APT trends report Q2 2019

securelist.com/apt-trends-report-q2-2019/91897

Authors

Expert GReAT

For two years, the Global Research and Analysis Team (GReAT) at Kaspersky has been publishing quarterly summaries of advanced persistent threat (APT) activity. The summaries are based on our threat intelligence research and provide a representative snapshot of what we have published and discussed in greater detail in our private APT reports. They aim to highlight the significant events and findings that we feel people should be aware of.

This is our latest installment, focusing on activities that we observed during Q2 2019.

Readers who would like to learn more about our intelligence reports or request more information on a specific report are encouraged to contact 'intelreports@kaspersky.com'.

The most remarkable findings

In April, we published our report on TajMahal, a previously unknown APT framework that has been active for the last five years. This is a highly sophisticated spyware framework that includes backdoors, loaders, orchestrators, C2 communicators, audio recorders, keyloggers, screen and webcam grabbers, documents, and cryptography key stealers; and even its own file indexer for the victim's computer. We discovered up to 80 malicious modules stored in its encrypted Virtual File System – one of the highest numbers of plugins we have ever seen in an APT toolset. The malware features its own indexer, emergency C2s, the ability to steal specific files from external drives when they become available again, and much more. There are two different packages, self-named 'Tokyo' and 'Yokohama' and the targeted computers we found include both packages. We think the attackers used Tokyo as the first stage infection, deploying the fully functional Yokohama package on interesting victims, and then leaving Tokyo in place for backup purposes. So far, our telemetry has revealed just a single victim, a diplomatic body from a country in Central Asia. This begs the question, why go to all that trouble for just one victim? We think there may be other victims that we haven't found yet. This theory is supported by the fact that we couldn't see how one of the files in the VFS was used by the malware, opening the door to the possibility of additional versions of the malware that have yet to be detected.

On May 14, FT reported that a zero-day vulnerability in WhatsApp had been exploited, allowing attackers to eavesdrop on users, read their encrypted chats, turn on the microphone and camera and install spyware that allows even further surveillance, such as browsing through a victim's photos and videos, accessing their contact list and more. In order to exploit the vulnerability, the attacker simply needs to call the victim via WhatsApp. This specially crafted call can trigger a buffer overflow in WhatsApp, allowing an attacker to take control of the application and execute arbitrary code in it. Apparently, the attackers used this method to not only snoop on people's chats and calls but also to exploit previously unknown vulnerabilities in the operating system, which allowed them to install applications on the device. The vulnerability affects WhatsApp for Android prior to 2.19.134, WhatsApp Business for Android prior to 2.19.44, WhatsApp for iOS prior to 2.19.51, WhatsApp Business for iOS prior to 2.19.51, WhatsApp for Windows Phone prior to 2.18.348 and WhatsApp for Tizen prior to 2.18.15. WhatsApp released patches for the vulnerability on May 13. Some have suggested that the spyware may be Pegasus, developed by Israeli company NSO.

Russian-speaking activity

We continue to track the activities of Russian-speaking APT groups. These groups usually show a particular interest in political activities, but apart from a couple of interesting exceptions we failed to detect any remarkable examples during the last quarter.

We did find a potential connection between Hades and a leak at the RANA institute. Hades is possibly connected to the Sofacy threat actor, most notable for being behind Olympic Destroyer, as well as ExPetr and several disinformation campaigns such as the Macron

leaks. Earlier this year, a website named Hidden Reality published leaks allegedly related to an Iranian entity named the RANA institute. This was the third leak in two months that disclosed details of alleged Iranian threat actors and groups. Close analysis of the materials, the infrastructure and the dedicated website used by those behind the leak led us to believe that these leaks might be connected to Hades. This might be part of a disinformation campaign in which Hades helps to raise doubts about the quality of the information leaked in other cases from earlier this year.

Zebrocy continued adding new tools to its arsenal using various kinds of programming languages. We found Zebrocy deploying a compiled Python script, which we call PythocyDbg, within a Southeast Asian foreign affairs organization: this module primarily provides for the stealthy collection of network proxy and communications debug capabilities. In early 2019, Zebrocy shifted its development efforts with the use of Nimrod/Nim, a programming language with syntax resembling both Pascal and Python that can be compiled down to JavaScript or C targets. Both the Nim downloaders that the group mainly uses for spear-phishing, and other Nim backdoor code, are currently being produced by Zebrocy and delivered alongside updated compiled AutoIT scripts, Go, and Delphi modules. The targets of this new Nimcy downloader and backdoor set includes diplomats, defense officials and ministry of foreign affairs staff, from whom they want to steal login credentials, keystrokes, communications, and various files. The group appears to have turned its attention towards the March events involving Pakistan and India, and unrelated diplomatic and military officials, while maintaining ongoing access to local and remote networks belonging to Central Asian governments.

We also recently observed some interesting new artifacts that we relate to Turla with varying degrees of confidence.

In April 2019, we observed a new COMpfun-related targeted campaign using new malware. The Kaspersky Attribution Engine shows strong code similarities between the new family and the old COMpfun. Moreover, the original COMpfun is used as a downloader in one of the spreading mechanisms. We called the newly identified modules Reductor after a .pdb path left in some samples. We believe the malware was developed by the same COMPfun authors that, internally, we tentatively associated with the Turla APT, based on victimology. Besides the typical RAT functions (upload, download, execute files), Reductor's authors put a lot of effort into manipulating installed digital root certificates and marking outbound TLS traffic with unique host-related identifiers. The malware adds embedded root certificates to the target host and allows operators to add additional ones remotely through a named pipe. The solution used by Reductor's developers to mark TLS traffic is the most ingenious part. The authors don't touch the network packets at all; instead they analyze Firefox source and Chrome binary code to patch the corresponding system pseudo-random number generation (PRNG) functions in the process's memory. Browsers use PRNG to generate the "client

random" sequence during the very beginning of the TLS handshake. Reductor adds the victims' unique encrypted hardware- and software-based identifiers to this "client random" field.

Additionally we identified a new backdoor that we attribute with medium confidence to Turla. The backdoor, named Tunnus, is .NET-based malware with the ability to run commands or perform file actions on an infected system and send the results to its C2. So far, the C2 infrastructure has been built using compromised sites with vulnerable WordPress installations. According to our telemetry, Tunnus's activity started last March and was still active at the time of writing.

ESET has also reported PowerShell scripts being used by Turla to provide direct, in-memory loading and execution of malware. This is not the first time this threat actor has used PowerShell in this way, but the group has improved these scripts and is now using them to load a wide range of custom malware from its traditional arsenal. The payloads delivered via the PowerShell scripts – the RPC backdoor and PowerStallion – are highly customized.

Symantec has also been tracking targeted attacks in a series of campaigns against governments and international organizations across the globe over the past 18 months. The attacks have featured a rapidly evolving toolset and, in one notable instance, the apparent hijacking of infrastructure belonging to OilRig. They have uncovered evidence that the Waterbug APT group (aka Turla, Snake, Uroburos, Venomous Bear and KRYPTON) has conducted a hostile takeover of an attack platform belonging to OilRig (aka Crambus). Researchers at Symantec suspect that Turla used the hijacked network to attack a Middle Eastern government that OilRig had already penetrated. This is not the first time that we have seen this type of activity. Clearly, operations of this kind make the job of attribution more difficult.

The international community continues to focus on the activity of Russian-speaking threat actors. Over the last 18 months, the UK has shared information on attacks attributed to Russian hackers with 16 NATO allies, including attacks on critical national infrastructure and attempts to compromise central government networks. In his former capacity as UK foreign secretary, Jeremy Hunt, recently urged nations to band together to create a deterrent for state-sponsored hackers. As part of this push, the UK and its intelligence partners have been slowly moving towards a 'name and shame' approach when dealing with cyberattacks. The use of the 'court of public opinion' in response to cyberattacks is a trend that we highlighted in our predictions for 2019. To help this new strategy the EU recently passed new laws that will make it possible for EU member states to impose economic sanctions against foreign hackers.

Researchers at the Microstep Intelligence Bureau have published a report on targeted attacks on the Ukrainian government that they attribute to the Gamaredon threat actor. Recently, the group launched attacks on a number of state organizations in Ukraine using

Pterodo, malware used exclusively by this group. Since February, the attackers have deployed a large number of dynamic domain names and newly registered domain names believed to be used to launch targeted attacks against elections in Ukraine.

Chinese-speaking activity

We found an active campaign by a Chinese APT group we call SixLittleMonkeys that uses a new version of the Microcin Trojan and a RAT that we call HawkEye as a last stager. The campaign mainly targets government bodies in Central Asia. For persistence, the operators use .DLL search order hijacking. This consists of using a custom decryptor with a system library name (e.g., version.dll or api-ms-win-core-fibers-l1-1-1.dll) in directories, along with the legitimate applications that load these libraries into memory. Among other legitimate applications, the threat actor uses the Google updater, GoogleCrashHandler.exe, for .DLL hijacking. Custom encryptors protect the next stagers from detection on disk and from automated analysis, using the same encryption keys in different samples. For secure TLS communication with its C2, the malware uses the Secure Channel (Schannel) Windows security package.

ESET discovered that the attackers behind the Plead malware have been distributing it using compromised routers and man-in-the-middle (MITM) attacks in April. Researchers have detected this activity in Taiwan, where the Plead malware has been most actively deployed. Trend Micro has previously reported the use of this malware in targeted attacks by the BlackTech group, primarily focused on cyber-espionage in Asia. ESET telemetry has revealed multiple attempts to deploy it.

LuckyMouse activity detected by Palo Alto involved the attackers installing web shells on SharePoint servers to compromise government organizations in the Middle East, probably exploiting CVE-2019-0604, a remote code execution vulnerability used to compromise the server and eventually install a web shell. The actors uploaded a variety of tools that they used to perform additional activities on the compromised network, such as dumping credentials, as well as locating and pivoting to additional systems on the network. Of particular note is the group's use of tools to identify systems vulnerable to CVE-2017-0144, the vulnerability exploited by EternalBlue and used in the 2017 WannaCry attacks. This activity appears to be related to campaigns exploiting CVE-2019-0604 mentioned in recent security alerts from the Saudi Arabian National Cyber Security Center and the Canadian Center for Cyber Security.

Last year, a number of Chinese hackers allegedly linked to the Chinese government were indicted in the US. In May, the US Department of Justice indicted a Chinese national for a series of computer intrusions, including the 2015 data breach of health insurance company Anthem which affected more than 78 million people.

Middle East

The last three months have been very interesting for this region, especially considering the multiple leaks of alleged Iranian activity that were published within just a few weeks of each other. Even more interesting is the possibility that one of the leaks may have been part of a disinformation campaign carried out with the help of the Sofacy/Hades actor.

In March, someone going by the handle Dookhtegan or Lab_dookhtegan started posting messages on Twitter using the hashtag #apt34. Several files were shared via Telegram that supposedly belonged to the OilRig threat actor. They included logins and passwords of several alleged hacking victims, tools, infrastructure details potentially related to different intrusions, the résumés of the alleged attackers and a list of web shells – apparently relating to the period 2014-18.

The targeting and TTPs are consistent with this threat actor, but it was impossible to confirm the origins of the tools included in the dump. Assuming that the data in the dump is accurate, it also shows the global reach of the OilRig group, which has generally been thought to operate primarily in the Middle East.

On April 22, an entity going by the alias Bl4ck_B0X created a Telegram channel named GreenLeakers. The purpose of the channel, as stated by its creator, was to publish information about the members of the MuddyWater APT group, “along with information about their mother and spouse and etc.”, for free. In addition to this free information, the Bl4ck_B0X actor(s) also hinted that “highly confidential” information related to MuddyWater would be put up for sale.

On April 27, three screenshots were posted in the GreenLeakers Telegram channel, containing alleged screenshots from a MuddyWater C2 server. On May 1, the channel was closed to the public and its status changed to private. This was before Bl4ck_B0X had the chance to publish the promised information on the MuddyWater group. The reason for the closure is still unclear.

Finally, a website named Hidden Reality published leaks allegedly related to an entity named the Iranian RANA institute. It was the third leak in two months disclosing details of alleged Iranian threat actors and groups.

Interestingly, this leak differed from the others by employing a website that allows anyone to browse the leaked documents. It also relies on Telegram and Twitter profiles to post messages related to Iranian CNO capabilities. The Hidden Reality website contains internal documents, chat messages and other data related to the RANA institute’s CNO (Computer Network Operations) capabilities, as well as information about victims. Previous leaks were focused more on tools, source code and individual actor profiles.

Close analysis of the materials, the infrastructure and the dedicated website used by the leakers, provided clues that led us to believe Sofacy/Hades may be connected to these leaks.

There was also other Muddywater activity unrelated to the leak, as well as discoveries linked to previous activity by the group, such as ClearSky's discovery of two domains hacked by MuddyWater at the end of 2018 to host the code of its POWERSTATS malware.

In April, Cisco Talos published its analysis of the BlackWater campaign, related to MuddyWater activity. The campaign shows how the attackers added three distinct steps to their operations, allowing them to bypass certain security controls to evade detection: an obfuscated VBA script to establish persistence as a registry key, a PowerShell stager and FruityC2 agent script, and an open source framework on GitHub to further enumerate the host machine. This could allow the attackers to monitor web logs and determine whether someone outside the campaign has made a request to their server in an attempt to investigate the activity. Once the enumeration commands run, the agent communicates with a different C2 and sends back data in the URL field. Trend Micro also reported MuddyWater's use of a new multi-stage PowerShell-based backdoor called POWERSTATS v3.

We published a private report about four Android malware families and their use of false flag techniques, among other things. One of the campaigns sent spear-phishing emails to a university in Jordan and the Turkish government, using compromised legitimate accounts to trick victims into installing malware.

Regarding other groups, we discovered new activity related to ZooPark, a cyber-espionage threat actor that has focused mainly on stealing data from Android devices. Our new findings include new malicious samples and additional infrastructure that has been deployed since 2016. This also led to us discovering Windows malware implants deployed by the same threat actor. The additional indicators we found shed some light on the targets of past campaigns, including Iranian Kurds – mainly political dissidents and activists.

Recorded Future published an analysis of the infrastructure built by APT33 (aka Elfin) to target Saudi organizations. Following the exposure of a wide range of their infrastructure and operations by Symantec in March, researchers at Recorded Future discovered that APT33, or closely aligned actors, reacted by either parking or reassigning some of their domain infrastructure. The fact that this activity was executed just a day or so after the report went live suggests the Iranian threat actors are acutely aware of the media coverage of their activities and are resourceful enough to be able to react in a quick manner. Since then, the attackers have continued to use a large swath of operational infrastructure, well in excess of 1,200 domains, with many observed communicating with 19 different commodity RAT

implants. An interesting development appears to be their increased preference for njRAT, with over half of the observed suspected APT33 infrastructure being linked to njRAT deployment.

On a more political level, there were several news stories covering Iranian activity.

A group connected to the Iranian Revolutionary Guard has been blamed for a wave of cyber-attacks against UK national infrastructure, including the Post Office, local government networks, private companies and banks. Personal data of thousands of employees were stolen. It is believed that the same group was also responsible for the attack on the UK parliamentary network in 2017. The UK NCSC (National Cyber Security Centre) is providing assistance to affected organizations.

Microsoft recently obtained a court order in the US to seize control of 99 websites used by the Iranian hacking group APT35 (aka Phosphorus and Charming Kitten). The threat actor used spoofed websites, including those of Microsoft and Yahoo, to conduct cyberattacks against businesses, government agencies, journalists and activists who focus on Iran. The sinkholing of these sites will force the group to recreate part of its infrastructure.

The US Cybersecurity and Infrastructure Security Agency (CISA) has reported an increase in cyberattacks by Iranian actors or proxies, targeting US industries and government agencies using destructive wiper tools. The statement was posted on Twitter by CISA director, Chris Krebs.

Southeast Asia and Korean Peninsula

This quarter we detected a lot of Korean-related activity. However, for the rest of the Southeast Asian region there has not been that much activity, especially when compared to earlier periods.

Early in Q2, we identified an interesting Lazarus attack targeting a mobile gaming company in South Korea that we believe was aimed at stealing application source code. It's clear that Lazarus keeps updating its tools very quickly. Meanwhile, BlueNoroff, the Lazarus sub-group that typically targets financial institutions, targeted a bank in Central Asia and a crypto-currency business in China.

In a recent campaign, we observed ScarCruft using a multi-stage binary to infect several victims and ultimately install a final payload known as ROKRAT – a cloud service-based backdoor. ScarCruft is a highly skilled APT group, historically using geo-political issues to target the Korean Peninsula. We found several victims worldwide identified as companies and individuals with ties to North Korea, as well as a diplomatic agency. Interestingly, we observed that ScarCruft continues to adopt publicly available exploit code in its tools. We also found an interesting overlap in a Russian-based victim targeted both by ScarCruft and DarkHotel – not the first time that we have seen such an overlap.

ESET recently analyzed a new Mac OS sample from the OceanLotus group that had been uploaded to VirusTotal. This backdoor shares its features with a previous Mac OS variant, but the structure has changed and detection is now much harder. Researchers were unable to find the dropper associated with this sample, so they could not identify the initial compromise vector.

The US Department of Homeland Security (DHS) has reported Trojan variants, identified as HOPLIGHT, being used by the North Korean government. The report includes an analysis of nine malicious executable files. Seven of them are proxy applications that mask traffic between the malware and the remote operators. The proxies have the ability to generate fake TLS handshake sessions using valid public SSL certificates, disguising network connections with remote malicious actors. One file contains a public SSL certificate and the payload of the file appears to be encoded with a password or key. The remaining file does not contain any of the public SSL certificates, but attempts outbound connections and drops four files: the dropped files primarily contain IP addresses and SSL certificates.

In June, we came across an unusual set of samples used to target diplomatic, government and military organizations in countries in South and Southeast Asia. The threat actor behind the campaign, which we believe to be the PLATINUM APT group, uses an elaborate, previously unseen, steganographic technique to conceal communication. A couple of years ago, we predicted that more and more APT and malware developers would use steganography, and this campaign provides proof: the actors used two interesting steganography techniques in this APT. It's also interesting that the attackers decided to implement the utilities they need as one huge set – an example of the framework-based architecture that is becoming more and more popular.

Other interesting discoveries

On May 14, Microsoft released fixes for a critical Remote Code Execution vulnerability (CVE-2019-0708) in Remote Desktop Services (formerly known as Terminal Services) that affects some older versions of Windows: Windows 7, Windows Server 2008 R2, Windows Server 2008 and some unsupported versions of Windows – including Windows 2003 and Windows XP. Details on how to mitigate this vulnerability are available in our private report 'Analysis and detection guidance for CVE-2019-0708'. The Remote Desktop Protocol (RDP) itself is not vulnerable. This vulnerability is pre-authentication and requires no user interaction. In other words, the vulnerability is 'wormable', meaning that any future malware that exploits this vulnerability could propagate from vulnerable computer to vulnerable computer in a similar way that WannaCry spread. Microsoft has not observed exploitation of this vulnerability, but believes it is highly likely that malicious actors will write an exploit for it.

Early in June, researchers at Malwarebytes Labs observed a number of compromises on Amazon CloudFront, a Content Delivery Network (CDN), where hosted JavaScript libraries were tampered with and injected with web skimmers. Although attacks that involve CDNs

usually affect a large number of web properties at once via their supply chain, this isn't always the case. Some websites either use Amazon's cloud infrastructure to host their own libraries or link to code developed specifically for them and hosted on a custom AWS S3 bucket. Without properly validating externally loaded content, these sites are exposing their users to various threats, including some that pilfer credit card data. After analyzing these breaches, researchers found that they are a continuation of a campaign from Magecart threat actors attempting to cast a wide net around many different CDNs. CDNs are widely used because they provide great benefits to website owners, including optimizing load times and cost, as well as helping with all sorts of data analytics. The sites they identified had nothing in common other than the fact they were all using their own custom CDN to load various libraries. In effect, the only resulting victims of a compromise on their CDN repository would be themselves.

Dragos has reported that XENOTIME, the APT group behind the TRISIS (aka TRITON and HatMan) attack on a Saudi Arabian petro-chemical facility in 2017, has expanded its focus beyond the oil and gas industries. Researchers have recently seen the group probing the networks of electric utility organizations in the US and elsewhere – perhaps as a precursor to a dangerous attack on critical infrastructure that could potentially cause physical damage or loss of life. Dragos first noticed the shift in targeting in late 2018; and the attacks have continued into 2019.

We recently reported on the latest versions of FinSpy for Android and iOS, developed in mid-2018. This surveillance software is sold to government and law enforcement organizations all over the world, who use it to collect a variety of private user information on various platforms. WikiLeaks first discovered the implants for desktop devices in 2011 and mobile implants were discovered in 2012. Since then Kaspersky has continuously monitored the development of this malware and the emergence of new versions in the wild. Mobile implants for iOS and Android have almost the same functionality. They are capable of collecting personal information such as contacts, messages, emails, calendars, GPS location, photos, files in memory, phone call recordings and data from the most popular messengers. The Android implant includes functionality to gain root privileges on an unrooted device by abusing known vulnerabilities. It would seem that the iOS solution doesn't provide infection exploits for its customers: the product seems to be fine-tuned to clean traces of publicly available jailbreaking tools. This might imply that physical access to the victim's device is required in cases where devices are not already jailbroken. The latest version includes multiple features that we haven't observed before. During our recent research, we detected up-to-date versions of these implants in the wild in almost 20 countries, but the size of the customer base would suggest that the real number of victims may be much higher.

Final thoughts

APT activity in the Middle East has been particularly interesting this quarter, not least because of the leaks related to alleged Iranian activity. This is especially interesting because one of those leaks might have been part of a disinformation campaign carried out with the help of the Sofacy/Hades threat actor.

In contrast to earlier periods, when Southeast Asia was the most active region for APTs, the activities we detected this quarter were mainly Korean-related. For the rest of the region, it was a much quieter quarter.

Across all regions, geo-politics remains the principal driver of APT activity.

It is also clear from our FinSpy research that there is a high demand for 'commercial' malware from governments and law enforcement agencies.

One of the most noteworthy aspects of the APT threat landscape we reported this quarter was our discovery of TajMahal, a previously unknown and technically sophisticated APT framework that has been in development for at least five years. This full-blown spying framework includes up to 80 malicious modules stored in its encrypted Virtual File System – one of the highest numbers of plugins we've ever seen for an APT toolset.

As always, we would note that our reports are the product of our visibility into the threat landscape. However, it needs to be borne in mind that, while we strive to continually improve, there is always the possibility that other sophisticated attacks may fly under our radar.

- APT
- Chinese-speaking cybercrime
- Instant Messengers
- Russian-speaking cybercrime
- Sofacy
- Targeted attacks
- Turla
- Vulnerabilities and exploits
- Zero-day vulnerabilities

Authors

GReAT

APT trends report Q2 2019

Your email address will not be published. Required fields are marked *