissance General Bureau as well as the group's involvement in the Sony Pictures hack, the Bangladesh Central Bank heist, and WannaCry. The timing of the indictment was opportune: Lazarus' new campaign against cryptocurrency exchanges, Operation Sharpshooter, was well underway. [35] Nevertheless, this campaign continued undeterred for several months after the indictment.

By the spring of 2019, Lazarus appeared to have both doubled down on ransomware attacks and expanded its targets. Researchers found the group targeting Russian firms[36] and exploring backdoors into MacOS.[37] It also ran several ransomware campaigns related to the Ryuk malware.[38] That same year, the U.S. government imposed sanctions against the group and its North Korean affiliates,[39] along with a pair of Chinese allies in 2020.[40] Shortly thereafter, several U.S. government agencies released a joint statement that Lazarus Group is offering its cybercrime services for hire across the world.[41]

home/librarydocuments/viewdocument?DocumentKey=b2b00f1b-e553-47df-920d-f79281a80269&CommunityKey=1ecf5f55-9545-44d6-b0f4-4e4a7f5f5e68&tab=librarydocuments

[31] DoJ Indictment (n. 2).

[32] id., 142.

[33] FireEye 2018 (n. 1), 27.

[34] DoJ Indictment (n. 2), 5.

[35] McAfee Labs and Advanced Threat Research, 'Operation Sharpshooter' (12 December 2018), *McAfee*: https://www.mcafee.com/enterprise/en-us/assets/reports/rp-operation-sharpshooter.pdf

[36] 'North Korea Turns Against New Targets?!' (19 February 2019), *Check Point*: https://research.checkpoint.com/2019/north-korea-turns-against-russian-targets/

[37] 'Operation AppleJeus: Lazarus Hits Cryptocurrency Exchange with Fake Installer and MacOS Malware' (23 August 2018), *Kaspersky*: https://securelist.com/operation-applejeus/87553/

[38] A. Hanel, 'Big Game Hunting with Ryuk: Another Lucrative Targeted Ransomware' (10 January 2019), *CrowdStrike*: https://www.crowdstrike.com/blog/big-game-hunting-with-ryuk-another-lucrative-targeted-ransomware/

[39] 'Treasury Sanctions North Korean State-Sponsored Malicious Cyber Groups' (13 September 2019), *U.S. Department of the Treasury*: https://home.treasury.gov/index.php/news/press-releases/sm774

[40] 'Treasury Sanctions Individuals Laundering Cryptocurrency for Lazarus Group' (2 March 2020) *U.S. Department of the Treasury*: https://home.treasury.gov /news/press-releases/sm924

[41] U.S. CERT 2020 (n. 1).

2007 Speculative formation of Lazarus Group.

2009 **Jul:** DDoS attacks against U.S. government and South Korean targets.

Oct 2009–Mar 2013: DDoS and espionage campaigns against South Korean targets.

2011 **Jul**: McAfee reports on Ten Days of Rain, outlining network artifacts.

2013 **Jun:** McAfee reports on Operation Troy, detailing TTPs and tools used in what is considered a years-long operation.

2014 **Nov**: Hacktivist group "Guardians of Peace" reveals the massive breach of Sony Pictures.

2015 **Jan**: $12 million stolen from Banco del Austro in Ecuador; incident remains confidential until a case filed against Wells Fargo in 2016.

Oct: $1 million stolen from an unnamed Philippine bank; details remain confidential.

Nov: Palo Alto Networks notes the return of a malware previously used in the 2013 DarkSeoul attacks.

Dec: Vietnam's TPBank successfully prevents a heist; details remain confidential.

2016 **Feb**: $81 million stolen from Bangladesh Central Bank through the SWIFT system; Novetta releases *Operation Blockbuster*, which christens the group "Lazarus" and details TTPs, malware, network artifacts, and other technical specifics.

May: Lazarus attempts a heist on another Vietnamese bank.

Aug: Lazarus targets an undisclosed bank in South East Asia.

2017 **Jan**: Lazarus targets Polish and other European Banks.

Apr: Lazarus commences a sophisticated spear-phishing campaign against Bitcoin exchanges; Kaspersky publishes *Lazarus Under the Hood*, which identifies the Bluenoroff subgroup and details the TTPs and malware employed in bank heists.

May: Lazarus releases WannaCry 2.0.

Oct: Lazarus steals $60 million from Far Eastern International Bank.

2018 **Jan**: Recorded Future reports on the spear-phishing campaign against cryptocurrency exchanges

	from April–October 2017, detailing network artifacts and TTPs. **Feb**: McAfee releases its own report on the cryptocurrency exchange spear-phishing campaign. **Jun**: U.S. DoJ indicts Park Jin Hyok, a member of Lazarus Group, for the Sony Pictures Hack, Bangladesh Central Bank Heist, and WannaCry 2.0. **Aug**: Kaspersky reports on Operation AppleJeus and provides details on the MacOS malware employed and other tools. **Oct–Nov**: Lazarus targets critical infrastructure and global defense across the world. **Nov**: FireEye promotes TEMP.Hermit to APT38. **Dec**: McAfee reports on Operation Sharpshooter, which occurred from October to November of that year.
2019	**Feb**: Check Point reports on Lazarus' apparent targeting of Russian firms.
2020	**Mar**: U.S. DoJ indicts two Chinese nationals for assisting North Korea in cryptocurrency scheme.

TYPOLOGY OF ATTACKS

The Lazarus Group typically begins its operations with a spear-phishing campaign to gain initial access to victim systems; through a combination of spear-phishing, decoy documents, and watering holes, the group usually succeeds in ensnaring a target. This pattern of behavior persists in multiple forms from Operation Troy in 2013 to the campaigns against cryptocurrency exchanges in 2018. The group conducts extensive research on its targets and ultimately becomes capable of passive persistence within a system for extended periods of times, such as in bank heists where it was present for eight months undetected.[42]

Knowing that its actions will be investigated, Lazarus puts significant effort into obfuscation and misleading investigators. This tendency was seen in 2013, when two separate groups claimed ownership of Operation Troy;[43] in 2014, when a fake hacktivist group also claimed responsibility for the Sony Pictures hack;[44] and in 2018, when Lazarus employed false-flag tactics to appear like Russia and China.[45]

2009-2013

As mentioned above Lazarus has used the same social engineering methodology since its first foray into espionage and disruption, namely 2013's Operation Troy. It spear-phished a victim within its target

[42] Kaspersky 2018 (n. 25), 8.

[43] Sherstobitoff *et al.* 2013 (n. 19), 4.

[44] Novetta 2016 (n. 1), 12.

[45] Check Point 2019 (n. 36).

organizations with a RAT,[46] though it remains unclear how far in advance of the operation this occurred—a matter of several weeks to months. In contrast to the DarkSeoul wiper, which was compiled in advance, the dropper used in this attack was compiled on the day of the attack.[47] After the target was infected with the RAT, the malware modified the registry property to allow remote connections. Lazarus then launched the wiper after its espionage objectives were completed or after the zombie machines had contributed to a DDoS campaign.[48]

Kaspersky noted an operational pause in the its activity between 2012 and 2013.[49] Geopolitical or structural changes within the North Korean regime perhaps best explain this gap: a disruption resulting from an information disclosure seems unlikely given that Operation Troy was well underway by 2013.

Knowing that its actions will be investigated, Lazarus puts significant effort into obfuscation and misleading investigators.

2014-2016

Novetta's *Operation Blockbuster*, published February 2014, identified forty-five different families of malware used by Lazarus, including wipers, uninstallers, spreaders, RAT, proxy, loaders, keyloggers, installers, and HTTP and DDoS servers.[50] The group nonetheless seemed to forego a retooling effort. Instead, it may be deduced that Lazarus sought to improve previous iterations of its malware. Reports from Kaspersky demonstrated the links between the malware used in the Sony Pictures hack (Destover) and the DarkSeoul malware,[51] while Palo Alto Networks similarly demonstrated the connections between the new TDrop malware and the dropper used in 2013's Operation Troy.[52] Despite the disclosure of TTPs and tools in *Operation Blockbuster*—not to mention US-CERT's various alerts regarding the msoutc.exe malware used in heists of TPBank in Vietnam and the Bangladesh Central Bank—BAE Systems was still able to opine some two years later that that Lazarus' malware exhibited 'the same unique characteristics,'[53] suggesting that *Operation Blockbuster* and the increased scrutiny from U.S. law enforcement did not cause Lazarus to retool.

[46] Sherstobitoff *et al.* 2013 (n. 19), 6.

[47] id., 3.

[48] id., 7.

[49] Guerrero-Saade & Raiu 2014 (n. 17).

[50] Novetta 2016, (n. 1), 24–7.

[51] K. Baumgartner, 'Sony/Destover: Mystery North Korean Actor's Destructive and Past Network Activity' (4 December 2014), *Kaspersky*: https://securelist.com/destover/67985/

[52] B. Lee & J. Grunzweig, 'TDrop2 Attacks Suggest Dark Seoul Attackers Return' (18 November 2015), *Palo Alto Networks Unit 42:* https://unit42.paloaltonetworks.com/tdrop2-attacks-suggest-dark-seoul-attackers-return/

[53] S. Shevchenko, & A. Nish, 'Cyber Heist Attribution' (16 May 2016). *BAE Systems:* https://baesystemsai.blogspot.com/2016/05/cyber-heist-attribution.html

2017-2019

Shortly before Lazarus Group released WannaCry 2.0, Kaspersky published *Lazarus Under the Hood*, which examined the tools used Lazarus' bank heists. In one observed incident, Lazarus infiltrated a system through an outdated version of Adobe Flash Player.[54] In another, Lazarus launched a watering hole attack on a government website to gain entry into European banks.[55] Its spear-phishing campaign against cryptocurrency exchanges in 2018 continued to employ established TTPs from 2016. The lures used in the cryptocurrency campaign targeted Korean-language users and exploited a vulnerability in the local Hangul Word Processor.[56]

During this interval, Lazarus' malware families began to overlap with each other. This development suggests that its developers have an extensive codebase from which they can cut-and-splice malware into new iterations. Consequently, identification and grouping of these malware types is only possible when their respective codes are compared.[57]

Lazarus Group also increased the scope of its targets in terms of both geography and industry.[58] During Operation Sharpshooter, the group began using Dropbox with an IP address in the United States to share documents containing malicious macros.[59] Once activated, these malicious documents retrieved the malware Rising Sun, which shares code with the Duuzer malware that Lazarus used in 2015.[60]

In a similar vein to the reuse of malware code, the Group continued to behave in predictable ways, particularly with regard to the process of securely deleting logs and covering its tracks.[61] Kaspersky observed that it wiped malware payloads from targeted systems if the intruder had reason to suspect that he had been discovered.[62] This behavior recalls previous reporting from 2009, 2013, and 2016 that showed that Lazarus self-extracted from systems when detected by defenders or antivirus programs.

DISCLOSURE EVENTS

Lazarus Group exhibited an almost wanton degree of operational carelessness since its first attack in July 2009. Nevertheless, multiple reports confirm and reaffirm the Group's ability to conduct extended reconnaissance on targets and subsequently to dwell in those infected systems unawares. Moreover, despite numerous disclosures of TTPs and

[54] Kaspersky 2018 (n. 25), 14.

[55] id., 22.

[56] Guerro-Saade & Moriuchi 2018 (n. 27), 4.

[57] J. Rosenberg & C. Beek, 'Examining Code Reuse Reveals Undiscovered Links Among North Korea's Malware Families' (9 August 2018), *McAfee*: https://www.mcafee.com/blogs/other-blogs/mcafee-labs/examining-code-reuse-reveals-undiscovered-links-among-north-koreas-malware-families/

[58] McAfee 2018 (n. 35), 3.

[59] id., 4.

[60] id., 17.

[61] FireEye 2018 (n. 3), 22.

[62] Kaspersky 2018 (n. 25), 5.

malware, the Group has maintained this capability without retooling.

McAfee's Operation Troy: June 2013

McAfee's report into the 2013 DarkSeoul incident made significant progress toward understanding Lazarus' TTPs. The report expounded Lazarus' extended attack cycle, which is notable because Lazarus did little to change it in the sequel: the general blueprint of the group's attacks has been both public and unaltered since the beginning of its major known operations. Similarly, the report provided a detailed breakdown of the its early malware—NSTAR, EagleXP, HTTP Troy, Http Dr0pper, and TDrop. Lazarus improved and featured several of these tools in its later campaigns, making this initial analysis critical. The report similarly served as an early marker for the steps Lazarus Group would take to mislead investigators by posing as different groups. In this case, it pretended to be the NewRomantic Cyber Army Team and the Whois Hacking Team.

Palo Alto Networks Unit 42 TDrop2 Attacks: November 2015

Palo Alto's report demonstrated how the clues left behind by Lazarus can illuminate its operations and targets. Unit 42 uncovered an updated version of TDrop2, which held direct ties to Operation Troy—another side-effect of Lazarus' decision not to retool. It is worthwhile to note here that, although Lazarus hacked Sony Pictures in 2014, most of the information concerning its involvement in the hack was not yet public. This gave the false impression that it had been dormant since March 2013.

Novetta's Operation Blockbuster: February 2016

A Novetta-led coalition of private industry partners published this landmark report in the wake of the Sony Pictures hack. The report named Lazarus as the actor, connected the Sony hack to incidents in South Korea, revealed Lazarus' forty-five families of malware, provided examples of decoy documents and other artifacts from past campaigns, and detailed the C2 infrastructure.

Kaspersky's Lazarus Under the Hood: April 2017

Kaspersky's report investigated two separate cyberattacks against banks in the aftermath of the Bangladesh Central Bank heist, focusing on the group's behavioral patterns. The stated motivation of this report was similar to that of *Operation Blockbuster*, namely to raise the costs for the group by revealing and disrupting its operations. The report revealed the existence of the Bluenoroff subgroup, which focused on cybercrime. Additionally, it provided details on malware, zero-days, the group's anti-forensic techniques, and infection vectors. Kaspersky's timeline of events suggests that the initial breach of one South Asian bank coincided with the group's heist from the Bangladesh Central Bank.

The report concluded that Lazarus was 'operating a factory of malware', and that 'this level of sophistication is something that is not generally found in the cybercriminal world. It's something that requires strict organization and control at all stages of the operation. That's why we think that Lazarus is not just another APT actor'.[63]

Recorded Future's South Korean Cryptocurrency Users and Exchange: January 2018

Recorded Future disclosed Lazarus' campaign against South Korean cryptocurrency exchanges shortly before a dialogue between the two Koreas began. The campaign was notable (to this report's authors, at any rate) because the spear-phishing stage targeted college students interested in foreign affairs. The malware used in the campaign had code similar to the Destover malware, which Lazarus used in both the Sony Pictures hack and the first WannaCry incident from February 2017. The malware also used a Ghostscript exploit that specifically targeted Hangul Word Processor users.

Notably, the report outlined Lazarus' effort to mislead investigators into concluding that a Chinese APT was responsible.[64]

McAfee's Lazarus Resurfaces: February 2018

McAfee announced the start of a new campaign by Lazarus Group in January 2018. The campaign used DropBox to disseminate a malicious document which, upon opening, deployed an implant that collected and transmitted system data to a C2 server. Yet again, researchers identified Lazarus as the culprit due to similarities with previous campaigns.

U.S. DoJ - Park Complaint: June 2018

The indictment identified an individual member of Lazarus Group, Park Jin Hyok. It detailed the group's activities and information on Hyok's activities in dark web forums and his work for a North Korean front company, Chosun Expo, which allowed him to visit Dalian, China. The indictment revealed additional heist campaigns that were undertaken in Africa but previously undisclosed. It painted a complete picture of the group's TTPs, including reconnaissance, spear-phishing, and ransom attempts.

Kaspersky's Operation AppleJeus: August 2018

Operation AppleJeus detailed Lazarus' foray into the development of malware and exploits for OSX. Lazarus sent out emails that recommended a third-party cryptocurrency trading platform, but the application to which it linked was a malicious version that infected systems with an old Lazarus tool, Fallchill.

McAfee's Operation Sharpshooter: December 2018

Lazarus Group's 2015 backdoor malware Duuzer returned in the form of Rising Sun. This updated version of the implant was used to target nuclear, defense, energy, and financial companies across the world, affecting eighty-seven organizations in total. The report outlined the methodology of the campaign, which used DropBox servers with an obfuscatory IP addresses.

[63] id., 24.

[64] Guerro-Saade & Moriuchi 2018 (n. 27), 8.

Check Point - North Korea Turns Against New Targets?!: February 2019

Check Point discovered a campaign against Russian firms that matched Lazarus' known TTPs. The campaign used KeyMarble, a malware previously affiliated with the group. Despite the gap after *Operation Sharpshooter*, Lazarus did not change its behavior.

U.S. DoJ - Indictment of Chinese National in North Korea Crypto Scheme: March 2020

Although the indictment did not specifically mention Lazarus or Hidden Cobra, it detailed the exfiltration of funds by North Korea's Chinese partners—Tian Yinyin and Li Jiadong—and explained Lazarus' focus on accruing wealth for the North Korean government.

APT28 (RUSSIA)

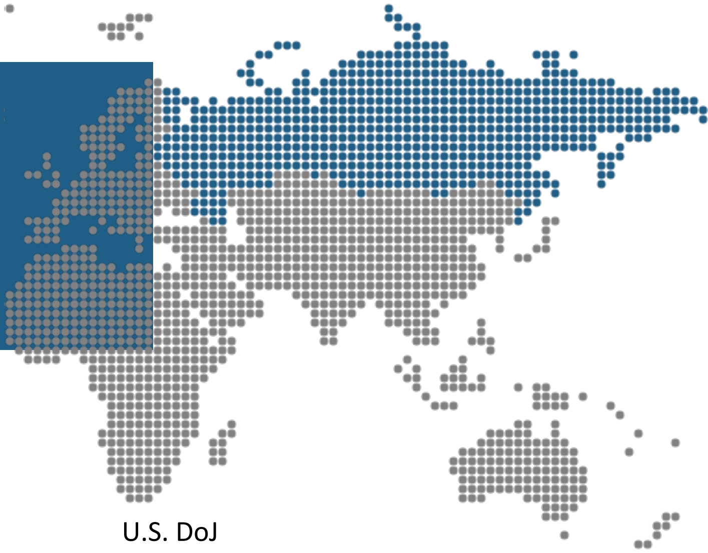

Timeline: Activity Levels and Disclosures

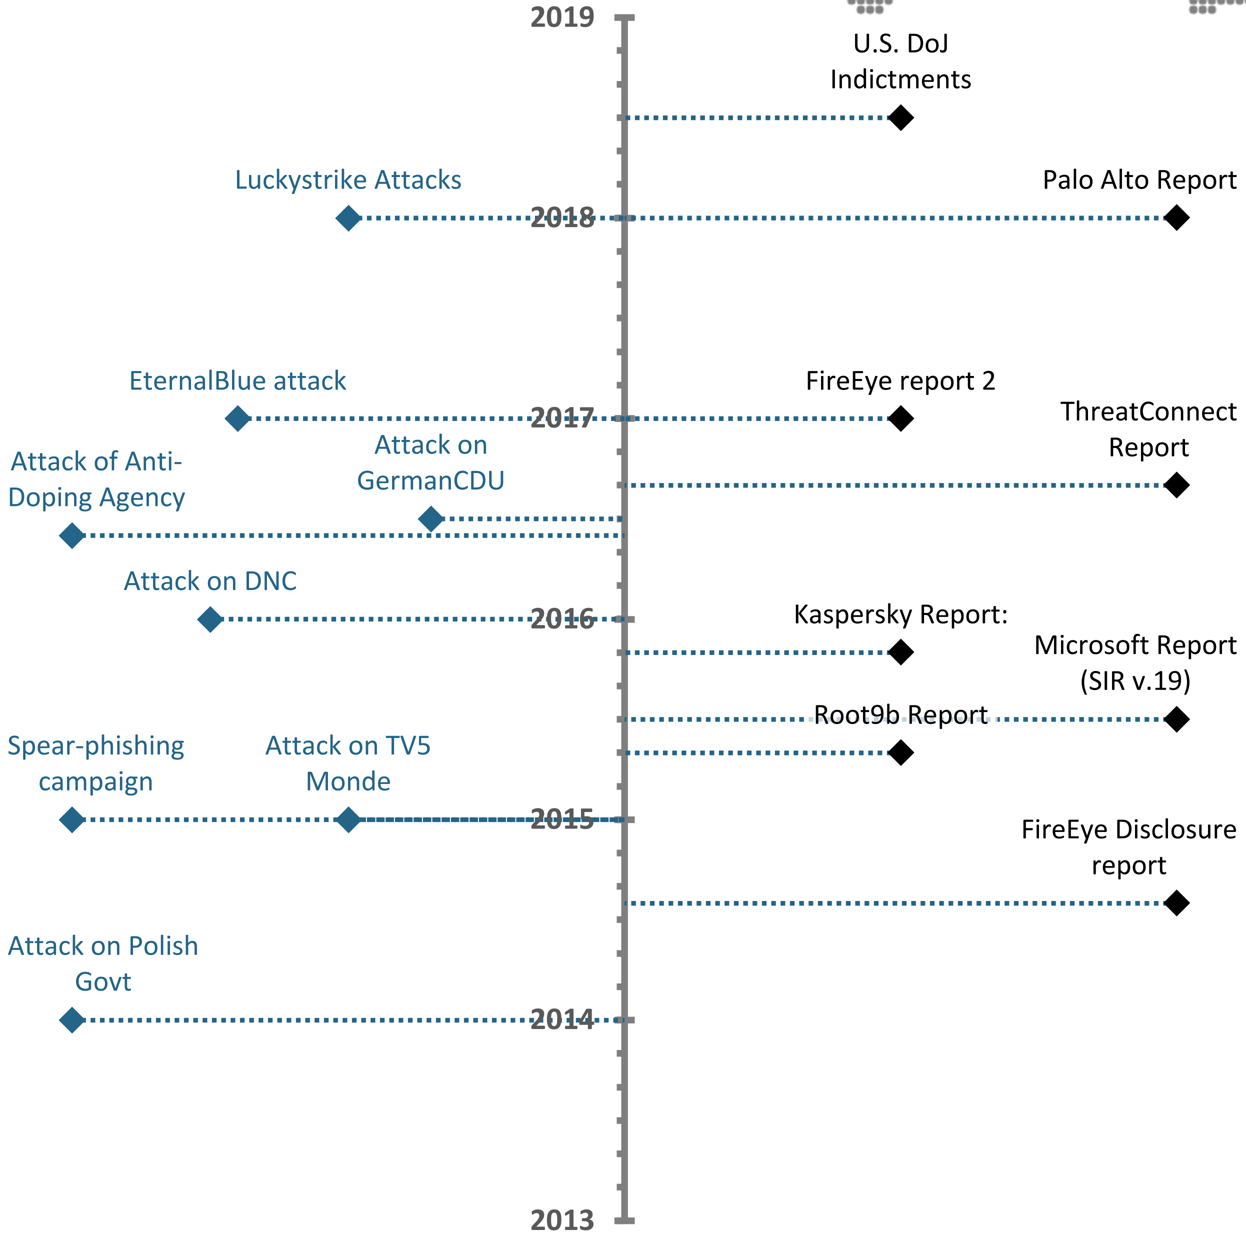

2019

U.S. DoJ Indictments

Luckystrike Attacks

2018

Palo Alto Report

EternalBlue attack

FireEye report 2

2017

ThreatConnect Report

Attack on GermanCDU

Attack of Anti-Doping Agency

Attack on DNC

2016

Kaspersky Report:

Microsoft Report (SIR v.19)

Root9b Report

Spear-phishing campaign

Attack on TV5 Monde

2015

FireEye Disclosure report

Attack on Polish Govt

2014

2013

IMPACT OF DISCLOSURES

INTRODUCTION

APT28 is a cyberthreat actor linked to the Main Intelligence Directorate of the Russian General Staff (GRU). Its campaigns have encompassed strategic espionage, property theft, and influence operations that align with Russian strategic and political objectives—specifically, providing the Russian government with decision advantage in diplomatic negotiations and undermining the West's trust in the liberal democratic principles and institutions. Numerous disclosures have revealed APT28's activities since its first operations in 2007. Nevertheless, APT28's campaigns continue, and it does so against an increasing range of targets, from military and defense espionage to information operations during democratic elections.

THE GROUP

APT 28 (Fancy Bear, Strontium, Tsar Team, Sofacy) is a Russian threat actor, active since at least 2007, whose objectives align broadly with Russian state interests. The group is known for its large-scale information operation campaigns against U.S. and European elections, Western security organizations and defense firms (especially NATO affiliated), and Eastern European countries and their militaries. Its activity primarily occurs during regular working hours for Moscow and St. Petersburg and is generally on Russian language platforms.[1]

As the number of APT28's campaigns have increased over time, so too has the number and diversity of its victims. Its initial campaigns began in 2007 and primarily targeted Eastern European nations, notably Georgia. From 2008 to 2014, it expanded into Western military and defense firms. [2] Most recently, APT28 has conducted major disinformation operations against the U.S. and European nations during the 2016 and 2017 election seasons as well as intrusions and disruptions against international organizations such as the World Anti-Doping Agency.

TIMELINE

APT28 was first observed in 2007 while conducting espionage against political and military targets in Georgia and Eastern Europe.[3] That would prove to be typical: APT28 primarily seeks defense and geopolitical intelligence to benefit the Russian state. During these early operations, APT 28 used the typical combination of spear-phishing emails and targeted malware to secure themselves a position on the networks they attacked. Its malware evolved over the next seven years to accommodate a revolving host of targets. Its focus fundamentally shifted toward the political in 2016, when it began conducting

[1] 'APT28: A Window into Russia's Cyber Espionage Operations?' (27 October 2014), *FireEye*: https://www.fireeye.com/content/dam/fireeye-www/global/en/current-threats/pdfs/rpt-apt28.pdf

[2] 'APT28', *FireEye*: https://www.fireeye.com/current-threats/apt-groups.html#russia

[3] FireEye 2014 (n. 1).

information operations against the citizenry of the U.S. (and later, France and Germany) in order to disrupt federal elections. The heavy activity which began in 2016 continued until the FBI indicted twelve GRU members for their role in disrupting the U.S. election in July 2018.[4] The pronouncement of this indictment coincides with a hiatus in their activity, which resumed in late 2018 with a series of attacks on think-tanks.[5] It continued operations through 2019 and into 2020 with attacks against the defense sector and political targets.

2007 APT28 is formed.[6]

2007-2014: Trend Micro publishes findings on Operations Pawn Storm, APT28's spear-phishing campaign against government and political organizations that used Sednit, a custom backdoor and information stealing malware.[7]

2014 APT28 conducts strategic web compromise attacks against the Polish government and energy company Power Exchange using Sednit, ultimately to deliver their custom Sofacy malware.[8]

2015 APT28's Coreshell malware is found on the networks of TV5 Monde following a website defacement incident for which the Cyber Caliphate claimed responsibility. Registration data associated with APT28 infrastructure is also found. This is believed to be an attempt at misattribution.[9]

Spear-phishing campaign against multiple German political parties; network of the Bundestag is compromised.[10]

[4] 'Grand Jury Indicts 12 Russian Intelligence Officers for Hacking Offenses Related to the 2016 Election' (13 July 2018), *U.S. Department of Justice, Office of Public Affairs*: https://www.justice.gov/opa/pr/grand-jury-indicts-12-russian-intelligence-officers-hacking-offenses-related-2016-election

[5] Think tanks included the Aspen Institutes in Europe, the German Council on Foreign Relations, and the German Marshall Fund. See S. Lyngaas, 'As Europe Prepares to Vote, Microsoft Warns of Fancy Bear Attacks on Democratic Think Tanks' (20 February 2019), *CyberScoop*: https://www.cyberscoop.com/european-think-tanks-hack-microsoft-fancy-bear-russia/

[6] FireEye (n. 2)

[7] Anomali Threat Research, 'APT28 Timeline of Malicious Activity', *Anomali*: https://forum.anomali.com/t/apt28-timeline-of-malicious-activity/2019

[8] P. Paganini, 'APT28: Cybercrime or State-Sponsored Hacking?' (4 June 2015), *Infosec Institute*: https://resources.infosecinstitute.com/apt28-cybercrime-or-state-sponsored-hacking/#gref

[9] S. Lyngaas, 'Lawmakers Call for Action Following Revelations that APT28 Posed as ISIS Online' (9 May 2018), *CyberScoop*: https://www.cyberscoop.com/lawmakers-call-action-following-revelations-apt28-posed-isis-online/

[10] A. Shalal, 'Germany Blocked Russian Hacking Attacks in 2016' (24 March 2017), *Reuters*: https://www.reuters.com/article/us-germany-elections-

Registration of *nato-news.com* and *bbc-press.org*; these domains host an Adobe Flash zero-day that is used against the Afghan Ministry of Foreign Affairs, Pakistani military, and NATO.[11]

2016

Spear-phishing campaign against Clinton campaign chairman John Podesta, the Democratic National Committee, and the Democratic Congressional Campaign Committee; emails subsequently leaked online via DC Leaks and WikiLeaks.[12]

Compromise of the World Anti-Doping Agency; medical data of athletes are released after accusations of doping against Russian athletes.[13]

Phishing campaign against members of the Christian Democratic Union; attempts made to gain account credentials through malicious macros and the Sednit malware.[14]

2017

Spear-phishing campaign against hotels in the Middle East and Europe; use of EternalBlue to move laterally through networks.[15]

2018

Use of open-source program Luckystrike to generate malicious documents for a global campaign against ministries of foreign affairs; Sofacy and Carberp used.[16]

2019

Scanning operation to locate vulnerable email servers, specifically targeting vulnerable webmail and Microsoft Exchange Autodiscover servers

Credential-harvesting campaign against Burisma

russia/germany-blocked-russian-hacking-attacks-in-2016-idUSKBN16V2FW

[11] Anomali Threat Research (n. 7).

[12] ib.

[13] A. Baldwin & J. Finkle, 'Anti-Doping Agency Says Athlete Data Stolen by Russian Group' (13 September 2016), *Reuters*: https://www.reuters.com/article/us-doping-wada-cyber/anti-doping-agency-says-athlete-data-stolen-by-russian-group-idUSKCN11J26T

[14] Thailand Computer Emergency Response Team, 'Threat Group Cards: A Threat Actor Encyclopedia' (19 June 2019), 212: https://www.thaicert.or.th/downloads/files/A_Threat_Actor_Encyclopedia.pdf

[15] '2018 Global Threat Report' (Feb/March 2018), *CrowdStrike*, 59: https://go.crowdstrike.com/rs/281-OBQ-266/images/Report2018GlobalThreatReport.pdf

[16] Z. Bederna & T. Szadeczky, 'Cyber Espionage Through Botnets', *Security Journal* 33 (2020), 53–6.

TYPOLOGY OF ATTACKS

APT28 consistently uses spear-phishing to gain initial access to networks. It has also used watering-hole style attacks to steal or spoof user credentials, most recently to help its influence campaigns in 2016 and 2017.

As a threat actor, APT28 is unlikely to be affected by public disclosures. APT28 develops its malware within a flexible framework that enables developers to alter code quickly and often enough to hinder the effectiveness of reverse engineering. [17] Moreover, it has a quality-control regime to ensure that its malware is systematically updated. These efforts render its malware more versatile than usual, which further complicates efforts to detect the malicious attachments in its phishing emails. Such software-developmental and operational sophistication is expected, given that APT28 is an established unit in the GRU and it functions to further Russian state interests.[18]

> *APT28 develops its malware within a flexible framework that enables developers to alter code quickly and often enough to hinder the effectiveness of reverse engineering.*

2007 through Mid 2013

Since its detection in 2007, APT28 has conducted a campaign dubbed Operations Pawn Storm, which broadly encompasses its geopolitical-themed spear-phishing emails that contain Sednit. This operation is thought to be ongoing. In 2013 it compromised the Georgian Ministry of Internal Affairs (MIA) with a malicious Excel file. Subsequent attacks against the MIA illustrated its sophisticated spear-phishing techniques: the malicious emails referenced legitimate aspects of the MIA network and were deceitfully signed by an actual sysadmin in Tbilisi.[19]

2014 through Early 2016

APT28 conducted a watering-hole attack that enabled it to infect the networks of the Polish government and the energy company Power Exchange with the Sofacy malware. The Sourface malware downloader also made an appearance as the payload of a malicious document that compromised the Georgian Ministry of Defense.[20]

FireEye published a report that implicated APT28 in the modification of DNS records to re-route the email traffic of the Kyrgyzstan Ministry of Foreign Affairs.

In 2015 APT28 used two zero-day vulnerabilities in Adobe Flash to compromise the Pakistani military, Afghan Ministry of Foreign Affairs, and NATO.[21] It also began to encrypt the files and data for exfiltration from a target network.[22]

[17] FireEye 2014 (n. 1).

[18] FireEye (n. 2).

[19] FireEye 2014 (n. 1).

[20] U.S. Department of Justice 2018 (n. 4).

[21] FireEye (n. 2).

[22] Anomali Threat Research (n. 7)

Early to Mid 2016

In coordination with APT29 and the broader Russian effort to disrupt the 2016 U.S. presidential election, APT28 conducted a large phishing campaign against individuals and institutions affiliated with, *inter alia*, the Democratic Party—the Clinton campaign's chairman John Podesta, the DNC, and DCCC.[23] APT28 used an URL-shortening service to trick recipients into visiting malicious sites, where their credentials were subsequently stolen. After compromising these secure networks and accounts, APT28 strategically released documents through the DC Leaks website and the threat actor Guccifer 2.0.[24]

Late 2016

APT28 compromised the World Anti-Doping Agency in retaliation for barring Russian athletes from participating in the Olympic games. A successful spear-phishing attack yielded APT28 a legitimate user account and other credentials that allowed them to log into the Agency's Administration and Management database. It then escalated its network privileges until it could access an International Committee account, whence it downloaded athlete data to publish.[25]

Late 2016 to Mid-2017

Following the success of the 2016 interference campaign, APT28 turned toward the upcoming European elections. It registered an email server that appeared to be associated with the Christian Democratic Union (a German political party) and from it they sent phishing emails to legitimate CDU members. These emails contained attachments with false hotel reservation information and malicious macros. Its Sednit malware and the open-source Responder tool facilitated lateral movement through CDU networks. APT28 continued to use public exploits, including EternalBlue, in a spear-phishing campaign against the European and Middle Eastern hospitality sector.[26]

Early 2018

APT28 continued to target Ministries of Foreign Affairs across the world with malicious Excel documents created by Luckystrike. Once enabled, the dropper installed and ran the primary payload—a variant of the group's custom SofacyCarberp backdoor.[27] The backdoor gathered system information and forwarded it to C2 servers, from where it was determined which additional malware would be needed to achieve its objectives.

In March 2018 it altered its spear-phishing tactics: it now sent Word documents with an Adobe Flash exploit. This new exploit required the victim to scroll through the entire three-page document before executing, as

[23] Paganini 2015 (n. 8).

[24] U.S. Senate Select Committee on Intelligence, 'Disinformation: A Primer in Russian Active Measures and Influence Campaigns' (30 March 2017): https://www.intelligence.senate.gov/sites/default/files/hearings/S%20Hrg%20115-40%20Pt%201.pdf

[25] Shalal 2017 (n. 10).

[26] ib.

[27] Anomali Threat Research (n. 7).

opposed to executing upon the document being opened.

By the middle of the year APT28 had launched a new phishing campaign that utilized the tool Zebrocy. Using the same attack vector of macro-enabled documents, Zebrocy prepared the infected machine with the second and third-stage malware needed to move laterally throughout target networks.[28]

Late 2018 to Early 2019

In September 2018 APT28 was found using a custom rootkit called Lojax, which is itself a variant of the legitimate anti-theft software Lojack.[29] This new rootkit targeted the unified extensible firmware interface and could maintain persistence on a machine even it shut down.

2019 to Present

Beginning in mid-November, APT28 used spear phishing and a watering hole attack to harvest credentials of Burisma employees. Burisma is the Ukrainian gas company on whose board Hunter Biden served; and APT28 targeted it to find sensitive information to use in operations against former Vice President Joe Biden's 2020 presidential campaign.[30]

APT28 also spent 2019 scanning the internet for vulnerable email servers. Specifically, it searched for vulnerable webmail and Microsoft Exchange Autodiscover servers on TCP ports 445 and 1433. Although its intent was unclear, APT28 likely used the compromised servers to harvest credentials and create more targeted and authentic spear phishing campaigns.[31]

DISCLOSURE EVENTS

Numerous cyber threat intelligence groups have reported on APT28's operations, TTPs, and tooling since 2014. Nevertheless, APT28 continues to operate with impunity and without significant disruption from these publications. Its TTPs and network/host artifacts have, however, evolved over time—albeit not for the desired reasons. These changes supervene on its general operational processes rather than any direct response to the disclosure event. That is to say: disclosures have had a limited impact because they encourage APT28 to do something that it had already planned to do. Namely, they encourage APT28 to change its TTPs and modify its malware. Furthermore, APT28's persistence in cyberspace and the growing scale of its campaigns offer additional signs that public disclosures have little to no effect. That makes sense—as a

[28] Baldwin & Finkle 2016 (n. 13).

[29] Not to be confused with the anti-theft platform for cars, though in essence they do the same thing.

[30] N. Perlroth & M. Rosenberg, 'Russians Hacked Ukrainian Gas Company at Center of Impeachment' (13 January 2020), *New York Times*: https://www.nytimes.com/2020/01/13/us/politics/russian-hackers-burisma-ukraine.html

[31] C. Cimpanu, 'APT28 Has Been Scanning Vulnerable Email Servers for More Than a Year' (20 March 2020), *ZDNet*: https://www.zdnet.com/article/apt28-has-been-scanning-and-exploiting-vulnerable-email-servers-for-more-than-a-year/

part of the Russian government, APT28 acts with the support of the Russian state. Resultingly, members of APT28 do not have to worry about a government crackdown or anything approaching a criminal investigation. Whether indictments and disclosures have a negative effect rather than the absence of a positive one remains to be seen.

FireEye Report: 2014

FireEye released the first major report on APT28. In it, they detailed the motivations, targets, TTPs, and timeline of the major intrusions that APT28 had conducted against military, defense, and foreign affairs networks. It contrasted APT28's campaigns with the objectives of the Russian government; and the contrast was slight. APT28 had gathered geopolitical intelligence from Eastern European nations and the Caucasus, military and defense firms, and NATO—all of which cultivated decision advantage for the Russian government. The report noted that over 89% of the malware attributed to APT28 was compiled during regular work hours for Moscow and St. Petersburg and in a Russian-language build environment. [32] Lastly, the report analyzed APT28's three key pieces of malware—Sourface, EvilToss, and Chopstick. It included a detailed appendix on the Sourface and Chopstick families of malware.

Disclosures have had a limited impact because they encourage APT28 to do something that it had already planned to do.

Root9b Report: May 2015

Root9b released a comprehensive report on APT28's campaigns against the financial sector. It listed the numerous domains registered for APT28's watering hole attacks and the hashes related to the zero-day exploits used in its attacks. The report also detailed the procedure used by APT28 to create fake identities for their internet registration activity.[33] Ultimately Root9b recommended that financial institutions specifically monitor their networks for spear-phishing campaigns.

Microsoft Report (SIR v.19): Summer 2015

Microsoft's Malware Prevention Center released further details on APT28's threat actor profile, including its TTPs and domains along with more technical indicators. The report covered APT28's use of spear-phishing and exploitation of zero-days, which appeared to be the group's two main vectors of attack. It also listed the legitimate domains hijacked by APT28 and its more plainly spurious domains. There was also a specific focus on APT28's use of Mimikatz and PassTheHash for credential theft and lateral movements between computers.[34]

Kaspersky Report: Winter 2015

Kaspersky published research on the zero-days in Office, Java, Adobe, and Windows that APT28 used in their campaigns. No less

[32] FireEye 2014 (n. 1).

[33] FireEye (n. 2).

[34] FireEye 2014 (n. 1).

important, Kaspersky analyzed the Azzy implant that hit air-gapped machines in the networks of high-profile defense contractors.[35] Notably, then, the report found that Azzy was delivered through a known piece of malware rather than a zero-day.[36]

ThreatConnect Report: Fall 2016

ThreatConnect released a report on APT28's attack against the World Anti-Doping Agency. The report included two of the domains used to compromise the agency's network, and further noted that these domains were the same as those used in the campaign against the DCCC.[37] The procedure used to register these domains was also found to be similar, if not identical, to how APT28 had historically registered domains.

FireEye Report: January 2017

FireEye released a second, more detailed report on APT28's activity of during the 2016 U.S. election and against the World Anti-Doping Agency. The report covered the four key characteristics of its TTPs: (1) a flexible framework to accommodate the evolution of its toolset over time; (2) the use of a formal coding environment to create and deploy custom modules within its backdoor programs; (3) the incorporation of counter-analysis capabilities; (4) that its malware was compiled during regular weekday working hours for Russia.[38] The report also covered APT28's strategic operations, though it omitted technical indicators such as domain names and IP addresses.

Palo Alto Report: Spring 2018

Palo Alto analyzed APT28's spear-phishing campaigns against the ministries of foreign affairs of various states. The report also outlined its use of the custom payload Sofaacy-Carberp for initial network reconnaissance and C2 server beaconing.[39] There was also discussion of APT28's use of Luckystrike to generate the malicious documents, and the report provided the usual list of indicators of compromise associated with these campaigns.

U.S. DoJ Indictment: July 2018

The U.S. Department of Justice indicted 12 members of the GRU who are affiliated with APT28 for actions taken in support of the interference campaign during the 2016 U.S. presidential election. The indictment carried eleven charges that ranged from identity theft and money laundering to conspiracy to commit computer crimes against U.S. government systems.[40]

[35] U.S. Department of Justice 2018 (n. 4).

[36] FireEye (n. 2).

[37] Anomali Threat Research (n. 7).

[38] Paganini 2015 (n. 8).

[39] U.S. Senate (n. 24).

[40] Shalal 2017 (n. 10).

APT29 (RUSSIA)

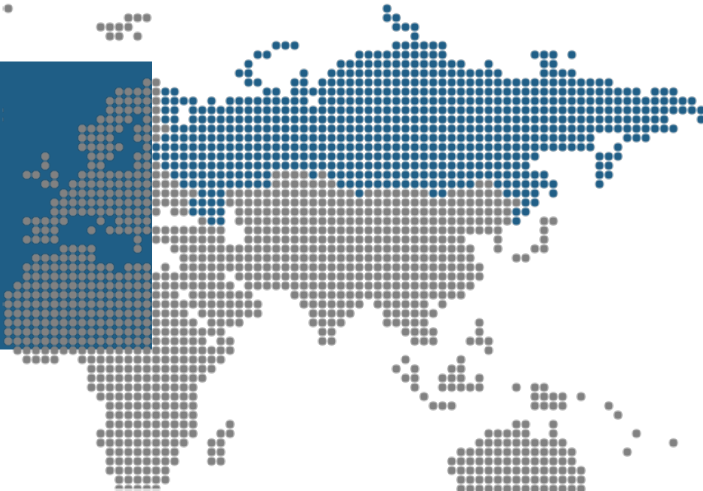

Timeline: Activity Levels and Disclosures

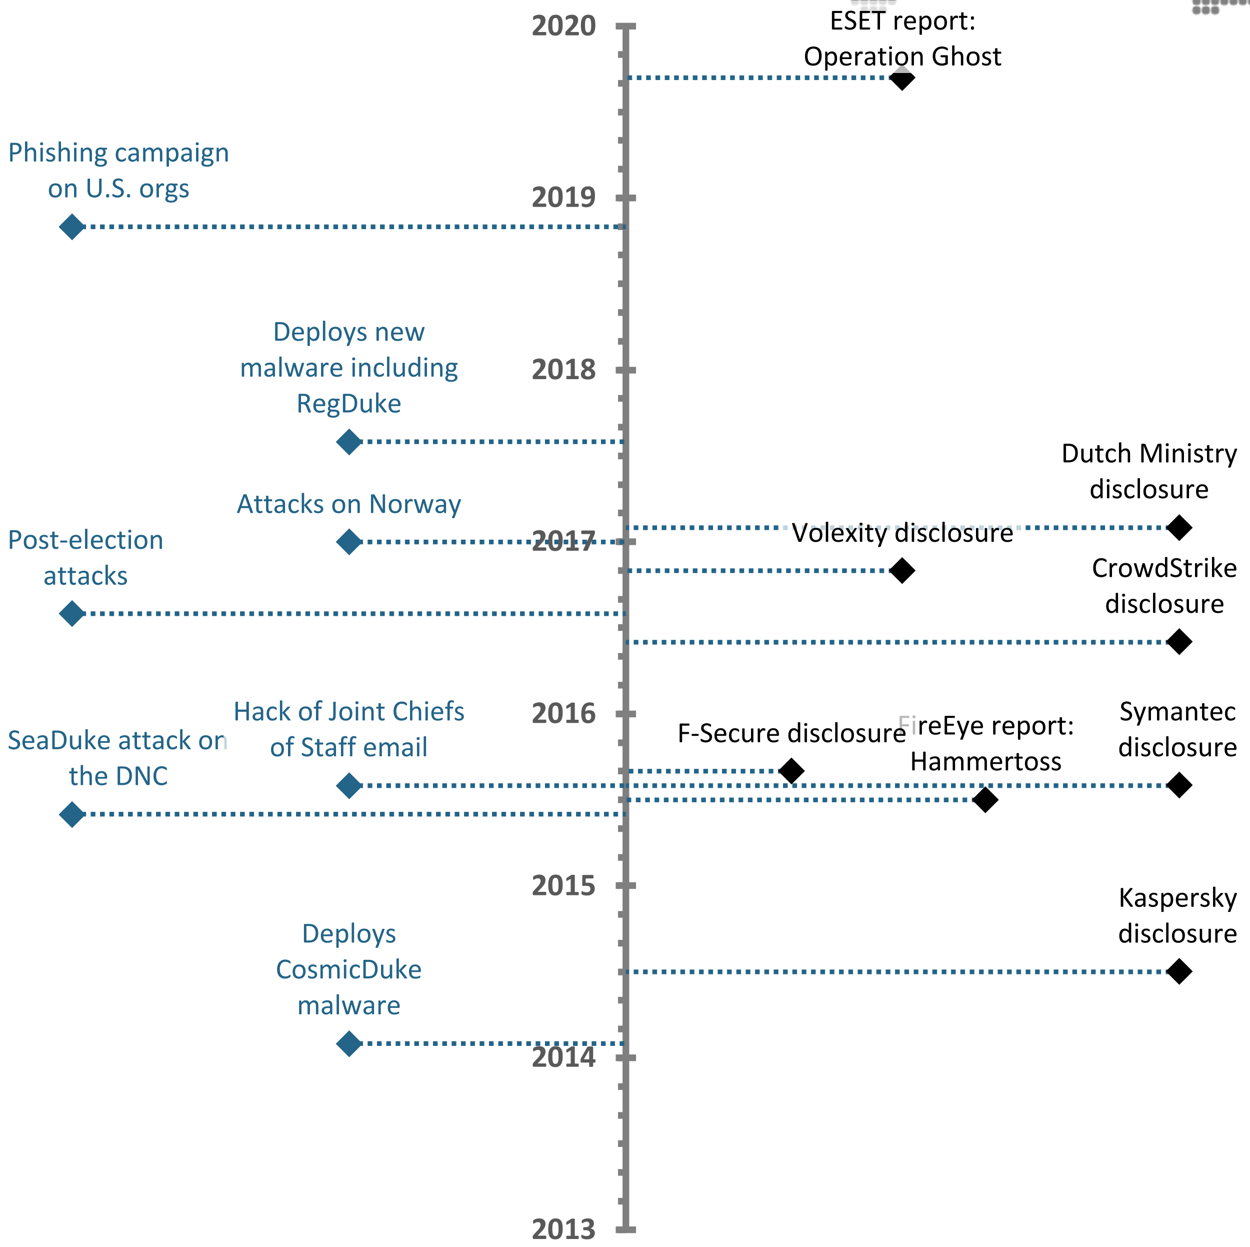

IMPACT OF DISCLOSURES

INTRODUCTION

APT29 is a Russian threat actor with suspected ties to the Russian Foreign Intelligence Service (SVR) or the Federal Security Service (FSB). APT29 has conducted multiple high-level campaigns that have primarily focused on thinktanks, government organizations, foreign ministries, and elections. It was a cause célèbre in 2016 for its involvement in the hack of Democratic National Committee servers and became the frequent subject of disclosure events. Most of these reports were issued from cyberthreat intelligence companies, and so they tended to focus on indicators of compromise; however, in 2018, Dutch intelligence (AIVD) attributed its actions with high-confidence to the Russian SVR. Yet, the effects of these disclosures are hardly forthcoming. Consistent with suspected Russian APTs, disclosures have failed to produce a discernible difference in APT29's behavior. The initiation and conclusion of its campaigns occur independent of disclosures: some campaigns begin before disclosures and carry through them, some start and end during, and every other possible variation.

THE GROUP

APT29 (Cozy Bear, Office Monkeys, CozyCar, The Dukes, Cozyduke) is a Russian threat actor known for campaigns against U.S. and European government departments and agencies, think tanks, research organizations, and NGOs. Although F-Secure has traced APT29's activity back to mid-2008, [1] APT29's first disclosure event occurred in 2014, when Kaspersky Labs reported on the evolution of the MiniDuke malware into the CosmicDuke backdoor. [2] Nevertheless, APT29 continued to elude major attention until the summer of 2015, when the DNC breach occurred and a FireEye report tied the group to the Hammertoss backdoor.[3]

The DNC hack evinced two highly disquieting qualities to find in a cyberthreat, aggressiveness and sophistication. The attack used a complex backdoor based on the SeaDuke implant that leveraged Windows Management Instrumentation to persist in the network. [4] Moreover, APT29 maintained a presence in DNC networks and servers for over a year, which is significantly longer than its counterpart APT28—perhaps suggesting an association with more persistent

[1] F-Secure Global, 'The Dukes: 7 Years of Russian Cyber-Espionage' (17 September 2015), *F-Secure*: https://blog.f-secure.com/the-dukes-7-years-of-russian-cyber-espionage/

[2] Kaspersky Global Research & Analysis Team (GReAT), 'Miniduke is Back: Nemesis Gemina and the Botgen Studio' (3 July 2014), *Secure-List*: https://securelist.com/miniduke-is-back-nemesis-gemina-and-the-botgen-studio/64107/

[3] 'Hammertoss: Stealthy Tactics Define a Russian Cyber Threat Group' (July 2015), Special Report, *FireEye*: https://www2.fireeye.com/rs/848-DID-242/images/rpt-apt29-hammertoss.pdf

[4] R. McCombs, 'Bear Hunting: Tracking Down COZY BEAR Backdoors' (27 September 2016), *CrowdStrike*: https://www.crowdstrike.com/blog/bear-hunting-tracking-cozybear-backdoors/

objectives. In the aftermath of this incident, CrowdStrike posited a connection to either the FSB or SVR,[5] whereas AIVD firmly attributed the group to the SVR.[6]

The timing of the attack is no less significant than the DNC hack itself. APT29's election interference efforts occurred after the disclosure of the Office Monkeys campaign just a year before. Its activity increased both during and after the DNC campaign, with multiple campaigns occurring concurrently against the Pentagon, think tanks, and Norwegian and Dutch ministries. Attributable activity disappeared after this string of campaigns, though it eventually resumed with a spear-phishing campaign during the second half of 2019.

To reiterate a previous point about APT29's sophistication and their tooling. Researchers have been taken aback by the insidiousness and quality of its TTPs and platforms, and its ability to develop advanced, bespoke toolsets in rapid succession. This trend or characteristic has been evident since OnionDuke, which quickly superseded CozyDuke and MiniDuke after Kaspersky reported on them in July 2014.[7] Two potential conclusions thus arise. APT29 could consist of an ingenuous team of hackers whose software development capacity renders disruptions nearly impossible—what one may expect from a military cyber unit or from a sophisticated criminal organization. At the same time, the group may prepossess an arsenal of toolsets, increasing in sophistication, that await deployment. Both possibilities, or a combination of the two, concede that highly capable adversaries are less vulnerable to disruption from disclosure events.

TIMELINE

Despite the frequent contemporaneity of APT29's campaigns, its activity profile neatly divides into three distinct periods from 2014 to the present.[8] The first block of activity revolved around the Office Monkeys campaign, which targeted a private research organization in March 2014.[9] The malware

[5] The Editorial Team, 'CrowdStrike's Work with the Democratic National Committee: Setting the Record Straight' (22 January 2020), *CrowdStrike*: https://www.crowdstrike.com/blog/bears-midst-intrusion-democratic-national-committee/

[6] R. Noack, 'The Dutch Were a Secret U.S. Ally Against Russian Hackers, Local Media Reveal' (26 January 2018), *Washington Post*: https://www.washingtonpost.com/news/worldviews/wp/2018/01/26/dutch-media-reveal-country-to-be-secret-u-s-ally-in-war-against-russian-hackers/; M. Faou, M. Tartare, T. Dupuy, 'Operation Ghost: The Dukes Aren't Back—They Never Left' (October 2019), *ESET Research White Papers*: https://www.welivesecurity.com/wp-content/uploads/2019/10/ESET_Operation_Ghost_Dukes.pdf

[7] Anomali Threat Research, 'APT29: A Timeline of Malicious Activity', *Anomali*: https://forum.anomali.com /t/apt29-a-timeline-of-malicious-activity/2480

[8] N.B. APT29 activity can be traced as far back as 2008, though without great significance given the lack of reporting.

[9] A. L. Johnson, '"Forkmeiamfamous": Seaduke, Latest Weapon in the Duke Armory' (13 August

used, CozyDuke, was a custom backdoor, and phishing was its primary method of distribution. Most notably, APT29 revised the code of CozyDuke within twenty-six days of its disclosure—a telling sign of minimal disruption.

Summer 2015 to the autumn of 2016 encompassed the second period of activity. In June 2015, the SeaDuke malware penetrated the DNC servers and propagated through the network via Mimikatz. This attack persisted for over a year. Major disclosures from Symantec, FireEye, and F-Secure all failed to disrupt APT29's activity: the email server of the Joint Chiefs of Staff was compromised just two months later in August 2015. That campaign continued to 2016, when it ended with a flurry of post-election spear-phishing attacks against U.S. think tanks and NGOs.

These attacks used a different toolset and TTPs, with Microsoft Word and Excel documents, RATs, and steganography becoming the preferred media to upload a new backdoor named PowerDuke.[10]

The third spell of activity is believed to stretch from January 2017 through the present. The difficulty in establishing this timeframe owes partly to the stealth and complexity of APT29's new toolsets, and partly to the uncertainty of whether it ever truly ceased activity in late 2016.[11] This round of activity began in January 2017 with a spear-phishing campaign against the Norwegian Government and the Norwegian Labor Party. There has been considerable difficulty since then in discerning where, or if, any new activity is taking place. But the appearance of four new families of malware—PolyglotDuke, RegDuke, FatDuke, and LiteDuke—would suggest that APT29 is far from retired.[12]

2014

Feb: Malware evolution occurs, deploying CosmicDuke, an upgrade from MiniDuke and CozyDuke.[13]

Mar: APT29 compromises a U.S. private research group.[14]

Jul: Kaspersky releases a report on APT29;[15] APT29 revises CosmicDuke code to bypass detection systems, twenty-six days after Kaspersky's abovementioned disclosure.[16]

2015), *Broadcom*: https://community.broadcom.com/symantecenterprise/communities/community-home/library documents/viewdocument?DocumentKey=6ab66701-25d7-4685-ae9d-93d63708a11c&CommunityKey =1ecf5f55-9545-44d6-b0f44e4a7f5f5e68&tab=librarydocuments

[10] S. Adair, 'PowerDuke: Widespread Post-Election Spear Phishing Campaigns Targeting Think Tanks and NGOs' (9 November 2016), *Veloxity*: https://www. volexity.com/blog/2016/11/09/powerduke-post-election-spear-phishing-campaigns-targeting-think-tanks-and-ngos/

[11] M. Faou *et al.* 2019 (n. 6).

[12] id., 5.

[13] GReAT 2014 (n. 2). N.B. This occurred before any significant disclosure.

[14] Johnson 2015 (n. 9).

[15] GReAT 2014 (n. 2).

[16] F-Secure Global 2015 (n. 1), 10.

Oct: APT29 develops SeaDuke malware.[17]

2015

Jun: APT29 uses SeaDuke to breach the DNC and Mimikatz for lateral movement.[18]

Jul: FireEye publishes a report on Hammertoss.[19]

Aug: APT29 hacks Joint Chiefs of Staff email server;[20] Symantec discloses APT29 activity.[21]

Sep: F-Secure discloses APT29 activity.[22]

2016

Jun: CrowdStrike discloses DNC hack details.[23]

Aug: APT29 post-election spear-phishing campaign is underway against U.S. based think tanks and NGOs.[24]

Sep: First known deployment of FatDuke.[25]

Nov: Volexity discloses the post-election spear-phishing campaigns.[26]

2017

Jan: APT29 conducts attacks against the Norwegian Ministries of Defense and Foreign Affairs and the Norwegian Labor Party.[27]

Feb: Norwegian police attribute attacks to APT29;[28] Dutch ministries disclose APT29 attacks against them.[29]

Mar: FireEye reveals APT29's new 'domain fronting' techniques.[30]

Aug: First in-the-wild sighting of RegDuke.[31]

[17] Johnson 2015 (n. 9).

[18] CrowdStrike 2020 (n. 5).

[19] FireEye 2015 (n. 3).

[20] B. Starr, 'Official: Russia Suspected in Joint Chiefs Email Server Intrusion' (7 August 2015), *CNN*: https://edition.cnn.com/2015/08/05/politics/joint-staff-email-hack-vulnerability/

[21] Johnson 2015 (n. 9).

[22] F-Secure Global 2015 (n. 1).

[23] CrowdStrike 2020 (n. 5).

[24] Adair 2016 (n. 10).

[25] ESET Research, 'Operation Ghost: The Dukes Aren't Back–They Never Left' (17 October 2019), *WeLiveSecurity*: https://www.welivesecurity.com /2019/10/17/operation-ghost-dukes-never-left/

[26] S. Adair 2016 (n. 10).

[27] 'Norge utsatt for et omfattende hackerangrep' (3 February 2017), *Norsk rikskringkasting*: https://www.nrk.no/norge/norge-utsatt-for-et-omfattende-hackerangrep-1.13358988

[28] ib.

[29] P. Cluskey, 'Dutch Opt for Manual Recount After Reports of Russian Hacking' (3 February 2017), *The Irish Times*: https://www.irishtimes.com/news/world/europe/dutch-opt-for-manual-count-after-reports-of-russian-hacking-1.2962777

[30] M. Dunwoody, 'APT29 Domain Fronting with TOR' (27 March 2017), *FireEye*: https://www.fireeye.com/blog/threat-research/2017/03/apt29_domain_frontin.html

[31] ESET Research 2019 (n. 25).

2018	**Nov**: APT29 conducts phishing campaign against several U.S. organizations.[32]
2019	**Oct:** ESET releases report on Operation Ghost, disclosing multiple new toolsets which suggest that APT29 never ceased activity.[33]

TYPOLOGY OF ATTACKS

APT29 is known to have a wide range of bespoke malware variants and complex C2 arrangements. This includes its ever-evolving "Duke" series of backdoors and C2 systems[34] and the hijacking of Twitter and GitHub cloud servers for use as storage devices.[35] APT29 primarily uses phishing and spear-phishing campaigns to deliver its malware variants. Lures are often malicious Microsoft Word and Excel documents that contain the usual falsities alongside authentic information harvested from other organizations.[36] The group sometimes employs steganography to mask these downloads.[37] Once inside the target network, APT29 generally uses Mimikatz to achieve lateral movement.

APT29 possesses a wide range of alternative techniques. In 2017 FireEye/Mandiant reported on a new technique being used, 'domain fronting', which spoofs outbound network connections to look like requests from commonly visited websites.[38] It also recompiles and modifies its code to bypass detection semi-regularly, and uses different C2 infrastructure for different victims, lest the compromise of one operation lead to the discovery of others.[39] That organizational decision allows it to conduct multiple campaigns without aggregating the usual risks involved in running several concurrent operations.

November 2008 – January 2010

APT29's initial campaigns began in November 2008 and used a bespoke trojan, PinchDuke, which it delivered via spear-phishing. Targets included Western organizations like NATO and certain think tanks. Without the impediment of a disclosure, PinchDuke evolved to another custom variant, GeminiDuke.[40]

Spring 2010

Attacks against organizations associated with foreign affairs continued into March 2010. The toolset evolved from the PinchDuke trojan to the CosmicDuke toolset focused on compromising and stealing information.[41]

2011

APT29 added MiniDuke and CozyDuke to its arsenal. MiniDuke was a simple backdoor

[32] ib.

[33] M. Faou *et al.* 2019 (n. 6), 5.

[34] For further information on the current versions of the "Duke" malware: ESET Research 2019 (n. 25).

[35] FireEye 2015 (n. 3).

[36] Anomali Threat Research (n. 7).

[37] ib.

[38] FireEye 2015 (n. 3).

[39] ib.

[40] F-Secure Global 2015 (n. 1).

[41] id., 6.

that enabled remote code execution. CozyDuke was a modular designed malware platform connected to a C2 server that was used to select which modules to deploy. CozyDuke was notably written in C++, signifying a major evolution in its design and complexity compared to past variants.[42]

February 2013 – July 2014

APT29 consistently delivered MiniDuke through spear-phishing emails containing a PDF exploit. As awareness of MiniDuke grew, OnionDuke appeared in an ostensibly separate campaign. OnionDuke's targets were not immediately clear because the vector of delivery was Torrents.[43]

July 2014

APT29 modified the code for CosmicDuke and the loader for both MiniDuke and CosmicDuke. It completed the recompilation, which allowed both to bypass security detection systems, three weeks after the disclosure event that precipitated such a change.[44] Other campaigns remained active while recoding was completed, which suggests the disclosure had a limited impact.

July 2014

APT29 began the Office Monkeys campaign, its first large operation with CozyDuke. This campaign did not use the standard lure of a malicious PDF document, instead opting for unique ones: a standard message mimicking an e-fax and a video of an advertisement from the 2007 Super Bowl.[45] That suggests a distinct evolution in TTPs.

October 2014 - February 2015

October 2014 marked the debut of SeaDuke, a new toolset developed solely in Python and designed to work both in Linux and Windows environments. In January 2015, APT29 unveiled yet another new technique to unload multiple malware variants on a target for enhanced network persistence. APT29 also added HammerDuke to its arsenal. This malware used an algorithm to change Twitter accounts periodically, which it leveraged to communicate with its C2 servers unsuspiciously.[46]

June 2015 – August 2016

Spear-phishing emails delivered SeaDuke to the DNC. These e-mails redirected recipients to a malicious website with a dropper that would subsequently download one of a series of RATs.[47] APT29 also returned to using malicious Microsoft Excel and Word documents that contained previously-harvested authentic information. These documents dropped another malware variant, PowerDuke. The attacks used steganography to deploy components of the backdoor.

November 2016

APT29 commenced a new spear-phishing campaign with PowerDuke. The lures were well crafted—somewhat the standard for

[42] id., 7.

[43] id., 9.

[44] id., 10.

[45] id., 11.

[46] id., 13.

[47] CrowdStrike 2020 (n. 5).

APT29—and they exploited the post-election environment with email subjects like 'The Shocking Truth about Election Rigging in the United States'. These emails contained a Microsoft shortcut file that in turn executed PowerShell commands to plant a backdoor.[48]

January 2017 – Present

APT29's crosshairs shifted to European elections and ministries, though it continues to target U.S.-based think tanks. Its techniques remain somewhat consistent, albeit with heightened sophistication. It developed an insidious four-stage system that drops increasingly harmful malware. First, it drops its new tool, Polyglot Duke, to communicate with C2 servers. The C2 servers then drop another new tool, RegDuke, which acts as a recovery tool. That, in turn, facilitates the deployment of a new but very similar version of MiniDuke. Finally, it deploys its newest backdoor, FatDuke. It again conducts lateral movement through credential harvesting. Well-crafted spear-phishing emails are the chosen attack vector. [49]

DISCLOSURE EVENTS

APT29 appears to operate irrespective of disclosure events, even going so far as to leave taunting messages for defensive operators in some cases. [50] It develops and modifies toolsets *ad libitum*, with new malware generally being built before any disclosure event has revealed its last campaign. As mentioned, when Kaspersky Labs released the first major disclosure of its activity in 2014, APT29 recompiled the code for CozyDuke within twenty-six days. Further, the group did not halt its ongoing campaigns.

APT29's resilience to disclosure events is especially apparent during the period of May 2015 to November 2016. This interval saw an interwoven series of attacks, disclosures, and toolset evolution; however, campaigns primarily leveraged the well-known CozyDuke and SeaDuke malware. That is telling. It may indicate a circumstantial preference for objective-completion over stealth; alternatively, it could mean that disclosures do not engender a significant enough amount of change to "burn" elements of the group's toolkit. At the very least, it raises doubts about the preclusive value of disclosures if the effect on the APT is null or negligible.

APT29's resilience to disclosure events is especially apparent during the period of May 2015 to November 2016.

This is not to suggest that all disclosures ineffectual and fruitless. Indeed, the F-Secure report from September 2015 apparently prompted a pause in activity. Nevertheless, it remains uncertain whether the disclosure directly or predominately caused this lull. The pause may well have occurred because resources were needed for another

[48] Adair 2016 (n. 9).

[49] Faou *et al.* 2019 (n. 6), 11.

[50] F-Secure Global 2015 (n. 1), 7.

objective, or because the campaign itself was over and another had already begun undetected. The present study of APT29, however, concludes that the F-Secure disclosure was not the likely cause for the break in their activity: the F-Secure disclosure was too similar to other disclosures which themselves failed to induce any clear behavioral change.

July 2014

Kaspersky released the first major disclosure against APT29. It included details on the MiniDuke malware and identified old and new variants. The report was written independently of F-Secure's report on the same malware (which they dubbed "CosmicDuke"), and it detailed all associated TTPs, and toolsets. [51] APT29's campaigns continued as normal.

July 2015

FireEye released the Hammertoss report, which disclosed the innerworkings of the tool and its associated TTPs, and which attributed it to Russia and APT29.[52] APT29 attacked the Pentagon with a different toolset the next month.

August 2015

Symantec released a report on MiniDuke and CozyDuke as well as the TTPs used in the Office Monkeys campaign. Symantec also uncovered and provided details on SeaDuke.[53] APT29 had already taken action on the related objectives.

September 2015

F-Secure released a detailed disclosure that chronicles APT29's tool development, TTPs, and campaigns from 2008 to 2015. It included detailed hashes, indicators of compromise, network infrastructure and artifacts, and a thorough timeline of campaigns. The report attributed APT29 to Russia.[54] It may be argued that this disclosure caused disruption, given the subsequent intermission in APT29 activity. But the apparent disruption more likely resulted from pure chance: several other campaigns also ended around that time.

June 2016

CrowdStrike released a forensic analysis of the DNC hack, which they attributed to APT29. They provided the indicators of compromise and TTPs and opined that SeaDuke compromised the DNC—despite all the reports that had been published beforehand. This suggests, if nothing else, that APT29 was confident enough in the sophistication of its malware to continue using it even after the world had been notified of its existence. It also attests to the lack of disruption caused by Symantec's disclosure on SeaDuke. Lastly, CrowdStrike offered attributive evidence to connect APT29 to the FSB or SVR.[55] A new campaign of post-election spearphishing began two months later.

November 2016

[51] GReAT 2014 (n. 2).

[52] FireEye 2015 (n. 3).

[53] Johnson 2015 (n. 9).

[54] F-Secure Global 2015 (n. 1).

[55] CrowdStrike 2020 (n. 5).

Volexity published a report that detailed the post-election spear-phishing campaigns. The disclosure provided details on the new PowerDuke toolset.

February 2017

Norwegian and Dutch governments revealed that APT29 attacked them, and they attributed the attack to the Russian SVR.[56] These events unfolded while APT29 developed a new toolset, and they may have precipitated or otherwise encouraged a recess in its activity. Yet the reality of the disruption remains unclear as before. It seems entirely possible that APT29 had already moved on to another campaign, making use of a new and advanced toolset. Moreover, instead of being halted by disclosures, APT29 often exploits the immediate post-disclosure environment, actively using the news and disclosures about itself to create new lures.

APT29 often exploits the immediate post-disclosure environment, actively using the news and disclosures about itself to create new lures.

March 2017

FireEye released a report naming the group as 'Russian nation-state actors APT29'. The report detailed an advanced backdoor that used 'domain fronting' by leveraging TOR, in addition to disclosing the expected TTPs and toolsets.[57]

October 2019

ESET published a major disclosure of APT29 activity. They asserted that APT29 had not been inactive or underground, as many believed. Instead, it was actively campaigning with a new set of advanced tools—PolyglotDuke, RegDuke, FatDuke, LiteDuke. Its targets and TTPs were found to be consistent with past objectives and techniques, with a focus on government entities and Western think tanks. An upgraded variant of MiniDuke also debuted as part of the abovementioned four-stage delivery system.[58]

[56] 'Norway Institutions "Targeted by Russia-Linked Hackers"' (3 February 2017), *BBC*: https://www.bbc.com/news/world-europe-38859491

[57] Dunwoody 2017 (n. 29).

[58] Faou *et al.* 2019 (n. 6).

DIFFERENCES BETWEEN ACTOR RESPONSE

Both Chinese groups, APT1 and APT10, disappeared completely. The Cobalt Group was not affected by disclosures, which always came after its operations had finished. Amongst the Iranian groups, APT33 appears to act independently of disclosure events, while APT34 was disrupted for a maximum of eight to ten weeks. The North Korean group, Lazarus, continued to operate using the same tools despite the disclosures and indictments. One Russian group, APT28, ceased operations for three months following an indictment, while the other Russian group, APT29, was not impacted by disclosures.

China

APT1, which was affiliated with the Chinese People's Liberation Army, was the subject of one high-profile private sector disclosure and a subsequent U.S. Department of Justice Indictment. The Mandiant APT1 report was the first of its kind; and following Mandiant's release of the report, APT1 halted activity for over a month. When it resumed operations, it took nearly six months to return to pre-report levels of activity. The efficacy of this disclosure is attributed to two main factors. First, APT1 relied on custom malware that needed to be rebuilt in its entirety. Second, the Chinese government had concerns about its appearance and reputation, and therefore wanted to manage any global fallout that might come with the report's release. After the DoJ indicted several of APT1's members fifteen months later, it ceased operating entirely. Although the indictment appears to be successful because the group was disbanded, remaining members and targets were likely reassigned to other groups. That renders the strategic impact of the indictment difficult to ascertain.

APT10, which was affiliated with the Chinese Ministry for State Security, was the subject of several high-profile disclosures. Following a 2013 report by FireEye, APT10 ceased operating for several months while it retooled. After its return, disclosures did not seriously affect the group again until the 2017 *Cloud Hopper* report and the 2018 U.S. Department of Justice indictment because it had continuously retooled and prepared for future campaigns. The *Cloud Hopper* report halted operations for eight to ten weeks while the group retooled, and its success may stem from the multi-stakeholder response that accompanied the report. APT10 disappeared entirely after the indictment.

Criminal Groups (Cobalt)

Cobalt, a criminal group that targets financial institutions, has been the subject of many disclosures but continues to operate. Disclosures on Cobalt's activities tend to be after-the-fact: the Group (and its dedicated development team) moves so quickly through operations, tooling, and tactics that it stays ahead of cybersecurity vendors. It relies on single-use, institution-specific, custom malware and an attack style that is difficult to detect and stop, coopting normal bank processes and workflows to siphon money out of banks. Driven by its profit-motive, Cobalt is unincentivized by disclosures that do not affect its ability to steal.

Iran

APT33, which is likely a government proxy group, has continued its operations despite several comprehensive disclosures by private cybersecurity vendors. Since its first major operation in 2016, it has consistently conducted a new campaign every few months; and pauses in its activity have not necessarily aligned with the timing of disclosures. It has grown stronger and more sophisticated over time, and it has increased the scope of its targeting. APT33's operations have continued for a variety of reasons, including its ability to switch between commodity and bespoke tools with relative ease, its sophisticated spear phishing tactics, and its status as a proxy group. Although Iran should, under international law, take responsibility for non-state actors operating within its borders, APT33's status as a proxy group raises the bar for the international community if it wants to issue indictments, retorsions, and countermeasures against the Iranian government—or otherwise incentivize Tehran to pressure the group to cease its operations. By using the contractor model to farm out its cyber operations, Iran can more plausibly deny its connections to APT33 and so absolve itself of responsibility.

APT34, which is likely a government proxy group, has been the target of over twenty-five major reports since 2016. Reports have impacted the group, as it continuously updates old malware, infrastructure, and tactics to defeat detection. It nonetheless continues to operate. The one notably effective disclosure was the large Lab Dookhtegan leak, which halted APT34's operations for a full eight to ten weeks. This leak distinguished itself from other disclosures through its sheer quantity of information, the possibility that it was an insider job, and the correlation of individual members to their cyber persona. APT34 can quickly resume its activities after each disclosure due to its use of bespoke tools and techniques that can be easily updated and its use of social engineering.

North Korea

APT38 is affiliated with North Korea's state intelligence apparatus, the Reconnaissance General Bureau. Despite several disclosures on its activities, it has conducted numerous high-profile attacks since the 2014 breach of Sony Pictures, including the 2016 Bank of Bangladesh heist and the deployment of WannaCry in 2017—a malware developed to raise money or hell, unclear which. APT38 appeared similarly unphased by the 2018 U.S. Department of Justice indictment that outlined its organizational details: it continued to target crypto exchanges and conduct ransomware attacks, though it did remain out of the news until the recent U.S. government statement that it is offering its cybercrime services for hire. The group remains resilient to public disclosures by frequently updating its custom toolset and spawning subgroups for different types of campaigns and targets. Additionally, although the group originally found notoriety for its politically motivated campaigns, APT38 has transitioned to a more cybercriminal model and now appears to be financially motivated. It has displayed an intent to remain in the realm of financial crimes so long as it can earn money for the North Korean government.

Russia

APT28 is affiliated with the Main Intelligence Directorate of the Russian General Staff (GRU). It has continued to operate despite several disclosures on its activities; and, in some instances, its activity level has increased following the release of reports. Its operations temporarily ceased following the U.S. Department of Justice indictment in 2018, though it resumed targeting in a few months. Over time, APT28 has increased the sophistication and scope of its campaigns. Its operations have continued for a variety of reasons, including a flexible operational and development framework that allows its malware developers to evolve quickly—not to mention encouragement by the Russian government to continue operating.

APT29 is affiliated with the Russian Foreign Intelligence Service (SVR) or the Russian Federal Security Service (FSB). Although it has been the subject of many public disclosures, APT29 appears unaffected. Reporting does not stop its campaigns, even when said reporting covers an active campaign. Moreover, disclosures sometimes appear to embolden APT29, as its activity level has increased following public reporting and taunting messages have been left in its malware for American investigators. APT29 is aided in these endeavors by a highly advanced arsenal of tools and a demonstrated ability to create new ones with remarkable rapidity.

CONTRIBUTING FACTORS TO SIMILARITIES AND DIFFERENCES

Disclosure by Private Vendor vs. by U.S. DoJ Indictment

Private cybersecurity vendors release private reports to paying customers on the activities of relevant threat groups, but these reports are not the focus of this research project. We are primarily concerned with public disclosures, and they do not appear to have had a lasting effect on any group. Disclosures have compelled each of these groups to become more sophisticated over time, and they have not caused the APTs under consideration to cease operation. Only APT1 and APT10 were stopped by vendor disclosures, albeit temporarily; and neither outfit shut down until the DoJ issued indictments. Even then, their personnel and targets were probably shifted to other Chinese groups. The APT1 report, being a public disclosure by a private vendor, likely had its particular effect because it was the first report of its kind: the equally comprehensive *Cloud Hopper* report on APT10, released a few years later, lacked the same long-lasting effect. Moreover, PwC UK issued that report in conjunction with the affected managed service providers, specifically to coordinate an effort to extirpate APT10 from their networks and force the group to resume operations *da capo*. And while the report does seem to have engendered several weeks of disruption, APT10 was operating at a normal capacity approximately ten weeks later—despite the coordinated actions of many stakeholders.

The effects of indictments on groups in other countries is unclear due to lack of pertinent evidence. In theory, indictments should have a more deterrent effect on contractors than they do on government employees, as contractors may have other commercial and travel interests that state agents lack. Indictments could hamper the ability of a contractor to do business, while civilian and military intelligence operators have little choice but to follow the orders handed down by their superior officers.

Indictments should have a more deterrent effect on contractors than they do on government employees, as contractors may have other commercial and travel interests that state agents lack.

At present, there is not a large enough sample size of non-state and proxy APTs that have been the subject of indictments to determine their effects on these actors. Cobalt may be affected by indictments, but it will likely continue to operate so long as it can make a profit and its members are not jailed.[1] Iran may be angered by indictments, but the Iranian proxy actors we reviewed have yet to be indicted.

Of the non-Chinese state actors included in this report, only North Korea operates well outside the bounds of normal diplomacy. The international community has exhausted sanctions against the North Korean regime, which has left them with little left to lose. Hence their lack of motivation to cease operations on account of indictments.[2] North Korean operators also do not travel to places that have mutual extradition treaties with the United States, and so they face little chance of ever seeing the inside of a courtroom. It also seems plausible that they are given little choice by the authoritarian government but to continue carrying out operations. The Russian groups present a similar story, insofar as Russian APTs are thinly veiled tools of the state. Nevertheless, Russian APTs have shown themselves to be unmotivated by fears of diplomatic backlash and will therefore likely remain unaffected by indictments in the long term.

Publicity of the Disclosure

[1] Interview with an unnamed source.

[2] As noted in the interviews with Jenny Jun (PhD candidate at Columbia University, North Korea focus) and Dmitri Alperovitch (co-founder, former Chief Technology Officer, CrowdStrike).

The high publicity of public disclosures appears to have affected both Chinese groups, but publicity did not seem to have lasting effects on any of the other APTs. This may be due to China's interest in saving face and its strategy of extending its global influence. The Chinese government does not want to appear out of control, so its groups are affected more by disclosures. APTs in other countries and contexts have also been the subject of small disclosures and of large reports, but their operations all continued because their motivations are different. Purely criminal groups, like Cobalt and North Korea's Lazarus (post-Sony), care primarily about making money. Iran wants to deter other threats, but not at the cost of plausible deniability, which it assiduously maintains even when disclosures receive great publicity.

The high publicity of public disclosures appears to have affected both Chinese groups, but publicity did not seem to have lasting effects on any of the other APTs.

One interview subject contends that Russia reacts differently to China in the face of highly public disclosures because it wants to relive its glory days as a world power. Public disclosures serve as free press, giving it the illusion of being bulletproof and unstoppable. [3] This behavioral difference could also be explained by Russia's muted presence in the global economy when compared to China. Beijing has a strong interest in promoting China as a good place for foreign investment and business. No less significant, Russian and Chinese APTs pursue different goals. Chinese groups have focused on economic espionage, which entails that stealth and a general dearth of conspicuity redound to their continued success: Russian APTs have engaged in large scale disruption and destabilization campaigns across Europe and the U.S., which have been somewhat louder affairs. And not without good reason. Such publicity inspires fear of Russian power and calls the legitimacy of democratic elections into question, which supports their objectives. Finally, the Kremlin might even welcome some hostility in order to bolster their narrative that Russia is under assault from a hostile West.

Financial motivation

Although Lazarus has a history of conducting strategic and political campaigns against North Korea's adversaries, much of its focus in recent years has been to generate cash. The profit motive may affect the efficacy of information disclosures because a group that is trying to make money would plausibly continue to operate until its costs exceed its revenue. It seems unlikely, then, that indictments would affect profit-motivated actors unless they are followed up by an arrest. Cobalt's ringleaders in Eastern Europe and APT38's technicians in North Korea have little to fear in that regard.

[3] Interview with an unnamed source.

Group Adaptability

Among the groups to continue operating despite disclosures, all increased in capability and technical prowess over time. This dynamic is circular and self-helping: heightened sophistication feeds into the ability to remain unaffected by public disclosures because these improvements manifest as the ability to adapt and modify infrastructure, TTPs, and toolsets that combine single-use, commodity, and bespoke tooling to suit different needs. The shift to continuous tool development seems to explain the effective difference of disclosures on numerous groups beginning around 2014, when public reporting generally stopped having serious effects on APTs. For instance, APT10 was disrupted by FireEye's first report in 2013, but later reports in 2015 and 2016—after it had diversified its tools and started the continuous development of new ones—had little effect. Lazarus' behavior also instantiates this trend.

The Human Factor

Many APTs have continued to operate in the face of numerous public disclosures because they exploit the human elements of the networks they attack. If network defenders make mistakes and do not implement proper controls on their networks, public disclosures do not matter because APTs can still gain access. Additionally, when members of targeted organizations are tricked by social engineering and unwittingly reveal credential or proprietary information, these groups are still able to access target networks after network defenders implement strong technical controls.

MEASURING THE SUCCESS OF DISCLOSURES

The success of a disclosure can ultimately be determined from its intent.[4] If the disclosures discussed in the case studies were intended to deter similar future behavior, they unequivocally failed. Each group continued to operate apart from the nominal exceptions of APT1 and APT10; and even then, other Chinese groups still target countries and companies globally. Were the intent to cause friction, then the hiatus seen in APT operations would suggest that disclosures were successful. But those were merely pauses, not definitive ends; and every group under consideration resumed operations in one capacity or another, generally not much longer than two months after a disclosure.

Naming-and-shaming disclosures that call out a group's low-level operators have been largely ineffective.

If the intent is to humiliate and shame the nation-state sponsoring or supporting the group, success is actor-specific. The reports and later indictments of Chinese groups embarrassed the Chinese government and should be considered successful, though not successful enough to convince the Chinese to halt operations completely. Naming-and-shaming disclosures that call out a group's low-level operators have been largely ineffective against non-Chinese groups, likely because public naming does not have a great effect on the least-senior technicians who merely carry out orders from above. In general, disclosures to embarrass have not been successful outside of the Chinese context. Iran is not embarrassed because its proxies give it plausible deniability. North Korea is beleaguered by financial sanctions and trade restrictions, so disclosures on its cyber activities carry little weight. Russia is politically focused and emboldened by the free publicity, excited to appear powerful to the rest of the world.[5]

While some disclosures may not have any meaningful impact on APT operations, their intent may be to add to the public record or to contribute to defensive cybersecurity.[6] They can be good for public relations, allowing the discloser to build relationships within the private sector and with the government. They may serve as proof that vendors or government entities are doing something to prevent cyberattacks, or to justify government sanctions on other governments, foreign entities, and individuals. Sanctions can be an important tool to influence other states' behavior, and they can be used to promote norms of conduct. Unlike denunciation and statements of attribution, sanctions must adhere to legal standards and present credible evidence of wrongdoing.

[4] As discussed in an interview with Max Smeets, Senior Researcher at the Center for Security Studies at ETH Zurich

[5] Interview with an unnamed source.

[6] Interview with an unnamed source.

Disclosures also make *post hoc* misrepresentation of APT operations difficult because they create a timeline of APT operations. Additionally, although these reports might not have any long-lasting effect on offensive operations against APT groups, they provide utility for network defenders, who can take the information released in the reports and use it to build stronger controls to keep adversaries out.

FOR PERSISTENT ENGAGEMENT AND FORWARD DEFENSE

Persistent engagement is the current U.S. Department of Defense strategy to counter adversaries in cyberspace. The strategy was designed to reflect the perspective that cyberspace has unique structural and operational characteristics that differentiate it from the other domains of conflict (land, air, and sea): namely, the characteristics of interconnectedness and of constant contact.[1] Operationally, persistent engagement has two objectives: first, to "disrupt or halt malicious cyber activity at its source," and second, to "provide public and private sector partners with indications and warnings (I&W) of malicious cyber activity."[2]

DoD's Cyber Command (USCYBERCOM) carries out persistent engagement by defending forward. Its cyber forces are in a constant state of contact in which they maneuver, outmaneuver, and react to adversaries who are also in their own constant state of maneuvering, outmaneuvering, and reacting. This means that Cyber Command is engaged with adversaries on their own networks, thus defending itself by acting in a space forward of its own perimeter. Additionally, Cyber Command engages in its own form of information disclosure, uploading malware indicators of compromise to websites like VirusTotal.[3] All of this action is intended to cause friction for adversaries, making it harder for them to target U.S. entities and interests, and occurs below the level of armed conflict.[4]

Persistent engagement is based on the concept of tacit bargaining. Explicit and overt bargaining between nations on norms and rules occurs through diplomacy and negotiations, in which all sides make clear what they want and then work together to reach a satisfactory agreement that is reflective of wants, needs, and power dynamics. Tacit bargaining theory states that all nations are self-interested and want to avoid conflict escalation, so they reach informal agreements on which actions are and are not acceptable by acting and reacting to each other. Over time, states' actions and reactions to others slowly delineate what they find to be acceptable behaviors. It is a process of trial and error. With persistent engagement, states

[1] 'Summary: Department of Defense Cyber Strategy' (2018), *Department of Defense*: https://media.defense.gov/2018/Sep/18/2002041658/-1/-1/1/CYBER_STRATEGY_SUMMARY_FINAL.PDF

[2] M. P. Fisherkeller & R. J. Harknett, 'What Is Agreed Competition in Cyberspace?' (19 February 2019), *Lawfare*: https://www.lawfareblog.com/what-agreed-competition-cyberspace

[3] See www.virustotal.com

[4] J. G. Schneider, 'Persistent Engagement: Foundation, Evolution, and Evaluation of a Strategy' (10 May 2019), *Lawfare*: https://www.lawfareblog.com/persistent-engagement-foundation-evolution-and-evaluation-strategy

communicate through their behavior in cyberspace as they maneuver and outmaneuver each other. The interactive process leads to all states having a tacit understanding of what each state will allow in cyberspace.[5]

In theory, information disclosures and indictments should complement persistent engagement and forward defense in three ways that are themselves complementary.[6] First, information disclosures by the U.S. government give private sector defenders more information on adversary operations by disclosing their actions, thereby giving them more chances to thwart adversary attempts to disrupt, degrade, destroy, or steal from their targets. Second, disclosures should create friction for adversaries, complicating their operations. By revealing their plans before they can be put into action, or by showing the world how they operate, these disclosures should slow down APT operations as they change their infrastructure, tools, and TTPs in order to not get caught. If information disclosures continue at a quick-enough pace, APTs may never get a chance to carry out a full attack. Indictments should play a role here too, serving as a form of formal information disclosure that also releases names and comes with diplomatic and legal consequences. Third, it contributes to the process of tacit bargaining as nation-states disclose behaviors of other actors that they find to be unacceptable.

Notably, declassified documents reveal that this is exactly how Cyber Command intends its public disclosures to work. They serve both offensive and defensive purposes. The documents note that Cyber Command's "objectives of VirusTotal uploads are to impose cost on adversar[ies]…and increase the resiliency of vulnerable networks." Both of these objectives are key to the strategy of persistent engagement.[7]

Our research suggests that public disclosures can be effective at achieving both of Cyber Command's objectives with respect to persistent engagement. Defenders benefit from any additional information they receive on adversary TTPs, tools, and indicators of compromise. Additionally, public disclosures do have some disruptive effect on APT operations. However, our research also suggests that some nuance is necessary. We find that the disruptive effect of public disclosures varies significantly based on a number of factors, including the scope of the disclosure and the disclosing actor. A simple public disclosure from a cyber threat intelligence firm or a posting of indicators of compromise VirusTotal is much less disruptive than an indictment. Further, the impact of indictments themselves is

[5] M. P. Fisherkeller & R. J. Harknett, 'Persistent Engagement and Tacit Bargaining: A Path Toward Constructing Norms in Cyberspace' (9 November, 2018), *Lawfare*: https://www.lawfareblog.com/persistent-engagement-and-tacit-bargaining-path-toward-constructing-norms-cyberspace

[6] Interview with Natasha Cohen, Senior Cyber Operations Planner and Strategist at DHS-CISA.

[7] J. Cox, 'Internal Docs Show Why the U.S. Military Publishes North Korean and Russian Malware' (25 February 2020), *Vice*: https://www.vice.com/en_us/article/5dmwyx/documents-how-cybercom-publishes-russian-north-korean-malware-virustotal

dependent on the APT in question. In most cases, regardless of disclosure type, an APT is only disrupted for a short period of time.

We must consider these questions of efficacy in direct relation to specific objectives. As mentioned, disclosures do have some disruptive efficacy and are thus somewhat useful in achieving the objectives of persistent engagement. The usually-short duration of this disruptive effect makes them much less useful for other purposes, such as deterring malicious cyber activity. The level of cost imposed by disclosure events is simply not high enough to change the risk calculus of adversaries conducting cyber activity below the level of armed conflict.

The impact of information disclosures on persistent engagement may also be country-dependent. Chinese APTs, for example, appears to be affected by highly public information disclosures from private vendors and by indictments, so information disclosure could be an effective part of a China-specific cyber strategy. APTs from other countries and contexts appear less affected, or unaffected, by information disclosures, implying that disclosures might not be an effective component of an offensive cyber strategy against them. In some cases, like Russia, disclosures might have the opposite effect and encourage further activity.

In most cases, regardless of disclosure type, an APT is only disrupted for a short period of time.

The efficacy of disclosures as a component of persistent engagement may be affected by both APT and network-defender sophistication. All of the groups outlined in our case studies are quite capable and sophisticated. There is a chance that information disclosure may be a more effective strategy against less capable groups that are unable to quickly adapt to new circumstances. Additionally, disclosures aimed at causing adversaries friction are less effective if network defenders do not take the time to implement new controls or update existing ones based on released information.

Ineffective information disclosures contribute very little to the concept of tacit bargaining, a key component of persistent engagement. While they do allow other nation-states and nonstate actors to gain a general sense of actions that the U.S. does and does not find acceptable, they do not appear to be able to stop those groups from acting. Tacit bargaining aims to outline what can and cannot be done with the intention that everyone then sticks to the set boundaries. If other actors in cyberspace do not respect those restrictions on action, the process of setting them cannot be considered a success. It is important to note here that public disclosures that are not consistent with, or backed up by, the disclosing state's behavior can therefore greatly undermine the norm-setting purpose of tacit bargaining. Thus, a state that indicts adversary personnel for economic espionage but then engages in economic espionage itself contradicts the very norms that it is trying to set.

Further, information disclosures on a single group at a time should not be considered a uniquely and independently effective component of persistent engagement, which is a nation-state strategy to counter other nation-states. Disclosures might be more effective if they involved coordinated releases of mass information on every known APT group an adversarial nation-state sponsors; otherwise, one group can make up for lost activity of another. In the case of Russian groups APT28 and APT29, both groups simultaneously targeted the DNC in 2016. Had a disclosure intended to get the Russians out of the DNC networks been operationally successful at halting the activity of one group, the presence of the other would render the disclosure a strategic failure. However, disclosures may be an effective tool to introduce friction within a state adversary's intelligence services if multiple APTs, particularly those serving different state agencies, are operating against the same target while unaware of their compatriots' activity. By revealing their activity to the others working the same target, disclosures can lead to infighting that pulls resources away from their operations.

With disclosures or indictments that reveal personas, the differing roles of compromised individuals might create differences in effect. It is worth exploring the potential differences in the disruptive effect of removing—or revealing the identities of—leadership versus operators, developers, or others. While the analogy is certainly not perfect, there has been a significant amount of research on the effect of leadership removal in the counterterrorism context. Austin Long, for example, argues that institutionalized organizations are resilient against leadership decapitation. He determines institutionalization based on "whether an organization exhibits functional specialization, hierarchy, and bureaucratic processes for conducting operations," [8] characteristics possessed by most APTs.

While institutionalized organizations should be more resilient to leadership removal, removing (or simply disrupting) key operational personnel may be more effective, particularly in a domain in which technical expertise is vital. Information disclosures that are a part of a broader strategy of persistent engagement should target each group's operational center of gravity. They should not be released without careful consideration of the cultural and political context in which the APT operates, and they should be designed specifically to maximize pain for that particular group. For example, numerous Chinese APTs might have centralized tool development, so disclosures targeting individual groups will be less effective than disclosures that target the activities of the high-level developers.[9]

[8] A. Long, 'Assessing the Success of Leadership Targeting' (November 2010), *CTC Sentinel*: https://ctc.usma.edu/assessing-the-success-of-leadership-targeting/

[9] 'Supply Chain Analysis: From Quartermaster to Sunshop' (November 2013), *FireEye*: https://www.fireeye.com/content/dam/fireeye-www/global/en/current-threats/pdfs/rpt-malware-supply-chain.pdf

FOR PRIVATE CYBERSECURITY VENDORS

From the standpoint of network defense, private vendors should continue to release information when necessary to boost the collective security of all. They should also explore information sharing channels that do not necessitate public disclosure.

However, private cybersecurity vendors that publicly disclose information as an offensive tool to cause pain to APTs need to consider context and intent. As discussed, the social, cultural, and political contexts in which APTs operate are complicated and vary by country, so disclosure efficacy will change for each group. Disclosures should be individually crafted to maximize friction and pain. Doing this requires deep non-technical understanding of each APT, so private vendors should retain non-technical talent with backgrounds in the social sciences. Additionally, vendors should consider the technological trends that reduce the value of malware disclosures over time, such as the use of commodity malware that groups can afford to lose and the use of social engineering, which cannot be stopped with signature-based threat detection.

Disclosures should be individually crafted to maximize friction and pain.

FOR THE FINANCIAL SECTOR

For the time being, the financial sector is not an offensive actor in cyberspace. Consequently, members of the financial sector should continue doing what they are already doing: monitoring public reporting and information disclosures on threat groups and working with private cybersecurity vendors as appropriate to defend their networks. When information is disclosed, they should promptly take the necessary steps to update their network defenses in line with cost and risk appetites. They should not grow complacent when a specific threat group is not targeting the financial sector, because capable APTs evolve and can shift their sights toward the financial sector in the future.

That being said, many private and commercial financial institutions maintain strong ties to governments around the world. They must work with regulators in the countries in which they operate to ensure the legality of their actions, and they sometimes collaborate with governments and central banks on finance and infrastructure projects. All financial institutions should consider aspects of their state collaboration that may make them a target for intrusive and disruptive attacks, both by the country with which they collaborate and by its adversaries. While financial institutions should meet regulators' disclosure requirements, they should also factor these considerations into their risk calculus when choosing to publicly reveal, or not to publicly reveal, attacks against their infrastructure and networks. Although such revelations may embarrass or stop the actions of some APTs, they could invite opportunist and retaliatory attacks by others.

ROOM FOR FURTHER RESEARCH

Our findings should be understood within the context of the scope of our project and our ability to measure APT activity. We measured activity through public disclosures. There were many periods for APTs when we did not record any ongoing campaigns, but that does not mean that these groups were not operating. Similarly, even in cases where we could not ascertain an impact resulting from a disclosure, groups may have still undertaken change at the command level.

More research with a larger and more representative sample is necessary to determine whether our findings hold true across larger sample sizes and to examine the causal mechanism behind our findings, as these outcomes are affected by the unique characteristics of each group itself, the country in which it is based, and the geopolitical situation in which it operates. Additionally, further research should include discussions with members of the financial sector on steps they might take to use information disclosures in an offensive manner in the future, should the sector choose to move in that direction.

ACKNOWLEDGEMENTS

The team would like to thank the many experts – Dmitri Alperovitch, Kris McConkey, Natasha Cohen, Max Smeets, Ben Read, Luke McNamara, Jenny Jun, and others that asked to remain anonymous – who agreed to be interviewed for this project and who shared and explained their unique insights into these numerous APT groups. This project could not have existed without the voluminous cyber threat intelligence collection and analysis conducted by the cybersecurity community over the years. Jason Healey's enthusiasm for and promotion of the project helped push us forward. Our greatest thanks go to our client sponsor Ken Wolf for making this project possible and our advisor, Neal Pollard, who provided valuable insight and suggestions throughout the research process.

ABOUT THE TEAM

Matthew Armelli

Matthew Armelli is a second year Master of International Affairs student at Columbia's School of International and Public Affairs. His studies focus on International Finance and Economic Policy as well as Cybersecurity. While attending SIPA, he interned on the Foreign Policy team at the Clinton Foundation and on the Online Security Team at Facebook. Prior to attending SIPA, he worked in Technical Security at the Cleveland Clinic after receiving his Bachelor's in German Studies at Miami University. Matthew was born and raised in Cleveland, Ohio.

Stuart Caudill

Stuart is a Master of International Affairs candidate studying international security and cyber policy. Prior to graduate school, Stuart spent over five years leading intelligence and cyber operations for the U.S. Army. Stuart co-authored the first public analysis of documents captured in the raid on Usama Bin Ladin's compound in Abbottabad, Pakistan, and his work on cyber policy has been published in *War on the Rocks* and *Strategic Studies Quarterly*. Stuart received a B.S. in International Relations and Arabic from the U.S. Military Academy at West Point and is a 2019 Tillman Scholar.

John Patrick Dees

John Patrick Dees studies International Security Policy and Data Analytics and Quantitative Analysis at Columbia University's School of International and Public Affairs. John received his commission into the United States Navy after graduating from George Washington University. He has served for nine years and was most recently stationed on board USS STETHEM (DDG 63) homeported out of Yokosuka, Japan.

Max Eager

Dr Eager is a candidate for the Master of International Affairs, with a concentration in international security policy and a specialization in business management. He has worked for the Information Security Division of the NYPD, and is presently Senior Researcher at GMG Strategic Advisers, a boutique technology and telecommunications consultancy based in NYC. Prior to SIPA, Dr Eager was College Tutor in Classics and Philosophy at Oriel College, Oxford. He holds a DPhil in Ancient History from the University of Oxford, and is the author of the monograph "Seneca's Influence with Nero" (Oxford, 2018). A native of New York City, Max has lived in Princeton, Rome, and Berlin. Currently he resides in Brooklyn with his partner Claire and chihuahua Banana.

Jennifer Keltz

Jennifer Keltz is about to graduate from SIPA with her Master of Public Administration. At SIPA, she studies International Security Policy and Technology, Media, and Communications. Immediately prior to graduate school, she was a Peace Corps Volunteer, first in Burkina Faso and then China. After finishing school this May, she will start working as a cyber risk consultant. She earned her undergraduate degree at the University of Virginia, where she majored in Government and minored in French.

Ian Pelekis

Ian is an accomplished geo-political analyst, formerly working with the Canadian government and a variety of think tanks, before deciding to focus on cybersecurity. Previously, Ian's research and work focused on proxy warfare and influence operations naturally leading Ian to the world of cybersecurity. During his time at SIPA, Ian has focused on cyber security, cyber conflict, and cyber threat intelligence while working with the New York City Cyber Task Force. Ian now works at Next Peak as a cybersecurity analyst, where he heads the creation of the Geo-Cyber Risk platform, a service combining geo-political and cyber risk to help identify and mitigate cyber risk.

John Sakellariadis

John Sakellariadis is a second-year Master of International Affairs candidate studying International Security Policy with a focus on cybersecurity. Prior to SIPA, John worked as a Research Assistant in U.S. Foreign Policy at the American Enterprise Institute. At SIPA, he serves as the Editor in Chief of the Journal of International Affairs, a student-run, peer-reviewed international affairs journal. Last summer, John interned as a Cyber Threat Intelligence Analyst at BlueVoyant. After graduation, John will begin a Fulbright research grant in Athens, Greece, where he will conduct research on the European Union's Agency for Cybersecurity (ENISA).

Virpratap Vikram Singh

Virpratap is a Master of International Affairs candidate focusing on cybersecurity and tech policy. Prior to graduate school, he worked as the digital media and content manager for Gateway House, a foreign policy think tank in India. His writing has been featured in *OODA Loop*. He holds a Bachelor in Liberal Arts from the Symbiosis School for Liberal Arts, majoring in Media Studies and minored in International Affairs.

Katherine von Ofenheim

Katherine von Ofenheim is a candidate for a Master of International Affairs, with a concentration in international security policy and specializations in technology, media, and communications and the Middle East region. Prior to SIPA, Katherine spent six years working in the humanitarian response and international development across the Middle East. She earned her undergraduate degree at the University of Oregon, where she majored in international affairs and geography. After graduation, Katherine will begin a Presidential Management Fellowship at the Department of State.

SIPA

Where the
World Connects