

KNOW YOUR ENEMIES 2.0

FEBRUARY 2016



A PRIMER ON ADVANCED PERSISTENT THREAT GROUPS
*THE MOST COMPLETE ENCYCLOPEDIA OF HACKTIVISTS,
NATION STATE AND MERCENARY HACKERS AVAILABLE*

AUTHORS:

JAMES SCOTT (ICIT SENIOR FELLOW - INSTITUTE FOR CRITICAL INFRASTRUCTURE TECHNOLOGY)

DREW SPANIEL (ICIT VISITING SCHOLAR, CARNEGIE MELLON UNIVERSITY)

Contents

Introduction:	4
China:	5
Blue Termite/ Cloudy Omega/ Emdivi	5
The Elderwood Platform.....	7
Axiom	9
Hidden Lynx / Aurora.....	11
Deep Panda / Black Vine / Pupa.....	13
PLA Unit 61398/ Comment Crew/ APT1.....	15
Putter Panda/ APT2/ PLA Unit 61486.....	16
Naikon / APT 30.....	17
Mirage.....	19
Iran:	20
Tarh Andishan/ Operation Cleaver	20
Ajax/ FLYING KITTEN/ Saffron Rose.....	23
South Korea	24
Dark Hotel/ Tapaoux/ Nemim/ Pioneer/ Karba	24
North Korea	27
Bureau 121/ Guardians of Peace/ Dark Seoul.....	27
Russia:	28
Energetic Bear/ Dragonfly/ Havex Crouching Yeti/ Koala Team	28
Uroburos / Epic Turla/ Snake / SnakeNet.....	31
APT 28/ Sofacy Group/ Sednit Group/ Tsar Team/ Fancy Bear/ Operation Pawnstorm.....	33
PinchDuke	34
GeminiDuke.....	35
CosmicDuke/ Tinybaron/ BotgenStudios/ NemesisGemina	36
MiniDuke.....	37
OnionDuke	37
APT29/ Hammertoss / HammerDuke.....	38
CozyDuke/ CozyCar/ CozyBear/ Office Monkeys/ Cozer/ EuroAPT	39
SeaDuke/ SeaDaddy/ SeaDask	40
CloudDuke/ MiniDionis/ CloudLook	41
Sandworm/ Quedagh/ BlackEnergy.....	42
Carbanak.....	43
Syria:	45
The Syrian Electronic Army (SEA).....	45
Global	47
Anonymous.....	47
America:	48

Butterfly Group/ Morpho.....	48
Regin/ Prax/ WarriorPride	50
Flame/ Flamer/ Skywiper	51
America’s Most Elite Line of Cyber-Defense: Tailored Access Operations (TAO).....	53
EQUATION Group	55
France	58
Animal Farm.....	58
Israel:	59
Duqu/ DQ	59
Unknown Nationality:	61
Hellsing.....	61
Moker	62
Shrouded Crossbow	63
Santa APT.....	64
Conclusion:.....	65
Appendix I:.....	67
Terms.....	67
Common Attack Vectors	69
Sources	71

Introduction:

Every system connected to the internet in every home, organization, and government entity is relentlessly subject to the attempts of malicious actors to steal their data or exploit their system. Cyber-attacks are prevalent in the digital age because computers (including mobile devices) are ubiquitous in society, because identification of an attacker and attribution is difficult, and because judicial rulings for cyber-crimes are nebulous. Most cyber-attacks are ` by basic security measures such as firewalls and antivirus applications. However, an elite percentile of the sea of cyber attackers is more persistent, more resourceful, and more sophisticated than the rest. These elite factions are known as Advanced Persistent Threats, and basic security measures are not enough to stop them from compromising some of the best-secured systems around the world.

Globally, at least a hundred advanced persistent threat groups are currently operational as criminal operations, mercenary groups, or nation-state sponsored divisions. Criminal operations typically target organizations or individuals for financial data or personal identifiable information for identity theft. Mercenary groups steal financial information or specific information from specific targets, as requested by their client. State sponsored groups may target organizations or governments to steal financial information, defense information, information that would grant a geopolitical economic or technological advantage, or any information that would be of use in intelligence or counterintelligence operations.

Nation state actors have been known to compromise enemy systems in order to plant malicious code that could enable the attacker to fully control the target networks or sabotage the systems altogether. While only limited application has been seen in the 2008 conflict between Russia and Georgia, the possibility of a joint cyber-physical war exists. An invading force can gain significant advantage over its enemy if it cripples their critical infrastructure prior to the attack. State actors such as Russia have entwined themselves within the networks of our Nation's critical infrastructure for surveillance purposes as they watch and wait for the appropriate time to cash in on their strategic positions.

Attribution of attacks is difficult. Researchers learn to think like attackers and to retrace the steps of a campaign. Security firms identify advanced persistent groups according to their tools, techniques, and characteristics. Advanced groups tend to develop and update their own sets of tools and malware. The keyboard language settings that remain in the code or the file names can reveal the attackers' nationality. Strings left in the code can reveal aspects of the development environment and the available resources. Who the group targets or does not target (especially in the case of nation state actors) can reveal the agenda of the

adversary. The infrastructure, domain, and network of attackers can reveal their location, resources, and sophistication. Finally, the specific information stolen and whether that data is sold, used, or stored can also assist in profiling an adversary.

Cyber-attacks deal significant damage to nations and citizens alike in the form of breaches and compromised systems. Most of the time, the adversary is an unknown phantom menace and security professionals analyze forensic data to attribute the crime to a specific actor. This primer aims to pull the veil from prominent actors and to assist in attribution attempts. This primer offers an introductory view of some of the most prolific advanced persistent threat groups in recent history. It also intends to give a top level view of the threat landscape and attack process. This primer is not a sole source of attribution information. Old actors falter about as quickly as new, identified groups emerge. Further, new groups often develop from or mimic their predecessors. It would be erroneous to rely upon a static document for attribution purposes. This primer cannot provide the whole view of an actor. Each actor is an extremely complex group. Even security firms rarely possess complete knowledge of an active adversary. Different firms identify different portions of actors and refer to them with different names. Malware can be copied or purchased by a new adversary. The adversaries update their malware and exploit kits, along with their attack vectors in order to stay ahead of detection efforts. Finally, different organizations and governments evaluate and fear APT groups differently. As a result, this primer will not serve as a ranking system just as it does not serve as a comprehensive list of all attackers. Instead, the information below is offered to raise awareness about the groups that populate the cyber landscape. We have categorized these APT groups by country so that the reader can more easily identify the characteristic similarities of the attack components.

China:

Blue Termite/ Cloudy Omega/ Emdivi

The Blue Termite malware campaign has targeted hundreds of Japanese organizations since its inception in 2011. According to Kaspersky, the malware is Chinese in origin. The C&C infrastructure is located in Japan, the primary target of the campaign. In a November 2014 report, Symantec indicated that the group might share communication channels or attack infrastructure with the Hidden Lynx APT group.

Over four years, the malware has stolen confidential information from government agencies, universities, public interest groups, financial institutions, media organizations, automotive companies, chemical organizations, healthcare firms, electrical companies, real

estate firms, technology firms, and other critical infrastructure organizations. The majority of the targets were based or located in Japan. Blue Termite is also allegedly responsible for compromising the personal data of 1.25 million Japanese citizens in a breach of the Japan Pension Service.

Initially, like most malware groups, Blue Termite relied on phishing campaigns to spread its malware. For instance, in 2013, it spread malicious emails relating to the Ichitaro product line. The content of the emails varied according to the target organization; however, many focused on political events. Opening the email attachment deployed a malicious payload. Usually, the attachment was an executable with a fake icon. Occasionally, the attachment would contain code to exploit a vulnerability in specific target software instead. A notable characteristic of the payload is that often the lure document would open and then the document reader would crash before reopening a clean document. The malware would be delivered at the time of the crash and the reopened document would no longer carry the malware.

Like Sofacy and many other threat actors, the activity of the group increased in July 2015 in response to the breach of the Hacking Team servers and the public disclosure of a number of valuable 0-day exploits and system vulnerabilities. In particular, the group began to use a Flash player exploit (CVE-2015-5119) to conduct drive-by-download malware attacks from compromised Japanese malware sites. The group altered its behavior to target individuals as well as organizations. The group also conducted watering hole attacks intent on infecting systems belonging to prominent members of the Japanese government. In other attacks, infected sites were configured to only infect visitors whose IP addresses belonged to target organizations.

Blue Termite's attack kit relies on the Emdivi family of malware. The group uses Backdoor.Emdivi, Backdoor.Korplug, and Backdoor.ZXshell to compromise a system and establish a persistence presence. The backdoor enables a remote adversary to execute commands from a C&C server via HTTP. The malware contains components to search files, delete files, upload files to C2 servers, execute code, acquire a list of running processes, steal auto-complete information and saved credential information from Internet Explorer, and steal the proxy settings of browsers such as Mozilla Firefox. Kaspersky noted that "One of the most interesting things about the malware used by the Blue Termite actor is that each victim is supplied with a unique malware sample that is made in a way that it could only be launched on a specific PC, targeted by the Blue Termite actor." Each variant has a unique version number and a type (Type S or Type T). The version number indicates that the group systematically deploys the malware as part of an organized campaign. The version and extra words are also used to generate a hash, which is used as an encryption key. Both types allow

the adversary to remotely execute code and to steal credentials stored in Internet Explorer. Both variants also share the same hardcoded C&C infrastructure. Type T, the more prevalent variant, is written in C++. Type T encrypts its C&C address and detection protection mechanisms as an anti-analysis technique to prevent debugging and analysis in virtual machines or sandbox environments. Type S is a .NET application that is based on Type T. Type S lacks the encryption and anti-analysis mechanisms. Type S also changes its file hash between versions by relying on Japanese sentences taken from the internet.

The Elderwood Platform

The Elderwood Platform is the name given to a set of zero-day exploits that is either used within a large organization or sold as a package to many attackers. The Elderwood platform was discovered by Symantec in 2009-2012, following the actor's 2009 compromise of Google with the Hydra (Aurora) Trojan. It is not clear whether Elderwood is a single criminal group that distributes its platform or if it is part of a major organization that distributes its platform to its subdivisions. In the former scenario, the Elderwood distributor may preferentially sell its platform to separate criminal entities at the same time. In most cases, the "buyers" receive the exploit around the same time. This could be an operational choice on behalf of the seller, a systematic choice (i.e. the "seller" sells once they find an exploit), or a procedure meant to obfuscate the activities of any one "buyer." In the latter scenario, Symantec theorizes that a parent organization may distribute the platform and it may task its subdivisions with targeting particular industries or sectors. Each subgroup then utilizes their own infrastructure to stage the attacks using the shared platform.

Zero-day exploits are rare and valuable and the Elderwood platform relies upon zero-day exploits to compromise its victims. Somehow the Elderwood platform has consistently been updated with new zero-day exploits since 2009. In fact, no other actor has been able to obtain and utilize as many zero-day exploits as the actor behind the Elderwood platform. This suggests that either the actor behind the Elderwood platform has a highly sophisticated technical team that is capable of farming zero-day exploits or that Elderwood is funded by a criminal organization or state sponsor that possess significant resources. Unless the technical team that farms the exploits is paid an extremely high sum, neither theory explains why the exploits do not appear on underground markets until long after Elderwood has used the exploit.

A hybrid theory is possible. Perhaps the Elderwood group sells their platform to a third party for one reason or another, and that party then resells the platform to smaller groups. The hybrid model could explain how the Elderwood platform continues to utilize new

and unique zero-day exploits because an exploit could be sold whenever the group feels that it has served its purpose and then they can purchase new exploits using the money received from selling the previous exploit to numerous other buyers. Alternately, perhaps a simpler solution exists. Zero-day exploits are juicy pieces of information. It is possible that Elderwood activity attracts the notice of other groups who watch the attacks and reverse engineer the exploits. The lower tier attackers would need inside knowledge of Elderwood activity and they would have to outpace cybersecurity response teams, else the exploits would be of little value.

In recent years, other notable campaigns have utilized the Elderwood platform or its exploits. Hidden Lynx used internet explorer exploits and its ZXshell backdoor in attacks against the defense industry. Vidgrab exploited internet explorer to install the vidgrab backdoor on systems belonging to Japanese users and it exploited Adobe flash to install the Jolob backdoor on systems belonging to Uyghur dissidents. Icefog exploited both Adobe Flash and Internet Explorer to install the Linfo and Hormesu backdoors respectively on systems in the manufacturing industry. Sakurel used multiple Internet Explorer exploits and an Adobe Flash exploit to compromise Aerospace engine manufacturer systems with the Sakurel Trojan.

The Elderwood platform is used against targets in a large number of sectors. Most frequently, the Elderwood platform is employed against organizations involved in defense, defense supply chain manufacturing, IT, and Human Rights. Organizations are attacked through watering hole attacks, spear phishing emails, and web exploits. Symantec believes that it is possible that manufacturers and tangential sector organizations and sites are compromised to target top tier primary targets. In this case, organizations in Manufacturing, Engineering, Electronic, Energy, Arms, Shipping or Aeronautics industries may be targeted as stepping stones to compromise Defense organizations. Additionally, Software or Financial firms might be targeted so the attacker can compromise NGOs.

The Elderwood platform predominantly targets United States organizations. Firms in Canada, China, Hong Kong, and Australia have also been frequently targeted. Organizations based in Taiwan, United Kingdom, Switzerland, India, and Denmark have been sporadically targeted. Victims are targeted for information and intellectual property contained on their systems. The lack of theft of financial information complicates the actor profile because a mercenary distributor is less likely to steal nation-state information over financial information. One could argue that information is stolen with the Elderwood platform to assist in other breaches; however, a mercenary group would likely not be able to analyze information as rapidly as had the Elderwood group.

Between 2009 and 2014, the Elderwood platform has featured numerous Adobe Flash and Internet Explorer zero-day exploits. Adobe Flash and Internet Explorer are notoriously vulnerable applications. Typically, Adobe Flash, Internet Explorer, or both are present on a system. The attack platform also contains a document creation kit which enables the attacker to combine a clean document with a Trojan of their choice to create a malicious document. These documents are then used in spear phishing campaigns. The platform also contains a Shockwave Flash file that ensures that Trojans are downloaded onto target machines in the correct locations. The platform could contain information gathering tools such as keyloggers, automated domain name and account generators, and an information analysis platform.

Axiom

The Axiom group is a Chinese, potentially state-sponsored, threat actor that compromises systems that contain information of value to advancing China's 12th Five Year Plan. Axiom was investigated in the October 2014 Operation SMN, a joint operation between private firms, led by Novetta which released information and led to the removal of Axiom malware from over 43,000 systems.

Since 2009, Axiom has been targeting networks in a broad range of sectors who possess confidential or classified information. Axiom campaigns share infrastructure, malware, or attack techniques with Operation Aurora (2009), the Elderwood Project (2009-2014), the VOHO campaign (2012), the Shell_Crew attacks on ColdFusion servers (2013), Operation Ephemeral Hydra (2013), Operation Snowman (2014), and 2014 attacks on American Middle Eastern Policy think tanks. Axiom could be connected to some of these other groups; however, it is more likely that Axiom advantageously adopts zero-day exploits or malware that are effective in other campaigns. It is possible that Axiom acquires its malware on deepnet or through underground trade.

Axiom is likely Chinese state sponsored, but there are no definitive links connecting it to the Third Department, which houses China's offensive threat groups Putter Panda and APT1. Axiom malware was configured to use simplified Chinese language settings and some of the filenames are in Chinese. Axiom is more sophisticated in its operations than the aforementioned Third Department groups. It utilizes different resources, and it may have a different mission than Third Department groups. Novetta hypothesizes that based on Axiom's domestic monitoring trends that it might be charged with domestic operations and targeting Chinese dissidents in other countries. Universities and research institutions in Hong Kong and mainland China have been targeted with Hikit malware for persistent operations. This could indicate state-sponsored concern over liberal academics and students.

Novetta has found that Axiom targets a wide variety of entities inside and outside governments. Axiom targets a wide variety of sectors, but it only targets specific entities in those sectors. Within Asian and Western governments, Axiom targets law enforcement, governmental records and communication agencies, environmental policy agencies, personnel management divisions, space and aerospace exploration and research entities, government auditing and internal affairs divisions. In the science and technology sectors, Axiom targets networks belonging to electronics and integrated circuitry manufacturers, networking equipment manufacturers, internet based service companies, software vendors, cloud computing companies, energy firms, meteorological service companies, telecommunications firms, and pharmaceutical companies. Additionally, Axiom has targeted journalism and media outlets, Human Rights NGOs, international law firms, international consulting and analysis firms, and high ranking United States academic institutions. Most of the target's organizations have been located in the United States, South Korea, Taiwan, Japan, and the European Union, with a majority of the breaches along the Eastern seaboard of the United States and Western Europe.

Axiom targeting coincides with interests reflected in China's 2006 and 2011 Five Year Plans, which push for advanced technology and advanced R&D efforts. As China shifts away from foreign technology, more organizations may be targeted by Axiom. The actor may target semiconductor and networking technology firms with offices in China because China wants to reduce its dependency on foreign technology. Western and Asian organizations may be targeted in intelligence and counterintelligence operations. Axiom targets NGOs concerned with international politics, environmental policy, pro-democracy movements, or human rights movements. In some instances, Axiom will target a satellite office and move laterally through the compromised network to the main office. Novetta theorizes that Axiom targets NGOs as a means of the Chinese ruling party keeping track of watchdog organizations and other groups who may publish claims that challenge the authority or "soft power" of the party. Targeting NGOs may also enable the party to suppress dissidents or intimidate whistleblowers.

Novetta believes that Axiom has a six stage victim lifecycle that uses a different team for each stage of the attack. This indicates large scale organization and coordination. Initially, the target is identified and the actor conducts reconnaissance. Then the system is compromised, confirmed to be a valuable target, and the network is surveyed. The actor laterally moves through the network and creates additional footholds. Compromised C2 infrastructure is connected to the victim network. Finally, valuable data is identified and exfiltrated.

Axiom initially compromises systems through web based attacks, targeted attacks against public facing infrastructure, zero-day exploits, watering hole attacks and phishing emails. Once a system is compromised, Axiom spends a few days determining whether it is valuable. If it is determined to contain useful information, then the group installs persistent malware platforms. Otherwise, the group tries to move laterally through the network to locate more valuable systems. Axiom has proven capable of compromising large pools of machines and sifting through them in hours or days to find the valuable ones. This indicates dedicated resources, possibly a dedicated targeting team and a deterministic set of criteria. After the initial compromise, Axiom begins reconnaissance to identify where they are in the target network and to identify any changes that have been made to the network. Axiom then escalates privileges using previously compromised administrative accounts, local exploits, or remote exploits as demonstrated in ZoxRPC malware. Then, over the course of minutes or months, they try to dump the latest user credentials and exfiltrate the data. Once inside the network, Axiom can also exploit Remote Desktop Protocol or exploit vulnerabilities in the custom tools designed by the organization itself. This allows Axiom to “fly under the radar” and not alert antivirus or IDS systems to the compromise.

As the campaign continues, Axiom may install additional families of malware as a mechanism of remaining in the system even if one malware is discovered by the target. Compromised systems have featured up to four layers of malware ranging from extremely common (Poison Ivy, Gh0st, ZXshell) to focused tools used by threat groups (Derusbi, Fexel) to custom Axiom malware (ZoxPNG/ZoxRPC, Hikit). Axiom routes its activity through compromised proxy infrastructure in the United States, South Korea, Taiwan, Hong Kong, and Japan to try to disguise its traffic as legitimate to casual observation.

Novetta observes that the Hikit malware is unique to Axiom and is only used on high value targets at the height of the victim’s operational lifecycle. Of the 43,000 compromised systems discovered in Operation SMN, only 180 systems were infected with the Hikit malware. Hikit is a late stage persistence and data exfiltration tool that is capable of uploading and downloading files, generating a remote shell, tunneling into the network, and connecting to other infected machines to generate a secondary network.

Hidden Lynx / Aurora

Hidden Lynx is a professional “hackers for hire” group that has operated since 2009 and that is believed to be based out of China. Hidden Lynx steals specific information from select targets from a wide range of sectors and governments. The 50-100 member group has proven themselves capable of breaching some of the best defended systems in the world.

The adversary can conduct multiple persistent campaigns concurrently against a variety of well defended targets. Hidden Lynx has been associated with 2010 Operation Aurora and the 2012 VOHO campaign.

In the past three years, Hidden Lynx has conducted hundreds of attacks against commercial organizations and governments across the globe. The sectors most targeted are the financial sector, the education sector, and government entities. Within the financial sector, investment banks and asset management agencies are the primary targets. In their 2013 report on the group, Symantec points out that “[t]he absence of certain types of financial institutions, such as those operating as commercial banks, clearly indicates that the attacks are focusing on specific areas.” With less frequency, the group has also targeted stock trading firms and indirectly attacked organizations that supply hardware, secure network communications, and specific services to the financial sector. Overall, the targets share the characteristics of possessing valuable information such as confidential financial data, specific knowledge of potential mergers or acquisitions, or other information that could give the client of the attacker a competitive advantage in the sector or specific knowledge of ongoing negotiations or business deals.

Outside of the financial sector, Hidden Lynx largely targets all levels of government and government contractors. Exfiltrated information from the defense industry sector or from an opposing government could grant a nation state the ability to close a technological gap or the ability to gear intelligence and counterintelligence efforts towards a specific country. Alternately, the information could allow private organizations to spy on competitors or to gain unfair competitive advantage by speculating on government technological research and interest. Microsoft claims that during Operation Aurora Hidden Lynx targeted databases containing court order emails.

Over half of Hidden Lynx attacks target United States organizations, while another quarter of the attacks target organizations in Taiwan or China. The broad range of targets accompanied by the specificity of the information targeted indicates the mercenary nature of the attacker. The information stolen is not processed by the attacker or used for direct financial gain, so it is likely that the information is stolen on behalf of a third party. The stolen information, predominately financial or technological in nature, would be valuable to corporations and nation states alike.

Hidden Lynx targets organizations and government entities in wealthy and technologically advanced countries. Most of the Lynx attacks originate from infrastructure located in China. The group initiates campaigns with a two pronged approach. Hidden Lynx usually infects compromised systems with multiple Trojans, a mass exploitation Trojan

(Trojan Moudoor) and a targeted Trojan (Trojan Naid). Each Trojan may be managed by a different team. Trojan Moudoor deploys the Moudoor backdoor, which is a modified version of the “Gh0st RAT” malware. The remote access Trojan is used to control machines in significant campaigns against multiple large companies across several sectors. The Moudoor team must be sizable because the attack vector requires attackers to breach individual targets and to extract valuable and specific data from compromised networks. Trojan Naid is used in limited attacks against valuable targets. Given its limited use and the sophistication of its application, each team behind it is likely a highly skilled special operations team within the overall group. In recent years, Hidden Lynx added the Gresim backdoor, the Fexel backdoor, the Hikit backdoor, and the Derusbi malware to their exploit kit.

The adversary regularly exploits zero-day vulnerabilities, which are purchased, discovered, or reworked from other groups’ attacks. Ultimately, Hidden Lynx is methodical and it tailors its exploit kit in each attack to its victim. Hidden Lynx adapts and it will develop custom tools or perfect new techniques if necessary. Most attacks begin as a watering hole attack or a spear phishing email; however, Hidden Lynx has also been known to attack public facing infrastructure or hack the supply chain in order to distribute their malware.

Deep Panda / Black Vine / Pupa

Deep Panda began attacking the healthcare, aerospace, and energy sectors around 2012. Deep Panda is believed to be a Chinese state sponsored group. Symantec believes that Black Vine may be affiliated with a Beijing IT security organization called Topsec. Topsec is a research institute with sites across China. Topsec focuses on information security research, training, auditing, and security products. It also hosts a hacking competition (from which they hire hackers). It is possible that some members of Topsec are affiliated with Deep Panda.

Deep Panda attacks tend to have massive impacts and they accrue proportional media attention. In order to conduct multiple sizable campaigns against United States Federal government agencies and major western health care providers for extended time periods, Deep Panda must have considerable resources at their disposal. In illustration, it is possible that Deep Panda was concurrently engaged in cyber-attacks against the United States Office of Personnel Management, the Anthem healthcare network, United Airlines, and other entities. In December 2015, the Chinese government announced that it had arrested the actors behind the OPM breach and that Deep Panda was not responsible. Many in the political and cybersecurity spheres remain skeptical that the arrests are legitimate.

Deep Panda conducts watering hole attacks; zero-day exploits, and spear phishing campaigns. The group also utilizes some of the exploits and tools from the Elderwood platform. A vast majority, ~80%, of Deep Panda targets are American. Deep Panda targets government agencies, the aerospace sector, the healthcare sector, financial organizations, technology firms, and energy entities (primarily gas and electric manufacturers).

In the United States health care sector, Deep Panda has attacked VAE, Anthem, Empire Blue Cross Blue Shield, and Carefirst. In the recent 2014-2015 Anthem breach, the group exfiltrated ~80 million patient records. Information exfiltrated from Anthem includes social security numbers and other personal identifiable information or personal health information. It is believed that the Axiom group also attacked Anthem at the same time as Deep Panda, but with a different malware and along different vectors. The attack appears as a coordinated effort. Further, enough similarities exist between the meticulous planning and malware employed by the two groups, that many security firms hypothesize that they are both part of the same group. There is a strong possibility that the groups are affiliated.

Deep Panda is also believed to be responsible for the two 2015 OPM breaches. The breaches resulted in the exposure of the personal information contained in the SF-86 forms of 22.1 million current and former United States Federal employees. 5.6 million fingerprint files were also stolen. Deep Panda breached United Airlines in 2015 and stole departure and destination records. The health, OPM, and travel records stolen by Deep Panda can be aggregated to catastrophically impact the United States government over time. The adversary or their parent nation state can build a database of US employees for espionage purposes. Further, the information can be used to identify United States agents in the country or to identify Chinese assets who assist United States intelligence efforts. Even though their systems were not compromised and their agents' information was not included in the breach, the CIA has already begun retracting agents from the field in response to the cyber-attacks. The CIA made this decision because State Department records were stolen in the breach and the attacker could thereby discover embassy employees who were not included in the State Department records and capture those individuals as spies or coerce their behavior. In this manner, Deep Panda has pushed forward the boundaries of cyber-warfare to achieve a measurable "physical" nation-state response. Further, physical warfare has been suggested in the United States in response to the cyber-attacks.

Deep Panda relies on the Sakurel Trojan, the Hurix Trojan, and the Mivast backdoor in its attacks. Deep Panda is believed to have developed all of the malware themselves. Characteristics in the malware code are shared between all three malware. Further, each malware is capable of opening a named pipe back door and contains tools to collect and

exfiltrate system data, the ability to execute arbitrary code, the ability to create, modify, and delete registry keys.

The malwares are similar in that they utilized droppers that masquerade as installers for legitimate software applications like Adobe Reader, Juniper VPN, and Microsoft ActiveX Control. In some cases, a loading bar displays and then the user redirects to a login page for the associated software. The malwares contain measures to avoid detection. The malwares self-obfuscate as technology related applications such as media applications or VPN technologies. The malwares establish persistent presence on the system, deploy remote access Trojans (RATs) such as the Derusbi malware, and feature tools to record and seize user sessions. Tools such as PwDump and Scanline are included to steal user credentials, to allow the actor to escalate their privileges, to let the actor create unmonitored accounts, and to assist the attacker in lateral movements to systems across the network. Symantec believes that all three malware belong to the same family and that they have been updated and differentially developed over time by the same team. The malware is usually signed by the DTOPTOOLZ Co. signature belonging to a Korean software company. Domains and C2 servers often feature the names of Marvel comic book characters as the register.

PLA Unit 61398/ Comment Crew/ APT1

The 3rd and 4th Departments of the People's Liberation Army (PLA) General Staff Department (GSD) supposedly houses China's electronic warfare operations. Unit 61398 is the Military Unit Cover Designator of the Chinese state sponsored advanced persistent threat that operates out of the 2nd Bureau of the 3rd Department of PLA GSD, located off Datong road in Pudong in Shanghai.

Unit 61398 is tasked with computer network operations. It operates on four large networks in Shanghai. Two of these networks serve the Pudong region. The Unit has a dedicated fiber optics connection that was paid for in the name of national defense. The 3rd Department employs over 130,000 employees. Unit 61398 consists of personnel who are proficient in English and trained in computer security and computer network operations. Members of Unit 61398 use Chinese (Simplified) keyboard settings. Most of the IP addresses and the infrastructure used in the attacks trace back to China.

Unit 61398 targets sectors that are of interest to China's 12th Five Year Plan. They are large enough and well-resourced enough that it can simultaneously compromise dozens of organizations. This adversary has breached over 150 organizations since its inception in

2002. The majority of victims are located in the United States. Information Technology organizations, Aerospace firms, Public Administration agencies and other technology heavy sector are targets for Unit 61398. The adversary targets intellectual property data and financial data. It exfiltrates intellectual property data, proprietary documents, business plans, emails, and contacts.

Attacks begin with spear phishing emails that contain a malicious file or a malicious link. The emails are personalized to the target and may not easily be distinguished from legitimate emails. Attachments are usually in the ZIP format. Once the victim system is compromised, the attacker establishes a persistent presence by installing a backdoor from the dropper delivered from the email. The backdoor initiates contact with the C2C infrastructure from inside the network so that the traffic can bypass internal firewalls. The actor typically relies upon WEBC2 backdoors, which are minimally featured beachhead backdoors. WEBC2 can only communicate with a C2C server through comments. Sometimes the BISCUIT backdoor is used if more functionality is needed. BISCUIT uses the HTTP protocol for communication and it features modules to capture screenshots, log keystrokes, record system information, modify processes, modify the registry, execute code, log off or shut down the session, and other features. Unit 61398 remains persistent on the compromised system and it may revisit the system over the course of months or years. The group remains on the network for 1-5 years. During this time, the group escalates their privileges using login credentials that it gathers from publicly available tools built into the initial malware. Next, they conduct network reconnaissance, by typing commands into the command shell. Finally, they laterally move across the network to infect new systems and they maintain their presence on the infected network. Unit 61398 compresses stolen data into multiple files with a RAR archiving utility and exfiltrates the data through their backdoor or through File Transfer Protocol (FTP).

Putter Panda/ APT2/ PLA Unit 61486

Putter Panda is connected to the People's Liberation Army's (PLA) Third General Staff Department (GSD) 12th Bureau Military Unit Cover Designator (MUCD) 61486. Unit 61486 supports China's space surveillance network. The group may be responsible for space based signal intelligence (SIGINT) collection. The group has been actively conducting attacks since at least 2007 and is based out of the Zhabai district of Shanghai, China. Unit 61486 shares some infrastructure with Unit 61398.

Putter Panda targets the United States Government, Defense sector, Research sector, and Technologies sectors. According to CrowdStrike, the United States Defense

industry, communication industries, and European satellite and aerospace industries are particularly targeted.

Putter Panda relies on spear phishing emails containing malicious PDFs and Microsoft Word Documents to infect its target. Putter Panda's exploit kit includes two droppers, two RATs, and two tools. One dropper delivers a payload, such as the 4H RAT, to the victim system and installs it. The other dropper exclusively delivers the PNGDOWNER tool. Putter Panda uses the 4H RAT and the 3PARA RAT. The 4H RAT can initiate a remote shell, enumerate running processes, terminate processes, list files and directories, modify timestamps, upload files, download files, and delete files. The RAT communicates over HTTP and the communication is obfuscated by an operation, 1-byte XOR with the key 0xBE. The 3PARA RAT is a second stage, failsafe tool that allows the attacker to regain control of the system if their initial access vector is removed. The 3PARA RAT creates a file map at startup to verify that there is not another instance of the RAT running. The RAT is capable of remaining dormant for prearranged or commanded periods of time. The RAT only has limited commands, which include retrieving file or disk metadata, changing the working directory of the current C2 session, executing a command, and listing the current working directory. The first tool, PNGDOWNER is a simple download and execute tool. The second tool, HTTPCLIENT is a backup tool. The 3PARA RAT communicates over HTTP and authenticates with a 256-byte hash and a hard-coded string.

Naikon / APT 30

The Naikon group is one of the most active APT groups in Asia. Since 2010, it has launched spear phishing campaigns into organizations surrounding the South China Sea, intent on harvesting geo-political intelligence from civilian and military government organizations in the Philippines, Malaysia, Cambodia, Indonesia, Vietnam, Myanmar, Singapore, Nepal, Thailand, Laos, and China. The actors speak native Chinese. Based on the choice of targets, the operating language, and the sophistication of the toolkit, there is a distinct possibility that Naikon is a Chinese state sponsored threat group.

Spear phishing campaigns begin with a lure email relevant to the victim that carries a malicious Microsoft Word document, which, according to Kaspersky Lab, actually contains "a CVE-2012-0158 exploit, an executable with a double extension, or an executable with an RTLO filename". One of its most prolific spear phishing campaigns was the March 2014 attacks targeting organizations from countries affected by the MH370 tragedy. Upon opening/execution, the malicious payload, an 8kb encrypted file and configuration data, is injected into the browser memory where it decrypts the ports and paths to the C2C server, a user agent

string, filenames and paths to relevant components, and hash sums of the user API functions. The malicious code downloads the main malware from the C2C server over an SSL connection and then it loads it independently of the operating system functions without saving it to the hard drive by assuming control of the XS02 function and then handling the installation in memory.

The main component of the Naikon platform is a remote administration component. According to Trend Micro, the RARSTONE backdoor (BKDR_RARSTONE.A) can obfuscate itself by “decrypting and loading a backdoor ‘executable file’ directly into memory without the need to drop the actual ‘executable file.’” The backdoor installs like a Plug X backdoor, injecting code into hidden instances of internet explorer. The module establishes a connection to the C2C server to receive and execute any of an estimated 48 commands from the adversary on the host. These commands include profiling the system, uploading and downloading data, executing arbitrary code, installing other modules, or executing commands via the command line. The backdoor routine also has the ability to get installer properties from Uninstall Registry key entries, which allow it to silently uninstall applications that interfere with the malware. The espionage malware collects email messages, monitors victims’ keystrokes and screens in real time, and monitors network traffic.

The command and control infrastructure is minimalistic and organized according to locations of victims and targets. Communication protocol varied according to target. Some systems connected directly to the C&C servers while other systems were routed through dedicated proxy servers. The proxy servers were victim hosts running the XSControl software, which accepted incoming connections and routed them to relevant C&C servers. The proxy server application also offered a GUI administration utility, logged client and operator activity, and transmitted logs to an FTP server. The operator logs contained an XML database of downloaded files (including a timestamp, the remote path and the local path), a database of filenames and victim registry keys, and a history of executed commands. Perhaps the largest news story involving Naikon was the report by Kaspersky Lab that a rival APT in the region, dubbed the Hellsing group, had attacked the Naikon group. In March 2014, Hellsing group received a spear phishing email from Naikon and Hellsing responded with a reply message containing a locked malicious RAR archive labeled “confidential data.” The archive contained two PDFs and a SCR file, a backdoor specifically customized to target the Naikon group. The backdoor can upload and download files, update itself, and uninstall itself.

Mirage

The April 2012 Mirage campaign targeted a high profile oil company in the Philippines, a military organization in Taiwan, an energy company in Canada, and organizations in Brazil, Israel, Egypt, and Nigeria. The Mirage attacks are attributed to the Chinese government or a state sponsored threat actor. The campaign was investigated while advanced persistent threat groups were still developing into their current structure; consequently, the campaign was not investigated to the same level of detail as modern threats. The most distinct commonality between victims was that all parties were involved in the contest for rights to survey natural gas and oil in the South China Sea. It is believed that the intent of the campaign was to exfiltrate confidential information, steal intellectual property, or to construct a botnet.

The actors began the campaign by targeting mid-level to senior-level executives with spear phishing emails that contain malicious droppers that install the Mirage malware. The droppers are disguised as PDF attachments. If opened, then the dropper is deployed and an embedded PDF of a news story, relevant to the target, opens. The dropper contains a copy of the Mirage malware, which executes and copies itself into either C:\ Documents or C:\ Windows. The copy starts and the original closes. The new Mirage establishes persistence in the event of reboot by creating registry keys. The malware obfuscates its presence through the creation of one or more files named svchost.exe, ernel32.dll, thumb.db, csrss.exe, Reader_SL.exe, and MSN.exe. The malware profiles the system (MAC address, CPU speed, memory size, system name, and user name) and sends the information back to the command and control infrastructure via a HTTP request over ports 80, 443, and 8080. It can implement SSL for added security. The first variant of Mirage communicated via a HTTP POST request and it transferred information that was lightly encrypted by adding each character's ASCII value to its offset from the start of the payload. The second variant of the malware communicated through HTTP GET requests and it encrypted data the same way as the former version except that the payload of the initial request is encapsulated in a Base64-encoded string. The Mirage toolkit consisted of a backdoor and a remote access trojan (RAT). At the time of its discovery in 2012, the command and control structure consisted of over 100 domains. By the end of 2012, the Mirage campaign went dormant. However, some of its infrastructure reappeared in the 2015 Hellsing campaign.

Iran:

Tarh Andishan/ Operation Cleaver

In April 2010, a worm called Stuxnet, allegedly jointly developed by the United States and Israel, targeted Siemens industrial control systems (ICS) in developing nations such as Iran (~59%), Indonesia (~18%), and India (~8%). Stuxnet contained a programmable logic controller (PLC) rootkit designed to spy upon, subvert, and in some cases sabotage Siemens supervisory control and data acquisition (SCADA) systems that regulated specific industrial systems. In particular, Stuxnet variants were deployed by a nation state actor against Iranian industrial facilities associated with its nuclear program, such as uranium enrichment facilities. The Stuxnet infection was discovered three months later, but variants continued to compromise Iranian systems through 2012. Iran's nuclear infrastructure and its oil and gas infrastructure was also targeted by the Duqu malware from 2009-2011, and the Flame malware in 2012. In response to adversarial cyber campaigns, Iran began rapidly developing its cyber infrastructure.

In December 2014, ICIT Fellow Cylance exposed Iranian threat actor, Tarh Andishan in the white paper of their 2-year Operation Cleaver investigation. Tarh Andishan was likely developed in response to the Stuxnet, Duqu, and Flame campaigns. Iran could be demonstrating to global targets that it is a major cyber power, capable of competing with countries such as the United States, China, and Russia, on the global cyber landscape. Cylance released Operation Cleaver early to allow potential targets the opportunity to mitigate the threat to their systems, so they estimate that they only discovered a portion of the activity of Tarh Andishan. Nevertheless, Cylance managed to build an impressive profile of Tarh Andishan's operation, including hacker profiles, domain names, internal infrastructure, and indicators of compromise.

The infrastructure used to host the attacks belonged to the corporate entity Tarh Andishan in Iran, after which the threat group is named. The infrastructure was hosted by an Iranian provider (Netafraz.com), and Autonomous System Networks (ASNs), IP source netblocks, and domains were registered in Iran. The netblocks utilized had strong associations to state-owned oil and gas companies that employ individuals with expert knowledge of ICS systems. Further, tools in the malware warn the attackers if their outward facing IP address traces back to Iran. The infrastructure utilized by the group is too robust and too centralized to have belonged to an individual or small "grass-roots" hacktivist group. This leads leading security firms, such as Cylance, to believe that Tarh Andishan is either state sponsored or a well-funded mercenary hacker group.

In Farsi, “Tarh Andishan” translates as “Thinkers”, “Innovators”, or “Inventors”. Tarh Andishan consists of at least 20 dedicated hackers and developers, believed to be located in Tehran, Iran. Additional, members or hired associates operate out of the Netherlands, Canada, and the United Kingdom. Persian names (Salman Ghazikhani, Bahman Mohebbi, etc.) were used as hacker monikers. Most targets of Tarh Andishan speak English as a primary language and it appears that members of the group are proficient in reading and writing in English. Different members of the group specialize in different malware, different malware development tools, different programming languages and different adversary techniques.

Tarh Andishan targets government entities and critical infrastructure facilities in Canada, China, England, France, Germany, India, Israel, Kuwait, Mexico, Pakistan, Qatar, Saudi Arabia, South Korea, Turkey, United Arab Emirates, and the United States. Specifically, Tarh Andishan has been known to target: military installations, oil and gas facilities, energy facilities, utility facilities, transportation facilities, airlines, airports, hospitals, telecommunication companies, technology firms, institutions of education and research, aerospace and defense facilities, chemical companies, and governments. The expansive range of targets across the globe indicates that the Tarh Andishan campaign is likely a mechanism for gaining geopolitical leverage and establishing Iran as a cyber-power. Iran may be demonstrating that it can retaliate against any country that compromises its cyber-security.

Academic institution networks are often targeted by malware because universities, especially those that work with their government in some capacity, sponsor valuable research. Universities often store sensitive PII documents and research on local servers. Yet, university networks are de-centralized and often poorly secured because different schools on campuses host different networks that are supported by different IT teams and each network needs to be accessible to thousands of users with varying needs. While the origins of Stuxnet have never been definitively confirmed, it is believed to have originated out of a university research program. Tarh Andishan targets university networks for research, but according to Operation Cleaver, it also attempts to steal student PII, student photos for identification cards, and passport information from universities in the United States, India, Israel, and South Korea. Student PII and photos could be used for identity theft, but it could also be used for intelligence purposes because the next generation of government recruits and security researchers are currently students.

Tarh Andishan targeted airlines, airports, and transportation networks in South Korea, Saudi Arabia, and Pakistan by compromising Windows Active Directory and physical internal infrastructure such as Cisco edge switches, and routers. From there, the attackers stole VPN credentials so that they could establish a persistent presence and so that they could remotely access the entire infrastructure and supply chain. Tarh Andishan used the

compromised credentials and VPN access to compromise airport gates, access security control systems, make fraudulent payments with Paypal and Go Daddy, and to infect other internal infrastructure. Overall, Operation Cleaver saw Tarh Andishan dangerously compromise airline networks without encountering major resistance. Information exfiltrated by Tarh Andishan could put airline passengers at risk if Tarh Andishan used its access to compromise airline ICS, SCADA systems, or other critical infrastructure. Further, Windows Active Directory, Cisco edge switches, and routers are components of networks in almost every organization in almost every sector. Given its success, Tarh Andishan may easily adapt this technique to attack networks in other sectors of its attack profile, if it has not done so already.

According to Cylance, Tarh Andishan's "Initial compromise techniques include SQL injection, web attacks, and creative deception based attacks – all of which have been implemented in the past by Chinese and Russian hacking teams." Tarh Andishan did not appear to utilize zero-day exploits. The SQL injection attacks were made possible by attacking vulnerable applications that failed to sanitize input prior to passing it to a database in an SQL query. Later, Tarh Andishan began spear phishing attacks, which involved sending victims an email with a malicious link. One such attack told targets that they had been selected to apply for a new position at an industrial conglomerate and the link directed them to a copy of a legitimate resume creation website. The resume tool was combined with a binder tool that loaded malware onto created documents. The malware runs in the background of the victim's system and logs keystrokes and the information entered into forms. After the malware infected a host, the attackers would leverage existing, publically available, exploits (such as MS08-067) to escalate their privileges on Windows systems. The malware then propagated through the network like a worm, to compromise other systems on the network. Tarh Andishan compromises Microsoft Windows web servers that run Internet Information Services (IIS) and Coldfusion, Apache servers with PHP, Microsoft Windows desktops, and Linux servers. The group also targets popular network infrastructure such as Cisco VPNs, Cisco switches, and routers.

Tarh Andishan's most utilized malware, TinyZBot, gathers information from infected systems and it establishes backdoors for persistent access. TinyZBot uses the SOAP sub-protocol of HTTP to communicate with the C&C infrastructure and it abuses SMTP to exfiltrate data to the C&C servers. Among other capabilities, TinyZBot can also take screenshots of the system, download and execute arbitrary code, detect security software, disable some anti-virus, and modify PE resources. Once the malware has infected the system, Tarh Andishan can use customized tools to poison ARP caches, encrypt data, steal credentials, create backdoors, create ASP.Net shells, enumerate processes, record HTTP and SMB communications, detail the network environment, query Windows Management

Instrumentation (WMI), log keystrokes, and more. Effectively, Tarh Andishan can customize their tools to suit any target. The Net Crawler tool, which combines popular attacker tools Windows Credential Editor, Mimikatz, and PsExec, was used to gather the cached credentials from every accessible computer on the infected network. Shell Creator 2 was used to generate an ASPX web shell to protect the attacker from revealing internal information such as location by human error. The Nbrute utility uses NMap to map the network and then it attempts to determine network credentials via brute force. The attackers can also use tools such as the PVZ bot tool to log keystrokes on specific botnet systems and save information on infected systems to specific locations.

Ajax/ FLYING KITTEN/ Saffron Rose

The Ajax group began in 2010 with website defacement attacks, but their activity escalated to cyber-espionage by 2013. The group's C&C infrastructure was set to Iran Standard Time and used the Persian language. The Ajax team consists of 5-10 members and it is unclear if the group is part of a larger movement such as the Iranian Cyber Army. The group may have been founded by members using the monikers "HUrr!c4nE!" and "Cair3x." The group uses custom malware, but they do not leverage software exploits. The lack of exploits indicates that the group is more likely a patriotic hacktivist group than a state sponsored threat.

Ajax primarily targets United States defense contractors, firms that developed technologies that bypassed the Iranian censorship policies, and Iranian dissidents. The group has also participated in attacks against Israel with the Anonymous group.

The group tries to lure victims into revealing login credentials or self-installing malware through basic social engineering instead of leveraging software exploits. These social engineering attacks proceed through email, instant messages, private messages on social media, fake login pages, and anti-censorship technology that has been pre-loaded with malware. Past messages have directed targets to a fake login or conference page. The page spoofs a legitimate organization or application and it collects user login credentials. After the user logs in, they are directed to a different page that tells users that their browser is missing a plugin or that they need to install proxy software, which is actually the malware. In some cases, the messages just send the user to the latter page. Iranian Internet Service Providers (ISPs) block "unacceptable content" such as pornography or sources of political dissidence. Ajax team has been infecting anti-censorship software, such as Psiphon and Ultrasurf, with malware and redistributing it.

Ajax relies on the Stealer malware which consists of a backdoor and tools. Using one tool, the attackers can create new backdoors and bind them to legitimate applications. Stealer collects system data, logs keystrokes, grabs screenshots, collects credentials, cookies, plugin information, and bookmarks from major browsers, and collects email and instant messenger information along with any saved conversations. Stealer also has components that acquire Remote Desktop Protocol (RDP) accounts from Windows vault and collects user browsing history. Data is encrypted using symmetric encryption (AES-256) using a hardcoded encryption key. The information is then exfiltrated using FTP with a built in client (AppTransferWiz.dll).

A new version of the Stealer malware, dubbed Sayad, surfaced in July 2014. The variant includes a dropper called Binder and new communication modules that allow it to exfiltrate data using HTTP POST requests. Binder checks the .NET runtime version of the target machine and drops the relevant version of the malware. The malware is now more modular and contains development files suggesting the future capability to exfiltrate files from the target system.

South Korea

Dark Hotel/ Tapaoux/ Nemim/ Pioneer/ Karba

According to Kaspersky Lab, the Dark Hotel group may have been stealing confidential documents out of the secured computers of travelling executives since 2007. Researchers believe that the group is Korean in origin (in part) because variants of the malware were designed to shut down and remove itself from the host system if the infected system code page was set to Korean. Further, the kernel mode key logger used in Dark Hotel attacks has Korean characters in its code and may be tied to a South Korean programmer. Since the group still targets North Koreans, one could suppose the Dark Hotel group originates in South Korea. The Dark Hotel attack campaigns use a sophisticated keylogger and extensive infrastructure to steal confidential information directly relevant to South Korea from employees of other nations. Consequently, there is a strong likelihood that Dark Hotel is a partially or fully state sponsored threat actor.

The Dark Hotel group targets high-profile executives, sales and marketing employees, R&D staff, and government employees from North Korea, Japan, India, and the United States. Notably, targets tend to be from the Asian nations with nuclear capabilities and the United States. Dark Hotel often targets guests staying at luxury hotels in Asia. A smaller number of hotels in the United States have also been infected. Overall, fewer than two dozen hotel

network compromises have been discovered, but it is possible that many breaches remain undiscovered or unreported. Hotels appear to be targeted based on the expectation that specific individuals will be staying there in the near future. Evidence suggests that the adversary possesses knowledge of the personal information of targets, at which hotel individual targets will stay, and the duration of their stay. The attacks may target specific individuals or all individuals who try to connect within a specific period. It is possible that the hotel attacks target those unlikely to fall for a spear phishing campaign. Specific targets may be located based on their Wi-Fi connection in the network, which is often secured with a password created from their surname and room number. Either the actor targets the hotel network directly or on occasion, it compromises the third party that manages the Wi-Fi for multiple hotels. The malware is distributed across the network either before the staff arrives to work or after they leave. When the target concludes their stay, the adversary removes all or most traces of the attack from the hotel network. Neither backdoors nor tools are left behind.

Upon connection to the hotel Wi-Fi, target users encounter a malicious iframe that redirects their browsers to fake update installers. Victims see a pop-up for a software update (Adobe Flash, Google Toolbar, Windows Messenger, etc.) that is actually a malicious executable piggybacking off a legitimate update installer. The installer delivers one of the group's backdoors to the victim system. Supposedly, the malicious download proceeds even if the user becomes suspicious and attempts to terminate their Wi-Fi connection. In 2015, the group may have begun to infect mobile devices through the same process. The malware remains dormant for an estimated six months before data collection and exfiltration begins. This precaution evades corporate IT efforts to scan a travelling executive's computer upon their return to the home network.

In addition to the hotel attacks, the group infects victims through spear-phishing attacks and P2P networks. The spear phishing attacks are used to target a specific victim at a specific hotel while the P2P campaigns infect as many hosts as possible with botnet malware. The spear phishing campaigns typically target the defense sector, NGOs, and government entities. The lure emails are titled with topics related to nuclear energy or weapon capabilities. If the target ignores the spear phishing email, the group waits (up to a month) and then tries to spear phish the target again. The emails contain links that redirect the victims to landing pages that deliver zero-day exploits. Sometimes an attachment containing an Adobe zero-day exploit is included instead. Recently, some of the emails have also relied on (former) zero-day exploits that were revealed in the Hacking Team breach or have delivered malicious code disguised as .hta files.

In the P2P campaigns, the adversary compromises a swath of users through infected torrented material. One example of how the adversary deploys malware along the P2P attack vector was caught by Kaspersky Lab in 2013 - 2014. In this case, the malicious actor seeded

Japanese explicit comic book sites with the Karba trojan so that the malware would be widely and wildly distributed when torrent users downloaded the pornographic material as RAR archives on torrent clients. The archive in question was downloaded over 30,000 times over a six-month period. Even if the attack was only marginally successful and the malware only installed on a fraction of victim systems, the attacker still gained a sizable botnet. Considering that the adversary could run numerous similar attacks simultaneously, it is safe to speculate that the actor can leverage an enormous botnet in attacks. If the infected host system contains interesting information, then the actor uses the botnet to install a backdoor and more sophisticated tool kit on the system so that they can exfiltrate documents and data.

The malware appears as legitimate software verified by legitimate certificates. The adversary did not steal certificates. Instead, the actors generated 22 certificates by exploiting a certificate authority, DigiCert Sdn. Bhd., that belonged to the Malaysian government and Deutsche Telekom, which was using weak 512-bit signing keys. To generate legitimate certificates, the actor just factored weak 512-bit RSA digital signature keys. Some recent malware and backdoors attributed to the group have featured SHA1 and RSA 2048-bit certificates, which may have been stolen or generated from a different source.

The group's toolkit predominately relies on a sophisticated 300 kb kernel mode keystroke logger, which operates at the system core instead of at the application layer. As a result, it bypasses most security and detection systems. The driver of the keylogger installs as the system kernel driver "Ndiskpro" service, a self-described Microcode Update Device. The keylogger retrieves data directly from the motherboard controller at port 0x60. A moniker in the source code of the keylogger attributes it to "Chpie," a South Korean coder. The data is transferred to the user mode component, where it is encrypted (similar to RC4) and written to a randomly named temporary (.tmp) file that is located in the same directory as the initial dropper, which maintains persistence across reboots by amending the HKCU run key. The toolkit also contains an information stealer, the Karba trojan, research environment detection mechanisms, and selective infectors, droppers, and self-injectors. The information stealer collects passwords and user credentials stored in browsers for email clients and social media accounts. The Karba trojan collects system data and information about installed anti-virus software. The primary dropper (recognized as Virus.Win32.Pioneer.dx) drops the selective infector (igfxext.ece) to disk and runs it. The selective infector, true to its nature as a virus, infiltrates and infects other computers through the network or through shared USB connections. It also collects information and sends it back to the C2 infrastructure. At least nine different backdoors have been used in conjunction with the toolkit.

No server level back doors were discovered on the hotel networks. Server logs show that the attacker compromised the servers (through a currently unidentified attack vector),

infected the target hosts, deleted traces of their presence, and then abandoned the system without leaving a backdoor or other malicious code behind. Since some attacks occurred over years, it is likely that either the attacker deleted their backdoor when they abandoned the server or that they had an insider in the target company.

Prior to discovery in October 2014, the C2C infrastructure consisted of over 200 servers containing malware, botnet logs, and stolen data. After the campaigns were revealed to the public in late 2014, much of the C2C infrastructure was shut down; however, the group remains active as of 2016 on new infrastructure.

North Korea

Bureau 121/ Guardians of Peace/ Dark Seoul

According to defectors, Bureau 121 is one of six divisions of North Korea's General Bureau of Reconnaissance that is charged with cyber-intelligence operations. The bureau was created in 1998 and it consists of ~1800 handpicked hackers who are allegedly the "most talented and rewarded personnel within the North Korean military" according to a Reuters interview with a defector known as Jang Se-yul. Students are recruited directly from the University of Automation and paid relatively significant sums.

North Korea uses cyber-warfare as a cost effective intelligence branch of their military. Many in North Korea see cyber-warfare as their strongest weapon. Bureau 121 most frequently targets South Korea, Japan, and the United States. Bureau 121 targets financial institutions and media companies. In one March 2014 attack, 30,000 South Korean servers associated with banking and media broadcasting outlets were damaged. These systems were infected with DarkSeoul malware and they displayed messages claiming that they were hacked by the Whois Team. In November 2014, Sony Pictures' email server was hacked by a group claiming to be called the Guardians of Peace, in response to the upcoming release of the movie "The Interview" because it portrays a story and portrayal that is unflattering to Kim Jong-un. An estimated 100 terabytes of data were exfiltrated from Sony before the Wiper Trojan was used to delete the servers. The information contained emails, unreleased films, employees' personal information and financial information. Threats were also made against Sony that contained imagery reminiscent of the September 11, 2001 attacks.

The FBI, Obama Administration, and the NSA have attributed the Sony breach to North Korea. Members of the press and some security researchers doubt the evidence

attributing the Sony attack to North Korea. North Korea may not have been capable of exfiltrating hundreds of terabytes of data.

The Whois Team and the Guardians of Peace attacks are very similar. Both attacks were relatively unsophisticated and both attacks offered a moniker of a previously unheard of group. The procedure of each attack was to install malware through phishing campaigns, steal data, lock down the infected systems, display a banner message claiming responsibility, and then using malware to wipe the system.

Russia:

Energetic Bear/ Dragonfly/ Havex Crouching Yeti/ Koala Team

Since 2011, Energetic Bear, an Eastern European threat actor, has targeted the Defense Industry, Energy Industry, and ICS equipment manufacturers, with highly technical prolonged attacks that are suggestive of a state sponsor. Energetic Bear's exploit kit features specialized malware, likely developed or adapted by the attackers, that was compiled during business hours (Monday – Friday, 9am – 6pm) UTC+4, which corresponds to working hours in Russia or Eastern Europe. Most security firms conclude that Energetic Bear is a Russian state-sponsored group because the group targets nation states who are politically opposed to Russia. Further, the malware primarily compromises petroleum and energy systems that compete with Russia's energy complex in the economical arena.

Based on its choice of targets and the malware deployed, Energetic Bear seems primarily interested in gathering intelligence on its victims or their country of origin and establishing persistent access to compromised systems. The sophisticated exploit kit could easily be used to sabotage targets' operations to cause damage or disruption in critical infrastructure sectors that depend on ICS and SCADA systems. So far, while the malware has been positioned ideally to sabotage ICS and SCADA systems, investigations by Symantec and other leading firms witness more uses of the exploit kit for espionage purposes than the sabotage purposes. The threat actors may prefer not to utilize this capability or sabotage campaigns may occur, appearing as system failures that are not investigated as cyber-attacks. More likely, Energetic Bear may be pre-positioning its malware in compromised systems to grant the greatest utility while allowing for every attack vector. Given its selection of targets and its exploit kit, both of which are detailed below, Energetic Bear is uniquely positioned to assist in a combination of Digital and Physical warfare for military or political purposes. Notably, Russia conducted such a campaign in its 2008 conflict with Georgia.

When Energetic Bear was discovered in 2011, the group targeted aviation and defense companies in the United States and Canada; however, in 2013, energy firms in the United States and Europe became the primary targets of Energetic Bear. In particular, the exploit kit targets the systems of ICS equipment manufacturers and petroleum pipeline operators. Energy grid operators, electricity generation facilities, and industrial equipment providers are also susceptible to compromise. By ingeniously targeting the smaller, less protected ICS manufacturing companies and antiquated SCADA systems, Energetic Bear is able to circumnavigate the massive state-sponsored cyber-security systems that typically protect critical infrastructure systems.

The exploit kit mimics the Stuxnet worm (which monitored and sabotaged the Iranian Nuclear program in 2011) in potential impact. If the sabotage potential of the malware were realized, then Energetic Bear could disrupt and seriously damage energy supply and regulation systems in countries such as: the United States, Spain, France, Germany, Turkey, and Poland. Consider the tragedy that a malicious actor could wrought with the ability to remotely destroy oilrigs, energy generation facilities, or electrical grids. The smallest city-wide power outage has the potential to result in many deaths related to loss in electricity needed for in-home medical care, heating, and other technologies that assist in citizens' daily lives. Even if an attack is controlled well enough or mitigated soon enough to prevent serious physical damage to the facility, imagine the economic ramifications that the actor could inflict upon a nation state through repeated targeted attacks on its energy systems. The gas price hikes of the mid 2000's might seem a minor inconvenience in comparison to the damage caused by a persistent sabotage campaign.

From February to June 2013, Energetic Bear launched a spam campaign against the United States and European energy sectors. Executives and senior employees in seven organizations received emails, sent from a Gmail account, containing a malicious pdf. If the pdf was opened, then the malware spread to the network. The emails were made to look as if they came from a known source (such as the victims' boss) and organizations were targeted with anywhere between 1 and 84 emails. In a more ambitious spear phishing campaign, emails containing remote access Trojans (RATs) were sent to personnel in three ICS equipment manufacturers who dominated their markets. The malware injected malicious code into the ICS software update bundles that were later posted for download from the manufacturer's website. The targeted equipment which would receive the update are used in a number of sectors, including energy. The Trojan managed to compromise the bundles of two companies and infect the programmable logic controllers of devices produced by those manufacturers, before the infection was discovered.

Later, watering hole attacks were added to the campaign. In these attacks, websites often visited by personnel of the target organization were compromised (usually with an injected iframe) and set to redirect victims to a site that delivered an exploit kit that installed the malware on the victim's PC. The development of additional attack vector(s) and the resources to compromise third party sites as "stepping stones" to desired targets suggests that the group is state sponsored. In either attack, the malware was configured to search victims' systems for ICS software and updates and to trojanize the software so that the adversaries could compromise guarded ICS systems the next time the software was downloaded or they were updated by trusted personnel.

The group employs two exploit kits (LightOut and Hello) and two malware (Trojan.Karagany and Backdoor.Oldrea). The exploit kits are used to initially compromise the system and install the malware. The malware is used for espionage, persistent access, or sabotage. LightsOut exploits vulnerabilities in Java or in Microsoft Internet Explorer to deploy the Karagany or Oldrea malware onto a user's system. In September 2013, the Hello exploit kit replaced the LightsOut kit. The Hello kit is combined with watering hole attacks to redirect victims to a landing page, where a JavaScript fingerprints their system to determine details such as operating system, browser, and installed plugins. The victim is then redirected to the site that contains the exploit most likely to achieve the adversaries' goals.

Trojan.Karagany and Backdoor.Oldrea are remote access Trojans (RATs) that are used to install additional tools or malware, to search the system for valuable data, and to exfiltrate data from the system. In an attack, the group uses either Karagany or Oldrea, but never both, because the malware serve the same purpose. The Karagany malware is only used in 5% of attacks. Karagany is a widely available exploit for purchase or source code recompilation on the internet underground because its code was leaked in 2010. Karagany features tools for indexing documents, taking screenshots of the system, and collecting passwords. At the adversary's instruction, it can also download new tools or files, run plugins or executables, or exfiltrate data to a designated C&C server. Oldrea, also widely known as the Havex malware, appears to be used in most attacks and it appears to have been written by or written for the attackers. Once installed, Oldrea profiles the system by collecting system information, harvesting outlook address book information, noting VPN configuration files, and indexing files, programs, and the root of available drives. The data is compiled into a temporary file, encrypted, and sent to an adversary C&C server. Oldrea features a control panel that the adversaries can use to authenticate to a C&C server and download a compressed copy of each specific victim's data. The servers hijacked by Energetic Bear to serve as C&C servers may have been compromised using the same exploit of content management systems.

Uroburos / Epic Turla/ Snake / SnakeNet

In 2008, malicious code known as Agent.BTZ was placed on USB drives that were dropped in the parking lots of defense facilities, such as a United States Department of Defense in the Middle East, in what was considered the “worst breach of U.S. military computers in history” at the time. Agent.BTZ infected systems running Microsoft Windows and allowed attackers to log personal information, cached credentials, and user keystrokes. The infection propagated and lasted in United States government systems for over a year. The Agent.btz infection led to the creation of the United States Cyber Command. The Uroburos malware, which appeared in 2011 (or earlier) and was discovered in 2014, scans for the presence of Agent.BTZ on target systems and remains inactive if Agent.BTZ is installed. Comments and code itself indicate that the authors of both Agent.BTZ and Uroburos are proficient in Russian. Some file names, encryption keys, and other technical indicators are shared between the Agent.btz and Uroburos malwares. Although other possibilities exist, Agent.BTZ and Uroburos were likely developed by the same group or associated groups.

The Uroburos rootkit is a very advanced and very sophisticated modular malware designed to infect entire networks and exfiltrate confidential data. The sophistication and flexibility of the Uroburos malware suggests that a highly skilled team, who had access to considerable resources, developed it. The significant monetary investment necessary to develop the Uroburos platform suggests that it was developed to target businesses, nation states, and intelligence agencies, rather than average citizens. Based on the exploit kit, the Uroburos group likely has a political or espionage agenda. The Uroburos malware typically infects 32-bit and 64-bit Microsoft Windows systems that belong to governments, embassies, defense industries, pharmaceutical companies, research and education facilities, and other large companies.

The Uroburos group uses spear phishing campaigns, drive-by-infections, watering hole attacks, and social engineering to push their malware onto target networks. In spear phishing campaigns, the target receives a tailored email containing an executable RAR self-extracting archive (SFX). If opened, then the malware unpacks and installs itself (a .SCR executable) on the user system. When the Uroburos rootkit infects a machine, it can: execute arbitrary code, hide its activity on a system, identify and exfiltrate information such as files, capture network traffic, and infect other systems on the network. Uroburos consists of a driver (.sys file) and an encrypted virtual file system (.dat file). The complex driver seems to be specifically designed to be discrete and difficult to identify.

Remote attackers use Uroburos to infect other machines on the network and to communicate between infected hosts using a peer-to-peer architecture. Uroburos

opportunistically propagates through the network. If Uroburos infects at least one system on a network that has an active internet connection and that host is connected to other systems within the network, then the attacker can infect as many systems as their resources allow. The malware spies on each system for useful information and uses the P2P architecture to relay information to the attackers. As such, information can be retrieved from air-gapped systems, transferred from infected host to infected host until it reaches a host with an active internet connection, and then exfiltrated to the adversary. This methodology allows the malware to bypass many security controls.

The Uroburos rootkit aspires to hide its elements and remain undetected and persistent on the compromised system. Upon installation, the malware establishes a service (usually Ultra3.sys) that automatically executes during the startup of the system. This driver is necessary to decrypt the malware's virtual file systems, create additional hooks, inject code into user libraries and applications, and manage communication between the adversary and the malware. The driver hooks the malware into the system by injecting code into a running process and then redirecting the rest of the running code to execute at the end of the malicious code. As non-technical simplification, this process, known as inline patching, can be visualized as inserting an extension cord (the malicious code) between another cord and a wall socket. By doing this, the malware can better remain undiscovered because malicious activity is attached to legitimate processes.

The rootkit consists of two virtual file systems (a NTFS file system and a FAT file system) that are encrypted with CAST-128 and stored locally on the user system. The encryption key is hardcoded in the driver file. The virtual file system (a .dat file) has a random name and it is stored with the driver file. The encrypted file systems function as a work environment for the attackers. Third party tools, post-exploitation tools, temporary files, and binary output are stored in the file systems. The NTFS file contains bat scripts which enable the attacker to map remote servers, execute netstat commands, gather system information, log output of tools, tools to steal documents, encrypt stolen documents, and RAR tools to compress and archive stolen documents for exfiltration. A queue and library injection tool, which acts as a buffer between the queue and the user system, can pcap or snapshot network traffic.

The virtual file system contains protocol information to exfiltrate information through HTTP (external website with GET and POST requests), through ICMP (ping), through SMTP (email), and through named pipe to another infected system. New libraries and tools can be added by adjusting the built in queue, without reinstalling the malware. Airgapped systems can be infected through named pipe connections or through USB devices. In the former case, an infected system serves as a proxy node and it appears passive as it spreads the infection to other systems on the network. Any infected system can serve as a proxy node, so even if one

point of infection is discovered, a tangential system can continue to infect the network as the new proxy node. The peer-to-peer modular design is resilient to removal, scalable on any network, and reliable. Further, the framework can be extended to include new features and perform further attacks against the infected host or networks associated with the infected network. The design of the malware as a driver and a multi-file virtual file system that can only work in combination is an elegant, but sophisticated design that complicates analysis efforts. Without the driver, the other two files cannot be decrypted. Without the files systems, the driver is innocuous. The design is too sophisticated and too expensive to develop to be common spyware.

APT 28/ Sofacy Group/ Sednit Group/ Tsar Team/ Fancy Bear/ Operation Pawnstorm

APT 28 is believed to be a state sponsored group that has been active since 2007. The majority of the APT 28 malware was compiled between Monday – Friday from 8 a.m. – 6 p.m. in UTC+4. This parallels working hours in Eastern Europe, Moscow, and Saint Petersburg. Over half the malware contained portable executable information that indicated that it was programmed with Russian keyboard settings, while the remaining samples were coded using English or Neutral keyboard settings.

Unlike Russian cyber-criminal groups, APT 28 does not exfiltrate financial information from targets and it does not sell the information that it gathers for profit. Instead, APT 28 gathers geopolitical information that would be specifically relevant to Russia and it uses the information to leverage future attacks. APT 28 uses spear phishing campaigns, sophisticated malware, and zero-day exploits to infiltrate systems belonging to European governments, NATO affiliates, militaries, security organizations, and media organizations with the intent of exfiltrating state information that could be used to influence policy decisions, public opinion, or geopolitical issues. Most of the activity has centered on targets “of specific interest to a European government,” focusing on the Caucasus region and countries along the eastern European border.

APT28 relies upon spear phishing emails or zero-day vulnerabilities to initially compromise victim systems. APT28 spear phishing emails often originate from a typo-squatted mail server and they typically contain either a decoy document relevant to the target or the link to a typo-squatted malicious domain. The least sophisticated aspect of APT28’s more popular attack vectors is its reliance on user error to deploy its malware. Unsuspecting users must be tricked into opening the attachment or following the malicious link. Decoy documents are tailored to the target and they often contain a user specific title, to entice the user to open the attachment, or confidential information, likely obtained through previous breaches, to lend credibility to the document. In fact, the titles of the decoy

documents submitted or found online are so specific that the targets can often be retroactively guessed by security firms, such as Trend Micro, using only contextual information. Variations in the distributed decoy documents suggest that the actors are fluent in multiple languages (at least Russian and English) however; grammatical mistakes indicate that English is not their native language. While all signs in the malware indicate that Russian is the actors' native language some Russian researchers at the 2013 PHDays conference in Moscow argued that the dialect is not native Russian. APT28 uses specialized information about its targets to focus its attacks and limit detection. Only a limited number of personnel of the target organization receive the decoy documents. In one notable case, spear phishing emails were sent to only three employees of a billion-dollar multinational firm, whose email addresses were not publicly available or advertised online.

The Sednit platform consists of the SOURFACE/ CORESHELL downloader, the EVILTOSS backdoor, and the CHOPSTICK modular implant. SOURFACE (also known as Sofacy) or CORESHELL performs runtime checks and reverse engineering counter operations before verifying that the infected machine matches the system profile of the target. If the target is verified, then the SOURFACE/CORESHELL dropper obtains a second stage backdoor from the C2 server and installs it on the victim's system. The backdoor, EVILTOSS, is used to steal credentials and execute shellcode. EVILTOSS uploads an RSA public key and encrypts the stolen data. Then the data is sent via email as an attachment. EVILTOSS then delivers CHOPSTICK to the victim's system and installs it. CHOPSTICK is comprised of custom implants and tools that are tailored to the target system. CHOPSTICK actively monitors the victim's system by logging keystrokes, taking screenshots, and monitoring network traffic.

PinchDuke

The PinchDuke campaign, which operated from November 2008 until summer 2010, is believed to be the first campaign of the Duke malware family. PinchDuke targeted political organizations in Georgia, Turkey, Uganda, and the United States. The PinchDuke campaigns began 11 days after President Obama's April 5, 2008 speech concerning the deployment of missile defenses in Poland. In 2009 the campaign targeted the Ministry of Defense in Georgia, the ministries of foreign affairs in Turkey and Uganda, a United States foreign policy think tank, organizations associated with NATO exercises in Europe, and the Georgian Information Centre on NATO. In 2010, the group also targeted Kazakhstan, Kyrgyzstan, Azerbaijan, and Uzbekistan. The political nature of the targets suggests that the campaigns may have been state sponsored. The selection of targets closely mirrors those of the later APT 28/ Sofacy campaigns, which is widely believed a Russian state sponsored threat actor.

Like the rest of the Duke family of malware, the threat actor is attributed to Russia because error messages in the malware are written in Russian. Though many regions in Eastern Europe use Russian as their primary language, time stamps in the code suggest that the malware was developed in the same time zone as Moscow. The PinchDuke Trojan samples contain a text string that may serve as a campaign identifier to help the attackers differentiate between associated Duke malware campaigns that were run in parallel using similar exploitation kits.

The malware was delivered via phishing emails containing spoofed news articles from the BBC website or articles concerning NATO. The malware consists of multiple loaders and an information stealer trojan. The trojan is based around the source code of the information stealing malware, LdPinch, which has been available on underground forums since the early 2000s. PinchDuke's information stealer targets system configuration files, user credentials, and user files that were created within a predefined timeframe or whose file extension corresponds to a predefined list. PinchDuke communicated with its C&C servers through HTTP(s). In early 2010, PinchDuke campaigns decreased as other Duke campaigns began. Afterwards, PinchDuke or its components were absorbed into other campaigns. Notably, its loaders were later associated with CosmicDuke and occasionally the newer malware would install PinchDuke in its entirety on a victim system as a redundancy infection.

GeminiDuke

GeminiDuke was developed and deployed around the same time as PinchDuke and CosmicDuke. Unlike its sister campaigns, the January 2009 – December 2012 GeminiDuke campaign focused on collecting system configuration information from infected hosts. Samples of the GeminiDuke malware were compiled in UTC+3 and UTC+4 (depending upon the season), which corresponds to Moscow Standard Time during Daylight Savings Time.

Like PinchDuke and CosmicDuke, GeminiDuke was designed around a core information stealer component. The malware consisted of a loader, an information stealer, and numerous persistence components. The information stealer used a mutex based around a timestamp to ensure that only one instance of the malware was running at a time. The information stealer enumerates: local user accounts, network settings, internet proxy settings, installed drivers, running processes, values of environment variables, programs that run at startup, programs previously executed by the users, programs installed in the Programs Files folder, the files and folders in the users' home folder, the files and folders in the users' My Documents folder, and recently accessed files, folders, and programs. The malware employs multiple persistence

components similar to those included in CosmicDuke. MiniDuke's backdoor component resembles the source code behind one of GeminiDuke's persistence modules.

CosmicDuke/ Tinybaron/ BotgenStudios/ NemesisGemina

CosmicDuke is believed to have been developed and deployed by the same team as PinchDuke. CosmicDuke was compiled on January 16, 2010 and was still active as of June 2015. CosmicDuke superseded the PinchDuke campaign and its toolkit surpasses the functionality of the PinchDuke exploit kit. Unlike PinchDuke, CosmicDuke appears to be entirely custom written to the adversary's specifications. The techniques that CosmicDuke uses to extract user credentials and detect analysis tools may be based on PinchDuke. At a high-level, CosmicDuke's persistence techniques resemble those of GeminiDuke. Despite the similarities to the other Duke malware, CosmicDuke does not share any code with its sibling campaigns. CosmicDuke was most famously deployed against individuals believed to be trafficking illicit substances in Russia. It is possible that Russia's law enforcement agencies used the malware as spyware in their war against drugs.

CosmicDuke deploys from a series of loaders and the malware is built around an information stealer that is augmented by persistence components and a privilege escalation tool. Early variants of the privilege escalation module attempted to exploit CVE-2010-0232 or CVE-2010-4398. The malware authors likely chose which persistence and escalation tools to include in each variant of the malware in order to exploit known vulnerabilities in the target environment. For instance, in 2014, after the exposure of MiniDuke, Kaspersky noted the appearance of a CosmicDuke variant that featured a backdoor and the ability to start via Windows Task Scheduler.

The information stealer contains components to log keystrokes, capture screenshots, copy the contents of the clipboard, copy cached user credentials from web browsers and chat clients, export cryptographic certificates and private keys, and exfiltrate user files whose file extension corresponded to a predefined list. Additionally, CosmicDuke occasionally infected hosts with PinchDuke, GeminiDuke, or MiniDuke; though, CosmicDuke code never interoperated with the redundant malware code. After execution, the two malware ran concurrently and independent of one another. Typically, the malware even utilized different C&C infrastructure. F-Secure postulates that CosmicDuke may have deployed the other malware to allow the adversary to field test CosmicDuke while relying on the redundant malware to capture mission critical data should CosmicDuke not function correctly on the infected machine. CosmicDuke can exfiltrate the stolen data to hardcoded C&C servers via HTTP(s), FTP, or WebDav.

MiniDuke

MiniDuke is a highly customizable malware platform that was uncovered by Kaspersky Lab in February 2013. The malware may have been developed as early as 2010. According to Eugene Kaspersky, MiniDuke is unique in that it resembles more complex old school malware; in fact, many of its components are written in Assembly, a complex low-level programming language. This could indicate that the Russian authors behind MiniDuke have significant experience in the field.

The initial MiniDuke campaign compromised government institutions in Ukraine, Belgium, Portugal, Romania, the Czech Republic, and Ireland. Additionally, a research institute, two think tanks, and a healthcare provider were compromised in the United States, as well as a research foundation in Hungary. Victims were targeted with spear phishing emails containing malicious PDF files. If opened, the malicious attachments exploited a zero-day vulnerability and dropped a small downloader (20kb) onto the victim system.

The malware drops in 3 stages that are designed to evade sandbox, virtual, and analysis environments. Checks are processed at each stage before the malware decrypted more of itself. The downloader appears to be unique to the victim system and contains a customized backdoor. The downloader determines the system fingerprint and it later uses the information to encrypt its communication with the C&C server. If the target system meets pre-defined requirements and if the malware successfully installs, then the malware will access Twitter as a background process and search for specific tweets from pre-made accounts. Similar C&C infrastructure via Twitter can be found in variants of OnionDuke, CozyDuke, and HammerDuke. The tweets, authored by the malware operators, contain tags that correspond to the encrypted URLs where the backdoors are stored. The URLs lead to the C&C servers that contain commands and backdoors as .GIF files. In the event that Twitter is inaccessible, then the malware will run Google search in the background to find the encrypted strings that lead to the next C&C server.

OnionDuke

In October 2014, Leviathan Security Group disclosed that a Russia based Tor exit node was attaching malware onto the files that passed through it by wrapping legitimate executables with the malware executable. The technique increased the attacker's chance of bypassing integrity check mechanisms. The malware campaign is believed to have been active from at least February 2013 through spring 2015. OnionDuke does not operate like the other Duke campaigns; however, it does share some C&C infrastructure with the MiniDuke attacks.

Moreover, unshared domains in both campaigns were registered using the same alias, John Kasai. As such, it stands to reason that OnionDuke is another Russian state sponsored APT group.

OnionDuke attacks target government agencies in Central Europe. However, because it is unlikely that European government agencies are accessing Tor from their high value systems, the secondary distribution vector of the malware remains unclear. The malware has also been found targeting pirated software. It is possible that the campaign distributes the malware through scattershot attacks via the Tor network and torrent sites and through another yet unobserved vector, such as phishing or wateringhole attacks. The infection of Tor files appears to fail if the victim uses a VPN channel that encrypts traffic. Systems infected with CozyDuke may be infected with OnionDuke if the former malware is used to deliver and execute the latter malware's dropper. It is possible that the OnionDuke attacks were conducted to infect a broad range of target to gather information for the other Duke campaigns and to build a botnet for the adversary.

When traffic passes through the infected node, the dropper, Trojan-Dropper:W32/OnionDuke.A, is appended onto the legitimate files. The dropper contains a PE resource which appears as an embedded .GIF image file. In actuality, the resource is a .DLL file, Backdoor:W32/OnionDuke.B, which is then decrypted, written to disk, and executed. Next, the DLL decrypts an embedded configuration file, which attempts to contact a hardcoded C&C domain through HTTP(s) or through Twitter (if HTTP(s) fails). The domains appear to be legitimate websites that were compromised to deliver instructions and additional components to the malware. OnionDuke, like CozyDuke, is built upon a modular platform that was designed for versatility. The toolset delivered from the C&C server contains the information stealer, a DDOS module, a password stealing module, an information gathering module, and a social network (VKontakte) spamming component.

APT29/ Hammertoss / HammerDuke

APT29 is a new threat actor that operates during UTC+3 work hours. APT29 targets government organizations in an attempt to collect geopolitical data that could be of interest to Russia. APT29 might be a state sponsored threat group; however, the group is too new to exhibit definitive signs of state sponsorship.

APT29 employs anti-forensic techniques, they monitor analysis and remediation efforts, and they rely upon compromised C2C infrastructure. Apt29 embeds the Hammertoss commands into images using steganography. APT29 programs Hammertoss to operate to blend into normal target network traffic and normal target network traffic patterns. The

group preconfigures Hammertoss to activate after a predetermined date and only communicates during specified hours.

There are two variants of Hammertoss, Uploader and tDiscoverer. Both variants receive their instructions from an embedded image. Uploader goes to a hard-coded C2C server address and downloads an image of a specific file size. tDiscoverer generates and visits a new Twitter handle every day from a preconfigured algorithm. It attempts to visit that page. If the actor has registered the handle, then it visits the page and looks for a tweet with a URL that indicates the location of its instructions and a hashtag that specifies the minimum size of the image file. After the number of bytes, the hashtag may also contain a string that the malware adds to its encryption key so that it can decrypt the data. If the actor has not registered the handle, then the malware waits until the next day and repeats the process with the next handle generated by the algorithm. The malware fetches the image from the URL. Uploader or tDiscoverer, decrypts the data hidden in the image, and processes the attackers' command. Commands include conducting reconnaissance on the victim system, executing commands via PowerShell, or uploading stolen data to a cloud storage service.

CozyDuke/ CozyCar/ CozyBear/ Office Monkeys/ Cozer/ EuroAPT

The CozyDuke group began attacking governments and associated organizations around 2011. CozyDuke is a very precise group and it has not been extensively profiled. It may have been developed or used by actors of the MiniDuke or OnionDuke APT groups. CozyDuke shares at least some infrastructure with these groups. Security firm F-Secure reports that CozyDuke, like the rest of the Duke family of malware, originates from a seven-year campaign that is affiliated with the Russian government. It is unclear whether the Duke family of malware is sponsored by the Russian government or developed and used by a mercenary criminal organization.

Cozy Duke attacks very specific governmental organizations and affiliated entities. CozyDuke mostly targets United States entities; however, government and commercial entities in Chechnya, Germany, South Korea, and Uzbekistan have also been targeted. CozyDuke is believed to be behind the late 2014 attacks on the United States Department of State and attacks against the White House.

Like most APT's, CozyDuke attacks typically begin with a spear phishing email. Sometimes the emails are loaded with malicious Adobe Flash video attachments. In the past, the videos have been funny animal videos. Other times, the emails contain malicious links that deliver the user to websites that the attackers created to look like real sites. Otherwise, the email contains a ZIP file containing a decoy PDF document and a self-extracting RAR file.

Once the user opens the attachment or visits the link, the initial dropper is installed on the system. The initial dropper checks the system for security products, and will not install further malware if a program on the system matches software on its list. The dropper also runs processes to check if it is being run in a virtual machine or sandbox environment. If either check indicates an analysis environment, then the dropper exits. Otherwise, the dropper delivers an encrypted configuration file, and installs the CozyDuke components. The CozyDuke malware is signed with fake Intel and AMD digital certificates so that it appears legitimate to some security solutions.

The CozyDuke malware is a modular platform that consists of a core component, the CozyDuke backdoor, and modules tailored to its target. The platform includes multiple malware droppers and additional custom and open-source spyware tools. The CozyDuke main component establishes a persistent beachhead on the victim system, gathers system information, communicates with the C2 infrastructure, and manages the accompanying modules and scripts. The main component adds a registry value that is executed at system startup. It also obfuscates itself as a Windows service or scheduled task. Variants of the main component may also hijack the registry entry of a COM object “SharedTaskScheduler” so that the malware loads with the COM object.

CozyDuke modules can execute arbitrary code, harvest victim credentials, gather system information, and take screenshots of the victim system. Some of the CozyDuke modules appear to have been developed in the same development environment as MiniDuke and OnionDuke. The platform also contains the CORESHELL and CHOPSTICK modules made popular in the Russian state-sponsored APT28 attacks. CORESHELL is a second stage backdoor that runs numerous anti-analysis procedures. CHOPSTICK is a modular implant that logs keystrokes, takes screenshots, and monitors network traffic.

Recent variants of CozyDuke deliver SeaDuke and HammerDuke. SeaDuke is a cross platform backdoor that is written in Python. This expands the attackers’ pool of victims to include Linux users. HammerDuke is a backdoor that connects to a Twitter account name and uses tweets from the account to locate C2 server addresses from which it receives commands or to which to delivers data.

SeaDuke/ SeaDaddy/ SeaDask

SeaDuke appeared in October 2014, after the disclosure of most of the Duke campaigns. Like the majority of the Duke family, SeaDuke exclusively targets government organizations. The main difference between Seaduke and its sister campaigns is that SeaDuke focuses on a small number of high-value targets. Additionally, of the Duke malware, SeaDuke alone is

programmed in python. This developers' choice could indicate that the group is expanding their victim pool to Linux systems as well as Windows hosts. The overall framework of the malware remains similar to CozyDuke.

SeaDuke is a highly configurable trojan and backdoor that is often installed onto victim systems through CozyDuke or via a compromised website. It has hundreds of possible configurations. According to Symantec, the threat actor behind CozyDuke may only deploy SeaDuke in systems belonging to "major government-level targets." SeaDuke primarily allows the attacker to upload, to download, and to delete files on the victim machine as well as to retrieve bot/ system information and to update the bot configuration. It is possible that the threat actor deploys the malware to remove the indicators of compromise from other campaigns after a successful breach. The trojan may also be used to conduct pass the ticket attacks on Kerberos systems, to steal emails from Microsoft Exchange servers using compromised credentials, to archive sensitive data, or to exfiltrate data through legitimate cloud services. The C&C infrastructure behind SeaDuke relies on over 200 compromised web servers and several layers of RC4 and AES encryption and Base 64 encoding techniques. These extra obfuscation measures may be an attempt to remain undiscovered and thereby remove the attention on the Duke campaigns. SeaDuke communicates with its C&C servers via HTTP(s).

CloudDuke/ MiniDionis/ CloudLook

Discovered in June 2015, CloudDuke is the most recent Duke campaign. The campaign may be a tactical shift in response to the widespread disclosure of the other Duke campaigns by security firms such as Kaspersky, Symantec, and F-Secure. CloudDuke relies on spear phishing emails that closely resemble those deployed in the CozyDuke campaign. The CloudDuke emails contain a self-extracting archive attachment that appears as an empty voicemail file (.wav) or a PDF file (often containing the word "terrorism"). If opened, then the second stage dropper executes. So far, the campaign has targeted European diplomatic organizations.

The CloudDuke malware is comprised of a downloader, a loader, and two backdoors, which download and execute from either web address or from a Microsoft OneDrive account. The malware maps a OneDrive cloud storage drive as a network drive using hardcoded credentials and then it downloads its backdoors to the local system. The downloader may also download and execute additional malware, likely another Duke malware, from a preconfigured location. CloudDuke's backdoor functionality resembles that of SeaDuke. One backdoor will contact a preconfigured C&C server while the other relies on a Microsoft OneDrive account. As per its name, CloudDuke uses cloud storage services for its command and control infrastructure as well as its data exfiltration method.

Sandworm/ Quedagh/ BlackEnergy

The Sandworm team is a Russian advanced persistent threat group that targets systems of political targets of interest to the Russian Federation. Sandworm is likely state-sponsored. The group's name originates from strings in their code and names of their C&C servers that reference the Dune fantasy book series.

Sandworm has targeted governments and political organizations since at least 2009; but the group also may have been behind the 2008 cyber-attacks against Georgia. The Ukrainian government, NATO, the European Union, the European Telecommunications sector, European Energy companies, and Poland are among the group's top targets. Attendees of the May 2014 Globesec conference were also targeted. Many of the decoy documents used to deploy the malware were spoofed news coverage of political or economic situations in Europe.

The new variant of the BlackEnergy malware, which is now capable of stealing documents from targets, has been used against government institutions in Ukraine and Eastern Europe. The initial appearance of the malware coincides with the conflict between Russia and Ukraine. Trend Micro discovered that the newest variant of the malware, customized by the group, can target ICS and SCADA systems. The group may have infected these systems to monitor or sabotage systems that compete with Russia's energy interests.

Sandworm delivers malware through spear phishing emails containing malicious documents, such as a Microsoft PowerPoint attachment. The attachments either deliver the initial dropper or exploit a zero-day vulnerability to install the malware. In some cases, legitimate applications were trojanized to perform the installation. Through zero-day exploits, the malware infects any system running a Windows Operating System ranging from Vista to Windows, including Windows server systems. The malware only infects the victim system if the current user is a member of the local administrator group. If the user is not an administrator, then the malware will attempt to re-launch itself as Administrator or exploit the Windows backward compatibility features to bypass UAC.

The BlackEnergy crimeware appeared for sale in underground Russian cyber-markets around 2007. The malware was designed to create botnets for Distributed Denial of Service attacks (DDoS), but it has since evolved to support other capabilities. BlackEnergy can create botnets to send spam emails for phishing campaigns and it has tools to harvest passwords and banking credentials from infected computers.

The BlackEnergy toolkit gained notoriety during the 2008 cyber-attacks on Georgia during the conflict between Russia and Georgia. The BlackEnergy malware is available for

purchase in cyber underground communities; however, the variant used in Sandworm attacks has been modified with custom code, incorporates a proxy server infrastructure, techniques to User Account Control and driver signing features in 64-bit Windows systems, and tools to collect documents. F-Secure notes BlackEnergy is used by a variety of criminal and cyber espionage groups; so, Sandworm's adoption of BlackEnergy, instead of writing custom malware, may have been an attempt to shirk attribution and blend into the crowd of nefarious actors to remain undiscovered.

The BlackEnergy toolkit features a builder application that generates the clients used to infect victim systems, it features server-side scripts to create C&C servers, and it includes an interface for the attacker to communicate with their botnet. F-Secure comments that the toolkit is simple enough and convenient enough that anyone can build a botnet without possessing extensive skills. The information stealing plugin of the toolkit gathers system information, session information, a list of installed applications, a list of registered mail, browser, and instant messaging clients, a list of network connections, and stored user credentials for online and offline accounts, and exfiltrates the information back to the C&C server via a HTTP POST request. New variants of the malware may also be able to capture screenshots and record audio. On December 23, 2015, a Sandworm campaign against the Prykarpattiaoblenergopower plant in Ukraine caused a severe outage. More significant than the immediate loss of power, the threat actor, who is likely backed by the Russian state, demonstrated that the malware, which has been regularly discovered on U.S. networks, can severely cripple a nation's critical infrastructure as part of a cyber-physical campaign.

Carbanak

The Carbanak group is a criminal advanced persistent threat group whose attacks against dozens (potentially hundreds) of global financial institutions resulted in an estimated \$1 billion in losses in the first half of 2014. Depending on the victim, the attacks are believed to have begun between December 2013 and June 2014. According to Kaspersky Labs, each victim bank lost \$2.5 million to \$10 million to the campaign. The victim financial institutions were located in Russia, the United States, Germany, China and Ukraine; additionally, the group may also have begun targeting organizations in Malaysia, Nepal, and Kuwait. The vast majority of victims (at least 52) are located in Russia. Overall, the group targeted at least 100 financial organizations at 300 IP addresses located in approximately 30 countries. Of the ~100 organizations targeted, Kaspersky believes that at least half suffered financial loss.

The Carbanak group is particularly significant because it demonstrates how the dangerous escalation of sophisticated cyber exploit kits, perpetuated by state sponsored groups and government agencies, has guided the development of complex and demonstratively effective criminal platforms that can financially harm private organizations and individuals alike. Consider that the Carbanak group stole an estimated \$1 billion in less than 6 months. The loss to the global financial institutions, though meager compared to the entire global economy, can still lead to cascading global economic impacts within and outside the victim organizations.

Like most APT groups, Carbanak attacks began with a spear phishing campaign. The malicious emails appeared as legitimate banking communique accompanied by attached Microsoft Word (97-2003) documents and Control Panel Applet (.CPL) files. The attachments infected victim systems with malware and with a backdoor based on the Carberp malware. It is also possible that some of the emails contained urls that redirected the victim to a landing page that delivered the malware in the background before forwarding the user to a familiar financial site. Analyzed malicious attachments reveal that the attackers exploited vulnerabilities in Microsoft Word 2003, 2007, and 2010 (CVE-2012-0158, and CVE-2014-1761). After successful exploitation of a vulnerability, the shellcode decrypts and the Carbanak backdoor is installed on the victim host. The Carbanak backdoor installs and then it re-installs itself into “%system32%\com” as a copy of “svhost.exe” with the system, hidden, and read-only attributes. The initial version (delivered by the exploit) is then deleted. After installation, the backdoor connects to its C2 server through HTTP (with RC2+Base64 encryption) and downloads a file (kldconfig.plugin) which details which process to monitor. The kit sets the Termservice service execution mode to auto to enable Remote Desktop Protocol (RDP). The backdoor provided access to the intranet of the victim organization. Next, the adversary probed the intranet for other vulnerable targets and specifically for critical financial systems. Typically, tens to hundreds of computers were infected before an admin system, with the necessary access, was compromised. If banking applications such as BLIZKO or IFOBS are discovered, then a special notification is sent to the C2 server to notify the adversary that financial systems were discovered.

Once the attackers discovered financial systems on the victim network, they deployed keyloggers, tools to hijack video capture, and screen capture tools to learn as much information as possible about the environment. The Carbanak tool kit typically logs keystrokes and takes screenshots every 20 seconds. The monitoring occurs by intercepting the ResumeThread call. The captured videos are recorded at low bandwidth and are used to help the attackers develop an operational picture of typical workflow, tool usage, and practices. In addition to training the adversary to transfer money, the monitoring also reduces the likelihood that the adversary will set off behavioral analytic systems. The remote administration tool, Ammyy Admin, might also

be installed on victim systems to ease remote access (the tool is whitelisted by legitimate system administrators in some corporate environments).

Attackers studied the financial tools and applications installed on the victim hosts in order to maximize the potential gain from the compromised system. Rather than searching for exploits and flaws in the security and financial applications, the adversary meticulously recorded the activity of administrators in order to learn the necessary information and procedures to transfer money. Files on captured C2 servers indicate that the adversary may also exfiltrate classified emails, manuals, cryptographic keys, and other information. When the adversary knew the necessary information and knew how to use the most powerful host applications, they would withdraw or transfer significant sums. The method of withdraw or transfer depended on the system, situation, and available resources (time, people, etc.). Observed methods of stealing cash include fraudulent online banking transfers, electronic cash transfers to banks in China and the United States, SWIFT transfers to compromised bank accounts, and remote commands to ATMs to spew cash onto the street at a specific date and time. In the instances where physical interaction with an ATM or bank personnel was necessary, the group would pay individuals to act as “mules” in the cash transfer.

The command and control infrastructure rotates every few weeks. It consists of Linux servers to issue commands, Windows servers used for remote connections, backup servers, and drop servers containing executables and additional components. Victim systems are catalogued in server logs according to the adversary’s categorization.

Syria:

The Syrian Electronic Army (SEA)

The Syrian Electronic Army is a public online political group that emerged in 2011 to support Syrian President Bashar al-Assad and his regime. The army arose days after Syria lifted its online ban of Facebook and YouTube. SEA was once managed by the Syrian Computer Society, which was headed by President al-Assad in the 1990s. The Syrian Computer Society, which regulates the internet within Syria, even registered the SEA website. The SEA may be partially or entirely supported by the Syrian government. At present, the SEA’s domain is no longer hosted by the Syrian Computer Society and it claims no ties to the government. Based on the aptitude at social media and the humor used on defaced sites, the army likely consists of young adult males. One “inside source” claimed that the group consisted of nine Syrian college students; however, no other sources have verified this claim.

By all appearances, the SEA conducts attacks to garner global attention rather than to steal data or financial information. The SEA primarily targets media outlets and journalists, political groups that oppose al-Assad's regime, human rights groups, and western organizations. Most SEA attacks target the websites and social media accounts of United States news organizations because it argues that the outlets spread anti-Syria propaganda. The SEA uses malware and phishing campaigns to actively monitor Syrian rebels and members of Human Rights groups.

SEA attacks begin with phishing through spam or spear phishing using detailed information obtained from previous campaigns. The SEA attempt to gain user credentials, which it then uses to seize control of the websites and social media accounts of prominent organizations. The army has attacked the websites and/or social media accounts of: "60 Minutes," Al-Jazeera, Associated Press, BBC News, CBC News, CNN, The Daily Telegraph, Financial Times, The Guardian, The Onion, National Public Radio, The New York Times, Reuters, Time, and The Washington Post. Once it has control, The SEA posts fake stories or news and collects any confidential information that could be useful in future attacks, such as contact names. When phishing attempts fail, SEA may resort to malware, website defacement through web exploits, or denial of service attacks leveraging botnets. If no attack vector succeeds, then the SEA resorts to bombarding the social media accounts of its target with pro-Syria messages.

Most attacks amount to a banner ad or redirection to a site that supports al-Assad; however, the attacks can have tangible impacts. When the SEA hacked the Associated Press Twitter account in 2013, they posted a message that the White House had been bombed and that President Obama was injured. The post resulted in a noticeable impact on the DOW Jones and the S&P 500 Index (~\$136.5 billion).

In their attack on the New York Times, the SEA demonstrated the ability to breach a major domain registrar, Melbourne IT, using stolen credentials and redirect internet traffic or seize ownership of domains, such as Twitter. The SEA has also compromised the GoDaddy domain registrar, social media management services, and third party applications that serve news articles. The attack on a registrar indicates that the SEA may begin to attack third party services and underlying infrastructure in order to compromise its target. Recently, the SEA has attacked larger targets such as Microsoft, Facebook, EBay, and PayPal through the underlying infrastructure.

Global

Anonymous

Anonymous is a collective of hacktivists and script kiddies which originated in 2003 on the website 4chan. In the traditional sense, Anonymous is more of a cyber-mob than an advanced persistent threat; however, the group's construction and global membership afford it significant influence and resilience to law enforcement efforts. Anonymous has established a brand name with the physical weight of a cohesive advanced persistent threat group. Anonymous has a decentralized command structure and it unites its members through anarchic ideology. Essentially, the loosely affiliated members or member groups work towards goals that they agree upon or remain inactive or split off, if they do not agree. Dissent is common within the group and one of the largest difficulties in profiling Anonymous is that the only absolutely unifying characteristic is membership in the group. Some members participate to deface websites and prank organizations while other members participate because Anonymous affords them a serious political activism platform. Most of the members support the foundational anti-censorship and anti-control platform and they target entities accused of censoring the people. Members, Anons, range from non-technical supporters to active blackhat hackers. Essentially, if an individual believes in the Anonymous cause or simply says that they are a member, then they are part of the collective. Anonymous members are told to neither reveal their identity or to discuss the group. The sense of membership and ease of access has allowed a few skilled hackers in Anonymous to hide amongst massive crowds of protesters.

Anonymous began by attacking the Church of Scientology, but its scope rapidly expanded. Since then, Anonymous has protested mass surveillance, anti-digital privacy efforts, governments, financial institutions, and individual users. More specifically, Anonymous has targeted the MPAA, the RIAA, Sony, the Church of Scientology, the Westburo Baptist Church, government entities in the United States, Canada, Israel, Tunisia, and Uganda, PayPal, MasterCard, Visa, and child smuggling and child prostitution rings. Anonymous supported the Occupy movement against large businesses, and it supported the Arab Springs movement against oppressive regimes in the Arab region. The media is the only sector that Anonymous members are prohibited from targeting.

Anonymous defaces websites and organizes distributed denial of service attacks (DDoS). Hacked websites may feature the pivotal picture of the Guy Fawkes mask, it may feature a manifesto claiming responsibility for the attack, or it may simply display an internet meme. DDoS attacks are conducted with Gigaloder, JMeter, or the Low Orbit Ion Cannon

(LOIC) applications. These tools flood a server with inbound TCP or UDP packets. Botnets belonging to members of the group are often added to DDoS campaigns. In some attacks, these botnets account for up to 90% of the malicious traffic.

America:

Butterfly Group/ Morpho

The Butterfly group performs corporate espionage campaigns against organizations containing proprietary intellectual property. Stolen information is likely sold for fiscal gain. The Butterfly group is organized and efficient. It is likely that the group consists of only a few individuals (~3-10 members). According to Symantec, “[t]here are some indications that this group may be made up of native English speakers, are familiar with Western culture, and may operate from an Eastern Standard Time (EST) time zone.” The emergence of the Butterfly group should remind organizations that corporate espionage groups and non-state sponsored APTs still exist. In fact, in certain aspects, they are more dangerous than state sponsored groups. Mercenary and espionage groups may possess specific knowledge of what information to steal or from what systems to steal data. This information may come from competitors or it may come from insider threats within the organization. APTs, like the Butterfly group, are more likely to profit from exfiltrated data and stolen intellectual property than an enemy nation state might. Auction of stolen information to a third party will likely occur immediately after a breach because the group maximizes their potential by realizing profit and redirecting their resources to the next target. Few concurrent campaigns were observed. Once information is sold to a third party, attribution of the attack becomes more difficult. The realized impact of lost financial data or stolen intellectual property could cripple the organization.

The Butterfly group has targeted pharmaceutical companies, technology firms, law practices, oil and precious metal mining organizations, Twitter, Facebook, Apple, and Microsoft. Since their creation in 2012, the group has compromised at least 49 organizations. There was only one government victim and they may have been collateral damage of a different campaign. Butterfly does not appear interested in nation state intelligence. After the attacks against Twitter, Facebook, Apple, and Microsoft in February 2013 drew the attention of security researchers, the group went dormant. They reemerged in August 2013 and have been gradually increasing their number of attacks per year. Of the 49 companies targeted, 17 are based in the United States, 12 are based in Europe, and 4 are based in Canada. The remaining 16 victims are located in Brazil, China, Hong Kong, India, Israel, Japan,

Kazakhstan, Malaysia, Morocco, Nigeria, Taiwan, Thailand, South Korea, and the United Arab Emirates.

In attacks against pharmaceutical companies, the attackers breached small regional offices and then slowly moved across the network to the main network. In late 2014, two natural resource organizations that specialize in gold and oil were compromised. In June 2015, a Central Asian global law firm was compromised and financial information and information about regional natural resources may have been targeted. This has led to speculation that the attackers may be focusing on information that is valuable in the commodities market. The behavior may also indicate direction from a third party client who is invested in the commodities market.

Attacks seem to be focused on specific systems that are of interest to the attackers, such as Microsoft Exchange or Lotus Domino email servers. The attackers may want to monitor emails or they may want to inject messages into the server. Content management servers, which index and store documents and digital assets, were also targeted. According to Symantec, these servers likely contained legal documents, internal policies, training documents, product descriptions, and financial records. The actor may gauge the value of a target based on training materials and presentations for related technologies under development at the organization. In at least one instance, the group hacked a Physical Security Information Management (PSIM) system which collects, processes, and stores data from physical security devices such as CCTV, magnetic card systems, HVAC, and building security systems. The actor could have been monitoring employees throughout their daily activities, or the system could have been compromised by mistake.

The Butterfly group exploits zero-day vulnerabilities from a water hole website. In February 2013 Twitter, Facebook, Apple, and Microsoft were attacked within a three-week period. The Butterfly group initiated their campaign with a Java zero-day exploit that was delivered from a popular iPhone mobile development website. For some of the attacks, F-Secure believes that the payload delivered after the breach may have been a Mac OS X backdoor, dubbed OSX Pintsized. Attacks against Windows systems likely featured the Jripbot backdoor. Symantec believes that the group may also exploit Internet Explorer 10 or an Internet Explorer plugin. At least one recent attack suggests that the group might also conduct SQL injection attacks.

After a network is compromised, the group carefully adapts to the environment and utilizes remote access tools and management systems to laterally move across the network. The adversaries have used native Citrix systems and the TeamViewer applications to move across some networks. The attackers are able to rapidly assess whether a system is valuable or whether they should move to a new system on the network. The Butterfly group uses a

unique set of tools, which seem to have been developed by or developed for the attackers. Symantec could not find any open source data on the tools. The tools all contain use documentation. One tool, bj.dat, (called "Banner Jack.") is used to locate vulnerable network servers, printers, routers, HTTP servers, or TCP servers. Banner Jack retrieves default messages from Telnet, HTTP, and TCP servers. Banner Jack accepts an input IP range and port and then it connects each IP address to a port. Then it retrieves and logs any data printed by the server. The Proxy.A tool creates a proxy connection so that the actor can route traffic through a proxy node to a destination node. The Eventlog tool parses event logs, dumps interesting logs and deletes incriminating logs. The tool can also end processes and delete itself. The Multipurpose tool edits event logs, dumps passwords, securely deletes files, encrypts files, enumerates the network, and assists the attacker in moving across the network.

The Butterfly group exhibits intense operational security. Many of their tools self-delete, and others are securely deleted by a GNU Shred tool used by the attackers. Event logs are modified or deleted to hide the intrusion. Uninteresting computers are fully purged of all traces of the attacker's presence. C&C domains are registered with disposable names and emails. Hosts of C&C servers are paid using the Bitcoin anonymous digital currency. Symantec observed that the group "uses encrypted virtual machines and multi-staged C&C servers" to make it more difficult to investigate their middle infrastructure. Symantec managed to track activity through proxies to a C&C server that was digitally sterilized. No activity was logged and the system featured Truecrypt and a Virtual Box virtual machine. Compromised systems were likely attacked from within the virtual machine; consequently, analysis is difficult when the image is not live.

Regin/ Prax/ WarriorPride

The Regin malware campaign targeted international organizations from 2008 to 2011 and from 2013-2014. The malware may have remained undiscovered for at least five years prior to 2008. The complexity of the toolkit suggests the investment of significant resources over several years. In support of this assumption, Symantec notes that Regin appears to be designed for espionage campaigns that last several years.

The malware is allegedly the product of a collaboration between the United States NSA and the British GCHQ. This allegation derives from a document leaked to Der Spiegel and the Intercept by Edward Snowden. The malware primarily targeted systems belonging to private individuals, small businesses, and telecommunications companies in Russia, Saudi Arabia, Mexico, Ireland, and to a lesser extent, India, Afghanistan, Iran, Belgium, Austria, and Pakistan.

Symantec notes that the framework has been used for mass surveillance against “government organizations, infrastructure operators, businesses, researchers, and private individuals. “ Nearly half of the attacks targeted private individuals. The quarter of the infections against telecommunication infrastructure was likely an attempt to gain access to the calls routed through the networks. Regin does not have a clear infection vector; though, Symantec suspects that some infections are the result of watering-hole attacks and zero-day exploits.

Regin consists of a trojan and a backdoor that are widely customizable to fit the target. The platform excels at remaining undetected and obfuscating its indicators of compromise. Regin is a modular platform, reminiscent of Flame, Duqu, and Stuxnet. The Regin backdoor is a five stage modular component and each stage after the first is hidden and encrypted. After each successful installation of a stage, the next stage is decrypted and installed. Each piece provides as little information as possible about the total component. If any stage fails then the installation terminates. The flexibility of the Regin platform means that the actor can customize the payload to the target. Consequently, Regin has dozens of discovered payloads and likely has many more that remain known only to the actor. In general, the platform features several remote access trojans (RATs), and tools to capture screenshots, log keystrokes, monitor network traffic, steal credentials, recover deleted files, and hijack the point and click functions of the mouse. According to Symantec, advanced payloads also contained “Microsoft IIS web server traffic monitor and a traffic sniffer of the administration of mobile telephone base controllers.” The platform also features anti-forensic capabilities, a custom-built encrypted virtual file system (EVFS), and RC5 encryption. Communication with the C&C servers occurs over ICMP/ ping, embedded commands in HTTP cookies, and custom TCP and UDP protocols.

Flame/ Flamer/ Skywiper

Flame is a modular malware discovered in 2012 by MAHER Center of Iranian National, Kaspersky Lab, Iran’s CERT, and CrySyS Lab of Budapest University of Technology and Economics. Flame may have been active for 2-5 years prior to its 2012 discovery. Initially, malware targeted the Microsoft Windows operating system of computers that supported the Iranian nuclear program. However, Iran discovered the malware after detecting a cyber campaign against its oil industry.

Flame is a large piece of modular malware, designed to map and monitor the target network. Flame is about 20 megabytes of code. For comparison, it is ~20 times the size of Stuxnet; though, Flame is entirely focused on espionage and is considered a predecessor to Stuxnet. The malware leverages the victim’s network to provide the adversary with a steady

stream of exfiltrated data that can be used to inform cyber and cyber-physical campaign decisions. Flame is too large and too complex to be anything except state-sponsored malware. Because of its alleged purpose, the Flame malware is attributed to a joint development program between the NSA, the CIA, and the Israeli military. Flame may have been part of a classified operation meant to monitor and slow Iran's nuclear program, code-named Olympic Games.

To the credit of the allegations, the Stuxnet malware was developed under similar circumstances and for similar purposes. In fact, Flame contains some of the same code as Stuxnet. According to Kaspersky senior researcher Roel Schouwenberg, "It's very likely it's two teams working effectively on the same program but using two very different approaches." Supposedly, the campaign against Iranian oil industry, which led to the exposure of Flame, was a unilateral operation launched by Israel, without informing their American counterparts.

Kaspersky detected Flame malware infections in Iran, Israel, Sudan, Syria, Lebanon, Saudi Arabia, and Egypt. Significantly fewer infected systems were detected in Europe or North America. Infected systems belonged to state-related organizations, educational institutions, and individuals. Systems were compromised via spear phishing attacks, infected websites, infected USB devices, and other infected systems on the local area network. Flame targeted emails, documents, AutoCAD drawings, instant messenger logs, and Skype conversations.

Flame is one of the first malware complex enough to be considered an attack toolkit. For years, Flame evaded detection by masquerading as a Microsoft software update. Flame creates its own backdoor, operates like a Trojan, and replicates across the local network and removable media like a worm. Flame contains many different libraries for compression (zlib, libbz2, and ppmd), for encryption (five methods total), for database manipulation (sqlite3), and a Lua virtual machine. The virtual machine is included to integrate components of Flame with C and C++ code on the host machine. Flame also contains local databases with nested SQL queries, Windows Incident Management scripting, batch scripting, and other features.

Flame set the precedent for the typical espionage malware capabilities. Flame can log keystrokes, it can activate microphones to capture audio, it can activate cameras to capture video, it can extract geolocation data from images, and it can screenshot the display. Recorded data is compressed via a public-source library and periodically sent through the malware operator's C&C infrastructure through a covert SSL channel. Other data is similarly exfiltrated. Flame is unique (at least for 2012) in that it can activate and use Bluetooth wireless to send and receive commands and data. Through Bluetooth, infected machines can be turned into beacons or used to detect nearby Bluetooth enabled devices. Like Stuxnet, Flame can infect

other systems on the network through shared connections such as printers. Flame can also spread to air-gapped networks via a USB drive.

The malware detects the antivirus on the host system and configures its modules and file names so that it has the greatest probability of remaining undetected. The malware also protects its modules with READ, WRITE, and EXECUTE permissions to make them inaccessible to user-made applications. Flame employed fake Microsoft licensing certificates to make discovered modules appear legitimate. Finally, Flame includes a Kill module that discretely removes the malware from infected systems. After public disclosure of the malware, the operators sent the kill command and removed the malware from many high profile hosts; thereby obfuscating the actual breadth of the campaign.

America's Most Elite Line of Cyber-Defense: Tailored Access Operations (TAO)

As the most targeted Nation in the world, The United States intelligence community has been continuously raising the bar to combat global bad actors. Tailored Access Operations is the largest operative component of the Signal Intelligence Directorate of the United States National Security Agency (NSA), consisting of over 1000 military and civilian cyber security professionals, hackers, technology specialists, and hardware and software designers. Approximately 600 of TAO's Computer Network Exploitation (CNE) operators work in rotating 24 hour, seven days a week, shifts out of the Remote Operations Center at Fort Meade.

The Office of Tailored Access Operations produces some of the best intelligence for the United States government and its work has been pivotal to the success of numerous operations. TAO is credited with delivering critical information to the 2007 U.S. Army operations in Iraq and in the 2007 operations to prevent Iran from obtaining nuclear weapons.

TAO is comprised of four main divisions. The Data Network Technologies Branch develops the infiltration and collection software utilized by the TAO. The Telecommunications Network Technologies Branch curates infiltration techniques. The Mission Infrastructure Technologies Branch combines the spyware and techniques to use in campaigns and they develop and build the computer and telecommunications hardware. The Access Technologies Branch, which contains personnel seconded by the CIA and FBI, performs "off-net operations." TAO is headed by U.S. Cyber Command and the director of the NSA.

The NSA describes TAO operations as computer network exploitation. TAO conducts counterterrorism and traditional espionage operations, but they also conduct cyber-attacks on behalf of the United States. Supposedly, TAO is able to compromise even the hardest targets. TAO is tasked with monitoring foreign entities, infiltrating their networks, and gathering information. It accomplishes its task through spyware or by compromising network devices such as routers, switches, or firewalls, and monitoring the network traffic. TAO is also tasked with developing malware or information profiles that would enable the United States to cripple foreign network infrastructure or telecommunications if directed to do so by President Obama.

The NSA is not authorized to conduct operations against domestic targets; however, some are concerned about the massive telecommunications monitoring programs that were revealed as a result of the Snowden leaks. The NSA monitors domestic traffic to capture communications in which at least one party originates from outside the United States. When CNE operators identify a network or system belonging to a nefarious foreign entity, they attempt to compromise its security, download a copy of its hard drive for analysis, and plant malware tools to monitor email and network traffic from the machine.

The main attack suite developed by the TAO and made public by the Snowden leak is dubbed QUANTUM. QUANTUM features a suite of attack tools that enable DNS injection attacks, HTTP injection attacks, and the ability to inject into MySQL connections. It also contains tools to hijack IRC and HTTP-based criminal botnets and tools to create phantom servers. The QUANTUMDEFENSE portion of the program searches tapped connections for DNS requests for NIPRnet addresses and initiates a packet-injection attack on a DNS reply to redirect the target to an NSA controlled site. This site may be a FOXACID server, which probes the victim's browser for weaknesses. The TAO can exploit any weaknesses with the QUANTUMINSERT program and seize control of the victim system. QUANTUMSMACKDOWN conducts packet injection attacks against attacks aimed at Department of Defense assets. QUANTUMCOOKIE is used to de-anonymize Tor users through web cookies and fetch requests. Finally, the QUANTUMSQRREL program lets TAO pose as any authenticated user on virtually any site by spoofing the IPv4 or IPv6 address of the host. Through this, TAO can monitor most digital communication, create posts from a "trusted" account, or pose as specific users in online transactions.

EQUATION Group

With operations predating at least 2001, EQUATION group is one of the most persistent and arguably, the most sophisticated threat groups in operation. EQUATION Group was discovered during Russian cyber-security firm, Kaspersky's investigation into the Regin threat group. Kaspersky attributes EQUATION Group to the United States National Security Agency; however, definitive evidence of attribution remains absent. EQUATION group's name derives from their employment of encryption and obfuscation strategies throughout their operations. The RC5 encryption algorithm is deployed throughout the malware and additional encryption algorithms RC6, RC4, and AES are added in other modules. Some of the attribution of the group to the United States comes from similarities between the malware platform and exploits to Stuxnet and the Gauss malware.

EQUATION Group has globally targeted more than 500 victims in over 30 countries including Iran, Russia, Syria, Afghanistan, Kazakhstan, Belgium, Somalia, Hong Kong, Libya, United Arab Emirates, Iraq, Nigeria, Ecuador, Mexico, Malaysia, United States, Sudan, Lebanon, Palestine, France, Germany, Singapore, Qatar, Pakistan, Yemen, Mali, Switzerland, Bangladesh, South Africa, Philippines, United Kingdom, India and Brazil. Targets are affiliated with government institutions, diplomatic organizations, the telecommunication sector, aerospace firms, energy companies, nuclear research facilities, oil and gas companies, military systems, nanotechnology research facilities, Islamic activists and scholars, mass media outlets, the transportation sector, financial institutions, and companies developing cryptographic technologies. It is possible that even more infections remain undiscovered. Kaspersky estimates that EQUATION Group attacked 2000 targets per month in 2008; although, the estimate seems generous. EQUATION Group's known C&C infrastructure spans more than 300 domains on over 100 servers.

The EQUATION Group compromises systems by using zero-day exploits, by infecting physical media (USB stick, CDs, etc.), through web-based exploits, through the self-replicating Fanny worm, and through robust customized malware platforms. The zero day exploits targeted Microsoft Windows, Internet Explorer, Java, the Firefox 17 browser, and the TOR browser. Attacks incorporating infected physical media utilize interdiction, a technique where an attacker intercepts shipped goods, such as software, and replaces it with a version containing malware or backdoors, before sending it to the buyer. EQUATION Group has been known to exploit vulnerabilities in Java on popular websites to facilitate the delivery of one of its validator-style Trojans, DOUBLEFANTASY and TRIPLEFANTASY. The Fanny worm was created around 2008 and it was used to gather information from targets in the Middle East and Asia. According to Kaspersky, 59.36% of Fanny infections were in Pakistan, 15.99% of Fanny infections were in Indonesia, 14.17% of Fanny infections were in Vietnam, and 4.05% of Fanny

infections were in China. Networks are typically infected with the Fanny worm via infected physical media. Fanny resembles Stuxnet in operation, but it may actually predate Stuxnet and tie the EQUATION Group to the Stuxnet Group. Some variations of Fanny feature the Stuxnet LNK exploit and other exploits that were deployed in Stuxnet, and the Flame malware; however, it appears that the exploits were used in the Fanny worm prior to their inclusion in Stuxnet or Flame. Considering that Stuxnet and Flame were so effective because they employed zero-day exploits that were unknown to the public, there is merit to the theory that Stuxnet was created by or in collusion with the developer of the Fanny worm. Fanny is used to map air-gapped networks. USB devices (and other writable media) that are plugged into infected systems, are corrupted to store Fanny in a self-hidden partition. When the device is plugged into an air-gapped system, say for updates, basic system information or data is stored in the hidden partition. The information is exfiltrated to a C&C server the next time the device is plugged into a system with an internet connection. EQUATION Group can also store commands on the device while it is connected to the internet. Fanny will execute the commands the next time the device is connected to the air-gapped system. This process allows the group to map the network infrastructure and it allows the group to compromise air-gapped systems, which tend to contain more sensitive information. These systems are often less defended because their administrators equate their isolation to security.

The EQUATION Group developed unique malware and malware platforms. Typically, a zero-day exploit or a web exploit was used for the initial compromise of the target system. Next, a validator-style Trojan, dubbed DOUBLEFANTASY scans the infected system and uses input criteria to determine if the host is the intended system or if the characteristics of the system indicate that its data would be interesting to the attacker. DOUBLEFANTASY acts as a backdoor into the target system. If the target matches the criteria, then a malware platform, EQUATIONLASER, EQUATIONDRUG, or GRAYFISH is delivered and installed on the system. For example, in one campaign, EQUATION Group exploited a vulnerability in the PHP script used in an online Islamic Jihadist discussion forum. However, only systems belonging to users who were logged into accounts and whose traffic originated from a specific IP address range corresponding to Jordan, Turkey, and Egypt, were infected with malware installers. More recently, DOUBLEFANTASY has been upgraded into a more robust backdoor, called TRIPLEFANTASY.

The EQUATIONLASER platform was used from 2001 – 2003 to infect Windows 95 and Windows 98 systems. The EQUATIONDRUG platform replaced EQUATIONLASER in 2003, and was used until at least 2013. EQUATIONDRUG supports modular plugins, which can be dynamically uploaded and unloaded by remote attackers. EQUATIONDRUG installs with a cadre of modules that give full control of the operating system to the attacker. Further, it supports the addition of new plugins to increase its functionality. So far, at least 35 different

plugins and 18 drivers have been discovered. EQUATIONDRUG was designed to compromise Windows 95, Windows 98, and Windows ME. Since the malware does not have a trusted digital signature, it may not be able to run on a modern operating system. Legacy systems, prevalent in the public sector, are still at risk. Information gathered by EQUATIONDRUG tools is stored in fake fonts folders under the Windows/ Font file directory. If EQUATIONDRUG does not receive commands from an adversarial C&C server after a specified time, usually a month or two, then it deletes itself from the system.

Sometime between 2008 and 2012, EQUATIONDRUG appears to have been phased out in favor of the GRAYFISH malware platform. GRAYFISH is the most sophisticated Equation Group malware platform discovered. Upon delivery of the installer via TRIPLEFANTASY, a GRAYFISH bootkit is injected into the registry of the operating system. When a computer first powers on, the operating system code executes (booting up) and it enables the majority of the functionality of the system. When an infected system is powered on, GRAYFISH injects code into the boot record so that it can control every stage of the Windows launch process. GRAYFISH, its virtual file system, its stolen information, and its functional modules are stored in the registry of the system. Because everything is stored in the registry and GRAYFISH and its modules are dynamically decrypted and executed by the bootkit, there are no malicious executables contained in the user's filesystem. This means that the user cannot detect the GRAYFISH malware on the system; at least not with traditional anti-malware tools. During the bootup process, GRAYFISH processes through 4-5 layers of decryption where each layer triggers the execution of the next layer of decryption. If all of the layers successfully decrypt, then GRAYFISH executes its code and the malware silently runs on the machine. If even one layer fails to decrypt during launch, then GRAYFISH proceeds to delete itself from the system. This technique confounds analysis and makes GRAYFISH infection difficult to discover because the malware might delete itself the moment the user detects anomalous behavior and begins diagnostic procedures.

One reason that EQUATION Group is considered far more sophisticated than any other advanced persistent threat actor is the capability of modules contained in the EQUATIONDRUG and GRAYFISH platforms to reprogram hard-drive firmware. This allows for unprecedented persistence. Security firm F-Secure notes that this rarely seen module might be Tailored Access Operations IRATEMONK program which affects hard-drives produced by Seagate, Maxtor, Western Digital, Samsung, IBM, Micron, and Toshiba.

France

Animal Farm

Animal Farm is the first French speaking APT detected. It is worth noting that French is the official language of 29 countries. According to slides referencing Operation Snowglobe, released by Edward Snowden and Der Spiegel in January 2015, Animal Farm is a cyber threat group sponsored by France. The group is suspected to be a component of the French Directorate-General for External Security (DGES), which is France's external intelligence agency. The group began development of its toolkit in 2007 and it has been actively launching attack campaigns since 2009. The purpose of the group is to conduct cyberespionage and denial of service campaigns against political targets using traditional cyber-attack vectors, 0-day exploits, and a custom multitier malware platform.

Animal Farm targets government entities, activists, private companies, journalists, media outlets, and defense contractors in Syria, Iran, Malaysia, the United States, China, Turkey, the Netherlands, Germany, Great Britain, and Russia with spear phishing and watering hole attacks.

The Animal Farm trojans can be grouped into six families. The NBot malware is a standard botnet kit capable of enslaving systems and leveraging their resources in aggregate to conduct DDoS attacks. The EvilBunny trojan and its variants are validator trojans that were used in spear phishing attacks in 2011. The trojans were delivered through malicious PDF files through the 0-day exploitation of a vulnerability in Adobe reader. The trojan checks whether an emulator is running, what directory it is running from, whether its payload timestamp has been changed, and what time the API hook was detected. Bunny is designed as an execution platform for the attacker to inject Lua scripts into victim system processes.

The Casper and Tafacalou trojan families are also validator trojans. Casper is designed to persist and to track victim online activity. Casper is delivered via watering-hole attacks while Tafacalou may be delivered through spear phishing or watering-hole attacks. The Tafacalou malware is the used to deliver either the Dino espionage platform or the Babar espionage platform onto the victim host.

Babar is a spyware toolkit capable of logging keystrokes, monitoring web activity, taking screenshots, capturing audio, copying clipboard data and eavesdropping on online conversations that are conducted over popular messaging platforms (Skype, MSN, Yahoo messenger, etc.). Babar obfuscates its activity by hooking into the APIs of remote processes

through a series of named pipes. Babar may have been used to spy on Iranian nuclear research facilities and European financial institutions.

Dino is a modular malware capable of executing C2C commands and Windows batch commands, searching for specific files, uploading and downloading files from the C2C infrastructure, scheduling its own command executions, killing processes, and removing itself from the victim system. The PSM module is the encrypted on-disk copy of Dino's components. The CORE module stores configuration and the ENVVAR module stores environment variables. The CRONTAB module schedules tasks. Meanwhile, the CMDEXECQ module stores the queue of commands executed by the CMDEXEC component. Finally, the FMGR module manages file uploads and downloads.

Israel:

Duqu/ DQ

Duqu was discovered on September 1, 2011 by CrySys Lab of the Budapest University of Technology and Economics in Hungary. The code of the malware is very similar to Stuxnet and it is believed to be either the product of a sister-project or a derivative of the Stuxnet source code. In particular, the kernel driver of the malware is practically the same as the kernel driver of Stuxnet (commonly named JMINET.SYS and MRXCLS.SYS respectively). The former case implies that the malware was developed and deployed by a state sponsor, likely the United States or Israel. Meanwhile, the latter case expands attribution to practically any well-resourced actor on the internet.

Unlike Stuxnet, Duqu was not meant to sabotage the host systems; instead, like most modern malware, its purpose was covert information exfiltration. The Duqu malware was found on similar target systems as Stuxnet, so it is reasonable to conclude that it was likely developed and deployed to collect information pertinent to current events or information necessary to launch future espionage or sabotage campaigns. Duqu primarily targeted the industrial infrastructure of system manufacturers, and the industrial sector in Middle Eastern countries. The adversary exfiltrates confidential documents such as design specifications and network information, likely to aid in future attack campaigns.

The original Duqu components exploited a zero-day Microsoft Windows 32k TrueType font vulnerability (CVE-2011-3402). The vulnerability permits the attacker to execute code at the highest privilege level. A portion of the malware, dubbed the "Duqu framework" by Kaspersky, appears to be written in C with a custom object oriented framework and compiled

in Microsoft Visual Studio 2008. Duqu consists of an installer, a driver file, a DLL with embedded files, and a configuration file. The installer registers the driver as a service that starts at system initialization. The driver injects the DLL into the Windows process services.exe. From there, the DLL extracts the other components and injects them into other Windows processes. Sometimes the driver file is signed with a valid digital certificate to avoid detection.

Duqu was typically configured to run on an infected machine for 30-36 days. Unlike Stuxnet (a worm), Duqu (a Trojan) does not replicate and spread on its own. It only spreads through additional breaches and targeted installation. In service to its espionage function, Duqu's components mostly log keystrokes and system information. According to Kaspersky Lab, the Duqu operators were particularly intent on collecting passwords, stealing documents, and taking desktop screenshots. The Infostealer component collects information and then stores it in a local encrypted and compressed file. At regular intervals or upon request, the file is attached to a dummy .jpeg file and uploaded from the infected host.

Duqu communicates with its C2C infrastructure through HTTP(s). Each attack, in at least eight different countries, used a different C&C server. The servers, likely proxies, forwarded all port 80 and port 443 traffic to other servers, which in turn forwarded traffic to other servers, and so on. The servers also contained at least three different DLLs and the infostealer component used to collect information from the infected hosts. Most of the known infrastructure went inactive when the malware was exposed.

In late 2015, Kaspersky reported the reemergence of the Duqu malware, targeting western countries as well as the Middle East and Asia. Many of the targets were affiliated with the P5+1 events and venues associated with the Iran nuclear deal negotiations. An event honoring the 70th anniversary of the liberation of Auschwitz-Birkenau may also have been targeted. The recent attacks leveraged new 0-day exploits including CVE-2015-2360, which targets the Windows kernel. Duqu 2.0 runs as kernel level code. The updated malware survives in the system memory of infected servers and re-downloads onto desirable hosts upon reboot. Duqu 2.0 was built on top of the original code; however, it now loads from an MSI file, and has at least 94 more plugins. The known plugins allow the adversary to customize the toolkit to the target environment to circumvent system security and incompatibilities. Newer components appear to be written in the C++ programming language.

It is unclear if the new version is deployed by the original team. Strings in the code suggest that the malware was developed by English speakers; however, a few minor spelling errors could suggest the involvement of non-native speakers. Additionally, Kaspersky observed a target system that was infected by both Duqu 2.0 and Equation group malware. This suggests a lack of coordination and possibly competing interests. Since the source code of

Duqu was never made public, the revised version had to have been developed by the one of the original authors. If Duqu was developed by both the United States and Israel, and if Equation Group is not coordinating with Duqu, then one could postulate that Duqu 2.0 was developed by the Israeli development team in an effort to gain information about the US-Iran nuclear negotiations. Other explanations exist.

Unknown Nationality:

Hellsing

The Hellsing group targets government and diplomatic organizations in the APAC region, particularly organizations located in nations along the South China Sea. Most targets are from Malaysia, the Philippines, Indonesia, and India. Hellsing malware samples were primarily compiled in either UTC+6 or UTC+8. Typically, Hellsing infects targets through spear phishing emails containing password protected RAR, ZIP, and 7ZIP archives. The passwords are sent in the emails to the target. Locking the archives bypasses some security features such as Gmail scans. Hellsing was discovered when Kaspersky Lab was investigating the Naikon group and found that Hellsing had responded to a 2014 spear phishing email from Naikon with a custom backdoor. It is not clear whether Naikon intentionally targeted Hellsing or if Hellsing actually managed to infect Naikon; however, it is clear that Hellsing took the attempt as an attack and responded with an escalated attack. Hellsing responded to the spear phishing request for information with a series of inquisitive exchanges, pressing Naikon's assumed identity (as an employee of the secretariat division of the government of the assumed target nation) and fake credentials. The conversation demonstrates that the Hellsing members are more proficient in English than the Naikon group. Finally, Hellsing emailed back a "confidential" locked RAR and the accompanying password. The archive contained two PDFs and a malicious SCR file. The latter file was a backdoor specifically customized to target the Naikon group.

The backdoor can upload and download files, update itself, and uninstall itself. Each instance of the backdoor has a command and control server, a version number, and a campaign or victim identifier. The same Hellsing backdoor has been seen in attacks against ASEAN related entities in the South China Sea region. Some of the Hellsing infrastructure overlaps with an APT group tracked internally by Kaspersky, dubbed PlayfulDragon/ GREF, while other portions of the infrastructure coincide with the Mirage APT group and the Vixen Panda group.

After Hellsing establishes a variant of its backdoor, it deploys information-gathering tools. One tool, test.exe, gathers system information and tests available proxies. Another tool, xkat.exe, operates from the Dbgv.sys driver to delete files and kill processes. Kaspersky Lab

claims to have seen the tool used to remove malware from competitor groups from Hellsing victim systems.

Moker

In October 2015, Israeli cyber-security firm Ensilo discovered a remote access Trojan (RAT), dubbed Moker, inside the sensitive network of a customer. A RAT is not an APT. Malware is the tool that supports the APT campaign. However, Ensilo contends that the RAT is complex enough to suggest that it may be developed and deployed by an emerging APT group. The quality of the code is high. The code checks its return values, validates its pointers, handles its exceptions, and prevents buffer overflows. The malware also contains obfuscation measures to inhibit deconstruction and analysis attempts. Since the digital signatures of the malware did not register on Virus Total (a research tool for recognizing malware signatures), and because the malware itself contains features dissimilar to other campaigns, there is the possibility that the security firm either uncovered an undiscovered malware campaign or that they caught a threat as it emerged.

Neither the identity of the developer of the malware nor the infection vectors are known. The malware targets the operating system of Microsoft Windows hosts. The single sample of the malware discovered communicated with a domain that corresponded to a HTTP server in Montenegro. Based on its efforts to communicate with the C2 infrastructure, Ensilo postulates that the server is owned by the attacker who hosts C2 infrastructure via a Virtual Private Server (VPS) or a static IP rather than a hacked domain or a shared hosting server.

Moker is a remote access Trojan (RAT) capable of seizing complete control of the victim system. It generates a new administrative user account and it opens a RDP channel to allow the adversary to remotely access the infected system. If the remote desktop service is disabled, the malware will attempt to enable it as a background service. Moker establishes a persistent residence in the operating system files so that it appears a legitimate OS level process with system wide privileges and access to system settings. In operation, the malware injects its malicious code into the legitimate code of different system processes. In particular, it targets Explorer.exe, Svchost.exe, and csrss.exe. In order to execute code without the user's consent and at higher privileges, Moker either infects a program that already runs at elevated privileges or it exploits a flaw in the design of Windows to elevate the privilege of the DLL. In the latter case, Windows always loads certain DLLs from the system directory at escalated privileges; as a result, Moker writes a file named "ActionQueue.dll" into the "sysprep" directory so that the malware always runs with elevated privileges. Afterward, the malware modifies system sensitive files and system security settings so that it remains undetected for

as long as possible while maintaining access to the greatest amount of access to the system files. The malware itself is capable of recording HTTP(s) traffic, taking screenshots, logging keystrokes, and exfiltrating files. The malware also enables the attacker to use the infected machine as a proxy server, similar to a Socks server, so that the adversary can navigate the local network.

The malware contains a hidden control panel module which allows the adversary to direct it and access the malware without an active internet connection. Consequently, the malicious actor can exploit a VPN connection and legitimate, stolen user credentials to operate the malware on infected air-gapped systems. The local access panel may have been an intentional feature added by the developer so that malicious activity might be confused for the activity of a legitimate employee who VPN'ed into the system. Alternately, it could be a developer tool that was mistakenly left in the malware.

The malware features significant anti-analysis and anti-debugging techniques to inhibit deconstruction and investigation of the form and functionality of the malware. Moker bypasses or disables antivirus, renders Microsoft Windows user access controls ineffective, and confounds sandboxing and virtual machine analysis with encryption and multistage installation. The malware evades signature based anti-virus and network monitoring solutions by compressing its code. In an attempt to prevent sandbox and virtual analysis, Moker installs in two stages. The first stage dropper contains no malicious code and only delivers the malware infrastructure. Upon successful installation and validation that the infected environment is a legitimate target, the first dropper calls a C2 server or a local directory for the second stage delivery. The second stage of the malware is the malicious payload containing encrypted malware files and system monitoring tools. If the environment is confirmed a legitimate target, then the first stage malware decrypts the payload and injects the malware into the victim system processes. The malware also contains complex code and instructions that do nothing except deter deconstruction and analysis attempts.

Shrouded Crossbow

The Shrouded Crossbow group has been active since 2010, typically targeting companies that are close to governments and key industries in Asia. Common targets include government contractors, privatized government agencies, companies involved with consumer electronics, the computer industry, the healthcare sector, and financial industries. The malicious team is predicted to be about ten people, equipped with significant resources. Rather than develop its own attack kits and malware, the group uses its significant resources to purchase source code

and tools from other authors. Afterward, members of the group improve the code to suit their specifications.

The group employs the BIFROSE/ Bifrost, KIVARS, and XBOW backdoors in their attacks. As an indicator of resources available to the group, Trend Micro notes that BIFROSE has sold for more than \$10,000 on underground sites. BIFROSE has been around for about a decade and has been used in spam campaigns against NATO and United States government agencies. BIFROSE is a remote access Trojan (RAT) which establishes a persistent presence and then deploys tools to capture keystrokes, screenshots, and confidential information. Trend Micro actually believes that the group purchased the source code of BIFROSE, and then developed a new installer, created unique loader-backdoor pairs, and simplified the backdoor capabilities, thereby resulting in KIVARS. KIVARS is also available as a 64-bit variant. The group developed XBOW on their own, based on BIFROSE and KIVARS. The malware is delivered via spear phishing emails containing malicious .RAR files or .EXE. The email topics are generally breaking news, resumes, government data, or meeting requests. The malware corresponds to a C&C network of about 100 servers registered to free dynamic DNS or discrete IP addresses. The C&C servers appear to be organized according to the actor's use. IP address changes and renewal of domains happen according to an organized schedule. Trend Micro suspects that in addition to the 10-member development team, the malware group may employ separate teams to design and deploy the malicious emails and to maintain the C&C infrastructure.

Santa APT

Cloudsek, a Canadian Cyber Security firm, detected the activity of a suspected criminal advanced persistent threat group over the 2015 holiday season. The group, dubbed Santa APT because some of their malware masqueraded as Santa Claus applications, steals intellectual property for economic gain. Cloudsek believes that the malware developers are located in South Asia.

The group came to the attention of security professionals who noticed them selling information stealer malware, capable of jumping air gapped systems, on underground markets. The attackers were using the malware to steal classified data from software companies and government organizations. The malware collects files and screenshots and stores them in hidden files on any connected USB device. When the device is connected to an internet enabled system, the data is sent back to command and control infrastructure located in Germany. Empty voice recording and key log files on the C2C servers suggest that the malware is still under development. Cloudsek claims to have found the malware attributed to the group masquerading as Santa Claus mobile games, which had infected about 8000 systems. The malware stole contact lists, SMS messages, call records, location information,

calendars, pictures, video, environment readings, camera specifications, browser history, program information, sim card information, and device status. The mobile malware communicated with the C2C infrastructure via HTTP about once a minute. The C2C servers corresponding to the mobile infrastructure had separate login sections for user profiles and for administrative profiles. The victim information was organized according to user and then according to data type. The attacker could also arm the malware to send them an SMS alert if the victim left a regional area. This could allow the actor to track whether a victim has left home or the office, in real time. The adversary could also receive regular updates if a particular victim received an SMS message or phone call.

Cloudsek used passive DNS to track the group activity of a South Asian company that sells spy software to monitor employees. The company is recruiting mobile application developers for iPhone and Android.

Conclusion:

The conglomeration of hacktivists, state sponsored hackers and cyber mercenaries are continuously targeting American corporations, organizations, Universities and government networks. The malicious element is winning because the United States lacks proper cyber hygiene and has yet to expedite a path to a cybersecurity-centric culture. Metaphors matter as the language to describe cyberattacks today shape the legislative community's constitutional adherence in future policy. The reader is cautioned to be weary of the new cliché "Cyberwar" continuously being used as a kneejerk reaction in times of panic. As we experience warlike tactics being used by a wide variety of bad actors (with a multitude of motivations) in a cyber setting it will be important to distinctively separate what defines cyberwar from cyber conflict, cyberattack and cyber espionage as each term holds a different variant of retribution and/or penalties.

American industry as a whole is an easy target because seasoned adversaries are breaching virtually defenseless networks. Organizations are encouraged to follow, at a minimum the latest NIST Standards for Critical Infrastructure Cybersecurity. A vigilant approach to cyber and social engineering education, application patching, technical abnormality notifications etc. are paramount for organizations striving to minimize attack surface and maximize defenses.

Even with multi-factor authentication and cybersecurity protocols in place, breaches will happen. Optimized cybersecurity strategies will use early warning mechanisms such as behavioral analytics and behavioral biometrics coupled with multilayered encryption. This 'Tar Pit' method will slow down a breach, alert the proper administrator and minimize

threat. That said, even the most robust cybersecurity strategy is useless if the bad actor has obtained legitimate admin credentials, therefore education is essential. Spear Phishing is one commonality shared by a majority of hacking events. Spoofed URL's, watering hole attacks etc. are all dependent on getting the target to click on a link or open an attachment that carries the malicious code that will infect one's network; the objective is to train staff to identify these subtleties that will have catastrophic impact.

Targeted advanced persistent threats will continue to multiply and become more sophisticated. Optimal application of the most up-to-date defense technologies is the first step in demotivating hackers from attempting to breach one's network as attackers will typically pick the path of least resistance and complexity. Understanding the enemy and learning from past mistakes while planning for new threats based on reliable research must be part of every association's cyber strategy.

Appendix I:

Terms

Malware – Malware is the catch-all term for any malicious code. Malware can take the form of viruses, Trojan horses, worms, ransomware, spyware, adware, scareware, or malicious programs.

Virus – A computer virus is malicious code that replicates itself when executed, and may infect other programs or systems.

Trojan – A Trojan horse is a malicious program that tricks the user into installing it, by misrepresenting itself as a useful or desirable process or program.

Worm – A computer worm is a self-replicating malicious program that may spread to other computers and other networks.

Rootkit – A rootkit is malicious software that obfuscates its existence and that enables an attacker to access a system or its files.

Vulnerability – A vulnerability consists of a flaw in system (e.g. a flaw in the code), attackers' access to that flaw, and the attackers' ability to exploit the flaw.

Zero-day Vulnerability – A software flaw is that present at launch, but unknown to the software vendor. Often, zero-day vulnerabilities are repaired by the vendor through software patches.

Zero-day Exploit – Zero-day exploits are harmful vulnerabilities that are not discovered or are not repaired by the vendor. Zero-day exploits tend to be rare and expensive since the knowledge of their existence must remain secret lest vendors repair the vulnerability.

Backdoor – A hidden program, or program component, that allows unauthorized remote access to a computer.

Registry – The system registry is the database of system hardware information, profile information, installed programs, and settings.

Privilege Escalation – Privilege escalation is the exploitation of a bug, design flaw, or configuration oversight which grants elevated access to resources that are normally protected from an application or user.

Command and Control Server – Command and control servers are also referred to as a C2 server or a C&C server. C2 servers are the centralized system that issues commands and receives outputs from infected machines (a botnet).

Legacy System – Legacy systems are old or outdated systems that are not compatible with modern applications or programs.

SCADA System – A Supervisory Control and Data Acquisition (SCADA) system operates with coded signals over communication channels to provide control of remote equipment.

Virtual File System – A virtual file system is an abstract layer on top of the user file system. A virtual file system allows its user to access applications across multiple file systems.

Air-gapped System – An airgapped system is not connected to the internet and is not directly connected to any systems that are connected to the internet.

Exfiltration – The process of removing data from a system.

Adversary – The attacker, hacker, enemy nation-state, or malicious actor targeting a system.

Advanced Persistent Threat – A group of attackers or developers who are sophisticated, persistent, and who have access to significant resources.

Common Attack Vectors

Phishing (spam) – Most breaches are the result of human error, such as an employee opening a malicious email. Phishing campaigns consist of sending massive amounts of malicious emails which contain either malicious links or malicious attachments. The user system is infected with malware if they follow the link to a landing page or if they open the attachment. Even though most recipients of the email will ignore it, phishing is successful because it is cheap and the relative gain of even one infected computer resulting from millions of sent emails is high.

Spear Phishing – Spear Phishing is the process of sending tailored emails to specific targets of value to the attackers. Spear phishing emails require the attacker to know more information about the target and as a result, they can be very convincing, even to trained security professionals.

Watering hole Attack – In a watering hole attack, the adversary either infects or spoofs websites that are often visited by members of the target organization. When users visit the website, the adversary can infect their system with malware.

USB/ Air-gapped Attack – Airgapped network attacks are sophisticated techniques of infecting systems that are not connected to the internet, and in some cases not connected to other systems, with malware. Attackers can infect storage media in hopes that it will be plugged into the system, they can infect software updates to the system, or they can infect systems connected to the target and exploit the connection to install malware.

This Brief was authored by:

- James Scott (ICIT Senior Fellow – Institute for Critical Infrastructure Technology)
- Drew Spaniel (ICIT Visiting Scholar, Carnegie Mellon University)

Contact Information

Legislative Branch Inquiries:

- James Scott, Senior Fellow, ICIT (james@icitech.org, 202-774-0848)

Federal Agencies, Executive Branch and Fellow Inquiries:

- Parham Eftekhari, Senior Fellow, ICIT (parham@icitech.org, 773-517-8534)

Links

Website: www.icitech.org

Social Media:



<https://twitter.com/ICITorg>



<https://www.linkedin.com/company/institute-for-critical-infrastructure-technology-icit->



<https://www.facebook.com/ICITorg>

Sources

ARS Technica:

<http://arstechnica.com/security/2014/11/sony-pictures-hackers-release-list-of-stolen-corporate-files/>

<http://arstechnica.com/security/2015/06/us-army-website-defaced-by-syrian-electronic-army/>

<http://arstechnica.com/security/2015/02/how-omnipotent-hackers-tied-to-the-nsa-hid-for-14-years-and-were-found-at-last/>

<http://arstechnica.com/security/2015/03/new-smoking-gun-further-ties-nsa-to-omnipotent-equation-group-hackers/>

<http://arstechnica.com/tech-policy/2015/09/cia-officers-pulled-from-china-because-of-opm-breach/>

<http://arstechnica.com/security/2015/09/dhs-infosec-chief-we-should-pull-clearance-of-feds-who-fail-phish-test/>

<http://arstechnica.com/security/2015/09/us-counterintelligence-czar-tells-government-employees-raise-your-shields/>

<http://arstechnica.com/security/2015/08/china-and-russia-cross-referencing-opm-data-other-hacks-to-out-us-spies/>

<http://arstechnica.com/security/2013/03/the-worlds-most-mysterious-potentially-destructive-malware-is-not-stuxnet/>

<http://arstechnica.com/security/2015/09/seven-years-of-malware-linked-to-russian-state-backed-cyberespionage/>

<http://arstechnica.com/security/2015/09/how-highly-advanced-hackers-abused-satellites-to-stay-under-the-radar/>

The Atlantic:

<http://www.theatlantic.com/international/archive/2011/08/syrias-digital-counter-revolutionaries/244382/>

BBC News:

<http://www.bbc.com/news/technology-30189029>

Beta News:

<http://betanews.com/2015/04/22/anonymous-lulzsec-guardians-of-peace-a-guide-to-the-most-notorious-hacking-groups/>

Bloomberg Business:

<http://www.bloomberg.com/bw/articles/2013-05-23/how-the-u-dot-s-dot-government-hacks-the-world>

<http://www.bloomberg.com/news/articles/2013-06-25/s-korea-president-s-websites-closed-for-review>

Berkeley Varitronics Systems:

<https://www.bvsystems.com/WordPress/?tag=guardians-of-peace>

Breaking Malware:

<http://breakingmalware.com/malware/moker-part-1-dissecting-a-new-apt-under-the-microscope/>

<http://breakingmalware.com/malware/moker-part-2-capabilities/>

Cloudsek:

<https://www.cloudsek.com/announcements/blog/apt-malware-masquerade-as-christmas-apps-and-santa-claus/>

ComputerWorld:

http://www.computerworld.com/article/3014474/security/cyberspy-group-resurrects-12-year-old-bifrose-backdoor.html#tk.rss_all

CNet:

<http://www.cnet.com/news/us-army-website-offline-after-hack-by-syrian-electronic-army/>

CrowdStrike:

<http://blog.crowdstrike.com/ironman-deep-panda-uses-sakula-malware-target-organizations- multiple-sectors/>

<http://blog.crowdstrike.com/cat-scratch-fever-crowdstrike-tracks-newly-reported-iranian-actor-flying-kitten/>

Cyber War Zone:

<http://cyberwarzone.com/whitepaper-russian-cyber-espionage-campaign-sandworm-team-2014-free-download/>

Cylance:

http://www.cylance.com/assets/Cleaver/Cylance_Operation_Cleaver_Report.pdf.

Deepdotweb:

<https://www.deepdotweb.com/2014/11/20/apt-attacks-via-tor-network-onionduke/>

Electronic Frontier Foundation:

<https://www.eff.org/deeplinks/2014/03/new-nsa-slides-reveal-tailored-access-run-amok>

Ensilo:

<http://blog.ensilo.com/moker-a-new-apt-discovered-within-a-sensitive-network>

Enigma Software:

<http://www.enigmasoftware.com/mirage-removal/>

Epoch Times:

<http://www.theepochtimes.com/n3/1914960-china-security-in-cybersecurity-the-chinese-regime-has-become-the-boy-who-cried-wolf/>

F-Secure Labs:

<https://www.f-secure.com/weblog/archives/00002718.html?tduid=ff8c6c422cb66b85a8ae21edb2f35886>

https://www.f-secure.com/documents/996508/1030745/blackenergy_whitepaper.pdf
secure black energy

<https://www.f-secure.com/documents/996508/%201030745/CozyDuke>.

https://www.f-secure.com/documents/996508/1030745/dukes_whitepaper.pdf.

FireEye:

<https://www2.fireeye.com/rs/848-DID-242/images/rpt-apt29-hammertoss.pdf>.

Forbes:

<http://www.forbes.com/sites/katevinton/2015/06/08/syrian-electronic-army-claims-responsibility-for-hacking-army-website/>

GData:

https://public.gdatasoftware.com/Web/Content/INT/Blog/2014/02_2014/documents/GData_Urobuos_RedPaper_EN_v1.pdf.

HIS talk:

<http://histalk2.com/2015/02/09/readers-write-fact-and-fiction-about-anthems-breach/>

The Huffington Post:

http://www.huffingtonpost.com/2015/05/14/washington-post-hacked-syrian-army_n_7285382.html

IbTimes:

<http://www.ibtimes.com/us-confirms-blackenergy-malware-used-ukrainian-power-plant-hack-2263008>

Information Week Dark Reading:

<http://www.darkreading.com/security-companies-team-up-to-take-down-chinese-hacking-group/d/d-id/1317006>

<http://www.darkreading.com/attacks-breaches/with-operation-clever-iran-emerges-as-a-cyberthreat/d/d-id/1317861>

Infosec Institute:

<http://resources.infosecinstitute.com/equation-group-apt-tao-nsa-two-hacking-arsenals-similar/>

<http://resources.infosecinstitute.com/animal-farm-apt-and-the-shadow-of-france-intelligence/>

International Business Times:

<http://www.ibtimes.com/deep-panda-group-wasnt-behind-massive-opm-hack-other-chinese-hackers-were-fireeye-1975658>

<http://www.ibtimes.com/fbi-formally-blames-north-korea-sony-hack-chinese-involvement-under-investigation-1763579>

ISight Partners:

<http://www.isightpartners.com/2014/10/sandworm-team-targeting-scada-systems/>

<http://www.isightpartners.com/2014/07/weeks-threatscape-media-highlights-update-14/>

Kaspersky Lab

<http://www.kaspersky.com/about/news/virus/2015/equation-group-the-crown-creator-of-cyber-espionage>

<http://www.kaspersky.com/about/news/virus/2015/Blue-Termite-A-Sophisticated-Cyber-Espionage-Campaign-is-After-High-Profile-Japanese-Targets>

<http://usa.kaspersky.com/internet-security-center/threats/cosmicduke-malware-virus-definition#.VpYhnFLKM40>

http://www.kaspersky.com/about/news/virus/2013/Kaspersky_Lab_Identifies_MiniDuke_a_New_Malicious_Program_Designed_for_Spying_on_Multiple_Government_Entities_and_Institutions_Across_the_World

http://www.kaspersky.com/about/news/virus/2012/Resource_207_Kaspersky_Lab_Research_Proves_that_Stuxnet_and_Flame_Developers_are_Connected

<https://blog.kaspersky.com/billion-dollar-apt-carbanak/7519/>

<http://www.kaspersky.com/about/news/virus/2014/Kaspersky-Lab-sheds-light-on-Darkhotels-where-business-executives-fall-prey-to-an-elite-spying-crew>

Krebs on Security:

<http://krebsonsecurity.com/tag/the-elderwood-project/>

The New Yorker:

<http://www.newyorker.com/tech/elements/syrias-other-army-how-the-hackers-wage-war>

NCC Group:

<https://www.nccgroup.trust/uk/about-us/newsroom-and-events/blogs/2014/july/a-new-flying-kitten/>

Novetta:

http://www.novetta.com/files/5614/1329/6232/novetta_cybersecurity_exec_summary-3.pdf

Recorded Future:

<https://www.recordedfuture.com/russian-malware-analysis/>

Reuters:

<http://www.reuters.com/article/2015/06/21/us-cybersecurity-usa-deep-panda-idUSKBN0P102320150621>

<http://www.reuters.com/article/2014/12/05/us-sony-cybersecurity-northkorea-idUSKCN0JJ08B20141205>

SC Magazine:

<http://www.scmagazine.com/researchers-observe-animal-farm-group-using-variety-of-malware/article/402477/>

Schneier on Security:

https://www.schneier.com/blog/archives/2013/12/more_about_the.html

Secure List:

<https://securelist.com/blog/research/68750/equation-the-death-star-of-malware-galaxy/>

https://securelist.com/files/2015/02/Equation_group_questions_and_answers.pdf.

https://securelist.com/files/2015/06/The_Mystery_of_Duqu_2_0_a_sophisticated_cyber_espionage_actor_returns.pdf.

<https://securelist.com/analysis/publications/69953/the-naikon-apt/>

<https://securelist.com/blog/research/71443/minidionis-one-more-apt-with-a-usage-of-cloud-drives/>

<https://securelist.com/blog/research/71876/new-activity-of-the-blue-termite-apt/https://securelist.com/analysis/publications/36462/stuxnetduqu-the-evolution-of-drivers/>

<https://securelist.com/blog/incidents/34344/the-flame-questions-and-answers-51/>

<https://securelist.com/blog/incidents/32855/flame-bunny-frog-munch-and-beetlejuice-2/>

<https://securelist.com/blog/research/68732/the-great-bank-robbery-the-carbanak-apt/>

<https://securelist.com/analysis/publications/69567/the-chronicles-of-the-hellsing-apt-the-empire-strikes-back/>

https://securelist.com/files/2014/11/darkhotel_kl_07.11.pdf.

<https://securelist.com/blog/research/71713/darkhotels-attacks-in-2015/>

<https://securelist.com/blog/research/69114/animals-in-the-apt-farm/>

Security Affairs:

<http://securityaffairs.co/wordpress/29290/cyber-crime/security-firms-vs-hidden-lynx.html>

<http://securityaffairs.co/wordpress/36195/cyber-crime/cozyduke-russian-apt-group.html>

<http://securityaffairs.co/wordpress/34462/intelligence/babar-casper-french-intelligence.html>

<http://securityaffairs.co/wordpress/38204/cyber-crime/dino-malware-animal-farm.html>

Security Week:

<http://www.securityweek.com/cozyduke-apt-responsible-white-house-state-department-attacks-kaspersky>

<http://www.securityweek.com/blue-termite-apt-targets-japanese-organizations>

<http://www.securityweek.com/hacking-team-flash-player-exploit-used-target-japanese-organizations>

<http://www.securityweek.com/apt-group-uses-seaduke-trojan-steal-data-high-value-targets>

<http://www.securityweek.com/onionduke-apt-malware-distributed-malicious-tor-exit-node>

Secureworks:

<http://www.secureworks.com/cyber-threat-intelligence/threats/the-mirage-campaign/>

Sensor Tech Forum:

<http://sensorstechforum.com/bifrose-apt-backdoors-in-the-hands-of-shrouded-crossbow-group/>

Softpedia:

<http://news.softpedia.com/news/apt-group-upgrades-malware-from-the-black-market-into-dangerous-backdoor-497424.shtml>

Spiegel Online International:

<http://www.spiegel.de/international/world/the-nsa-uses-powerful-toolbox-in-effort-to-spy-on-global-networks-a-940969-3.html>

Symantec:

<http://www.symantec.com/connect/blogs/how-elderwood-platform-fueling-2014-zero-day-attacks>

<http://www.symantec.com/connect/blogs/dragonfly-western-energy-companies-under-sabotage-threat>

<http://www.symantec.com/connect/blogs/security-vendors-take-action-against-hidden-lynx-malware>

http://www.symantec.com/security_response/publications/whitepapers.jsp

https://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/

[butterfly-corporate-spies-out-for-financial-gain.pdf](#).

https://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/Dragonfly_Threat_Against_Western_Energy_Suppliers.pdf.

http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/hiddend_lynx.pdf.

<http://www.symantec.com/connect/blogs/forkmeiamfamous-seaduke-latest-weapon-duke-armory>

<http://www.symantec.com/connect/blogs/regin-top-tier-espionage-tool-enables-stealthy-surveillance>

<http://www.symantec.com/connect/blogs/operation-cloudyomega-ichitaro-zero-day-and-ongoing-cyberespionage-campaign-targeting-japan>

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

Tech Times:

<http://www.techtimes.com/articles/22160/20141215/operation-clever-is-bigger-threat-than-previously-thought-fbi-warns-us-businesses.htm>

TechWorm:

<http://www.techworm.net/2014/12/bureau-121.html>

Threatpost:

<https://threatpost.com/new-moker-rat-bypasses-detection/114948/>

Trend Micro TrendLabs Security Intelligence Blog:

<http://blog.trendmicro.com/trendlabs-security-intelligence/sandworm-to-blacken-the-scada-connection/>

<http://blog.trendmicro.co.jp/archives/11944>

<http://blog.trendmicro.com/trendlabs-security-intelligence/new-targeted-attack-group-buys-bifrose-code-works-in-teams/http://www.scmagazine.com/seaduke-and-cloudduke-detected/article/427939/>

http://blog.trendmicro.com/trendlabs-security-intelligence/bkdr_rarstone-new-rat-to-watch-out-for/

The Washington Post:

<https://www.washingtonpost.com/news/the-switch/wp/2015/05/14/the-syrian-electronic-army-just-hacked-the-washington-post-again/>

<https://www.washingtonpost.com/news/the-switch/wp/2014/12/29/a-qa-with-the-hackers-who-say-they-helped-break-in-to-sonys-network/>

<https://www.washingtonpost.com/news/the-switch/wp/2013/08/29/the-nsa-has-its-own-team-of-elite-hackers/>

https://www.washingtonpost.com/world/national-security/researchers-identify-sophisticated-chinese-cyberespionage-group/2014/10/27/de30bc9a-5e00-11e4-8b9e-2ccdac31a031_story.html

https://www.washingtonpost.com/world/report-ties-100-plus-cyber-attacks-on-us-computers-to-chinese-military/2013/02/19/2700228e-7a6a-11e2-9a75-dab0201670da_story.html

https://www.washingtonpost.com/blogs/checkpoint-washington/post/new-stuxnet-like-code-is-discovered/2011/10/19/gIQA8TTHxL_blog.html

https://www.washingtonpost.com/world/national-security/newly-identified-computer-virus-used-for-spying-is-20-times-size-of-stuxnet/2012/05/28/gJQAWa3VxU_story.html

https://www.washingtonpost.com/world/national-security/us-israel-developed-computer-virus-to-slow-iranian-nuclear-efforts-officials-say/2012/06/19/gJQA6xBPoV_story.html

Wired Magazine:

<http://www.wired.com/2014/10/russian-sandworm-hack-isight/>

<http://www.wired.com/2013/11/this-is-how-the-internet-backbone-has-been-turned-into-a-weapon/>

<http://www.wired.com/2014/03/quantum/>

<http://www.wired.com/2014/11/darkhotel-malware/>

