APT annual review 2021

By GReAT

Published: 2021-11-30 · Archived: 2026-04-05 18:36:05 UTC

In the Global Research and Analysis Team at Kaspersky, we track the ongoing activities of more than 900 advanced threat actors and activity clusters; you can find our quarterly overviews here, here and here. For this annual review, we have tried to focus on what we consider to be the most interesting trends and developments of the last 12 months. This is based on our visibility in the threat landscape and it's important to note that no single vendor has complete visibility into the activities of all threat actors.

Private sector vendors play a significant role in the threat landscape

Possibly the biggest story of 2021, an investigation by the Guardian and 16 other media organizations, published in July, suggested that over 30,000 human rights activists, journalists and lawyers across the world may have been targeted using Pegasus. The report, called Pegasus Project, alleged that the software uses a variety of exploits, including several iOS zero-click zero-days. Based on forensic analysis of numerous mobile devices, Amnesty International's Security Lab found that the software was repeatedly used in an abusive manner for surveillance. The list of targeted individuals includes 14 world leaders. Later that month, representatives from the Israeli government visited the offices of NSO as part of an investigation into the claims. And in October, India's Supreme Court commissioned a technical committee to investigate whether the government had used Pegasus to spy on its citizens. In November, Apple announced that it was taking legal action against NSO Group for developing software that targets its users with "malicious malware and spyware".

Detecting infection traces from Pegasus and other advanced mobile malware is very tricky, and complicated by the security features of modern OSs such as iOS and Android. Based on our observations, this is further complicated by the deployment of non-persistent malware, which leaves almost no traces after reboot. Since many forensics frameworks require a device jailbreak, this results in the malware being removed from memory during the reboot. Currently, several methods can be used for detection of Pegasus and other mobile malware. MVT (Mobile Verification Toolkit) from Amnesty International is free, open source and allows technologists and investigators to inspect mobile phones for signs of infection. MVT is further boosted by a list of IoCs (indicators of compromise) collected from high profile cases and made available by Amnesty International.

Supply-chain attacks

There have been a number of high-profile supply-chain attacks in the last 12 months. Last December, it was reported that SolarWinds, a well-known IT managed services provider, had fallen victim to a sophisticated supply-chain attack. The company's Orion IT, a solution for monitoring and managing customers' IT infrastructure, was compromised. This resulted in the deployment of a custom backdoor named Sunburst on the networks of more than 18,000 SolarWinds customers, including many large corporations and government bodies, in North America, Europe, the Middle East and Asia.

Not all supply-chain attacks have been that sophisticated. Early this year, an APT group that we track as BountyGlad compromised a certificate authority in Mongolia and replaced the digital certificate management client software with a malicious downloader. Related infrastructure was identified and used in multiple other incidents: this included server-side attacks on WebSphere and WebLogic services in Hong Kong, and Trojanized Flash Player installers on the client side.

While investigating the artefacts of a supply-chain attack on an Asian government Certification Authority's website, we discovered a Trojanized package that dates back to June 2020. Unravelling that thread, we identified a number of post-compromise tools in the form of plugins that were deployed using PhantomNet malware, which were in turn delivered using the aforementioned Trojanized packages. Our analysis of these plugins revealed similarities with the previously analyzed CoughingDown malware.

In April 2021, Codecov, provider of code coverage solutions, publicly disclosed that its Bash Uploader script had been compromised and was distributed to users between January 31 and April 1. The Bash Uploader script is publicly distributed by Codecov and aims to gather information on the user's execution environments, collect code coverage reports and send the results to the Codecov infrastructure. This script compromise effectively constitutes a supply-chain attack.

Earlier this year we discovered Lazarus group campaigns using an updated DeathNote cluster. Our investigation revealed indications that point to Lazarus building supply-chain attack capabilities. In one case we found that the infection chain stemmed from legitimate South Korean security software executing a malicious payload; and in the second case, the target was a company developing asset monitoring solutions, an atypical victim for Lazarus. As part of the infection chain, Lazarus used a downloader named Racket, which they signed using a stolen certificate. The actor compromised vulnerable web servers and uploaded several scripts to filter and control the malicious implants on successfully breached victim machines.

A previously unknown, suspected Chinese-speaking APT modified a fingerprint scanner software installer package on a distribution server in a country in East Asia. The APT modified a configuration file and added a DLL with a .NET version of a PlugX injector to the installer package. Employees of the central government in this country are required to use this biometric package to track attendance. We refer to this supply-chain incident and this particular PlugX variant as SmudgeX. The Trojanized installer appears to have been staged on the distribution server from March through June.

Exploiting vulnerabilities

On March 2, Microsoft reported a new APT actor named HAFNIUM, exploiting four zero-days in Exchange Server in what they called "limited and targeted attacks". At the time, Microsoft claimed that, in addition to HAFNIUM, several other actors were exploiting them as well. In parallel, Volexity also reported the same Exchange zero-days being in use in early 2021. According to Volexity's telemetry, some of the exploits in use are shared across several actors, besides the one Microsoft designates as HAFNIUM. Kaspersky telemetry revealed a spike in exploitation attempts for these vulnerabilities following the public disclosure and patch from Microsoft. During the first week of March, we identified approximately 1,400 unique servers that had been targeted, in which one or more of these vulnerabilities were used to obtain initial access. According to our telemetry, most exploitation attempts were observed for servers in Europe and the United States. Some of the servers were

targeted multiple times by what appear to be different threat actors (based on the command execution patterns), suggesting the exploits had become available to multiple groups.

We also discovered a campaign active since mid-March targeting governmental entities in Europe and Asia using the same Exchange zero-day exploits. This campaign made use of a previously unknown malware family that we dubbed FourteenHi. Further investigation revealed traces of activity involving variants of this malware dating back a year. We also found some overlaps in these sets of activities with HAFNIUM in terms of infrastructure and TTPs as well as the use of ShadowPad malware during the same timeframe.

On January 25, the Google Threat Analysis Group (TAG) announced a state-sponsored threat actor had targeted security researchers. According to Google TAG's blog, this actor used highly sophisticated social engineering, approached security researchers through social media, and delivered a compromised Visual Studio project file or lured them to their blog where a Chrome exploit was waiting for them. On March 31, Google TAG released an update on this activity showing another wave of fake social media profiles and a company the actor set up mid-March. We confirmed that several infrastructures on the blog overlapped with our previously published reporting about Lazarus group's ThreatNeedle cluster. Moreover, the malware mentioned by Google matched ThreatNeedle – malware that we have been tracking since 2018. While investigating associated information, a fellow external researcher confirmed that he was also compromised by this attack, sharing information for us to investigate. We discovered additional C2 servers after decrypting configuration data from the compromised host. The servers were still in use during our investigation, and we were able to get additional data related to the attack. We assess that the published infrastructure was used not only to target security researchers but also in other Lazarus attacks. We found a relatively large number of hosts communicating with the C2s at the time of our research.

Expanding our research on the exploit targeting CVE-2021-1732, originally discovered by DBAPPSecurity Threat Intelligence Center and used by the Bitter APT group, we discovered another possible zero-day exploit used in the Asia-Pacific (APAC) region. Further analysis revealed that this escalation of privilege (EoP) exploit had potentially been used in the wild since at least November 2020. We reported this new exploit to Microsoft in February. After confirmation that we were indeed dealing with a new zero-day, it received the designation CVE-2021-28310. Various marks and artifacts left in the exploit meant that we were highly confident that CVE-2021-1732 and CVE-2021-28310 were created by the same exploit developer that we track as Moses. Moses appears to be an exploit developer who makes exploits available to several threat actors, based on other past exploits and the actors observed using them. To date, we have confirmed that at least two known threat actors have utilized exploits originally developed by Moses: Bitter APT and Dark Hotel. Based on similar marks and artifacts, as well as privately obtained information from third parties, we believe at least six vulnerabilities observed in the wild in the last two years have originated from Moses. While the EoP exploit was discovered in the wild, we weren't able to directly tie its usage to any known threat actor that we currently track. The EoP exploit was probably chained together with other browser exploits to escape sandboxes and obtain system level privileges for further access. Unfortunately, we weren't able to capture a full exploit chain, so we don't know if the exploit is used with another browser zero-day, or coupled with exploits taking advantage of known, patched vulnerabilities.

On April 14-15, Kaspersky technologies detected a wave of highly targeted attacks against multiple companies. Closer analysis revealed that all these attacks exploited a chain of Google Chrome and Microsoft Windows zero-day exploits. While we were not able to retrieve the exploit used for remote code execution (RCE) in the Chrome web browser, we were able to find and analyze an EoP exploit used to escape the sandbox and obtain system

privileges. The EoP exploit was fine-tuned to work against the latest and most prominent builds of Windows 10 (17763 – RS5, 18362 – 19H1, 18363 – 19H2, 19041 – 20H1, 19042 – 20H2) and exploited two distinct vulnerabilities in the Microsoft Windows OS kernel. We reported these vulnerabilities to Microsoft and they assigned CVE-2021-31955 to the information disclosure vulnerability and CVE-2021-31956 to the EoP vulnerability. Both vulnerabilities were patched on June 8 as a part of the June Patch Tuesday. The exploit-chain attempts to install malware in the system through a dropper. The malware starts as a system service and loads the payload, a remote shell-style backdoor that in turn connects to the C2 to get commands. Because we couldn't find any connections or overlaps with a known actor, we named this cluster of activity PuzzleMaker.

Finally, late this year, we detected a wave of attacks using an elevation of privilege exploit affecting server variants of the Windows operating system. Upon closer analysis, it turned out to be a zero-day use-after-free vulnerability in Win32k.sys that we reported to Microsoft and was consequently fixed as CVE-2021-40449. We analyzed the associated malware, dubbed the associated cluster MysterySnail and found infrastructure overlaps that link it to the IronHusky APT.

Firmware vulnerabilities

In September, we provided an overview of the FinSpy PC implant, covering not only the Windows version, but also Linux and macOS versions. FinSpy is an infamous, commercial surveillance toolset that is used for "legal surveillance" purposes. Historically, several NGOs have repeatedly reported it being used against journalists, political dissidents and human rights activists. Historically, its Windows implant was represented by a single-stage spyware installer; and this version was detected and researched several times up to 2018. Since then, we have observed a decreasing detection rate for FinSpy for Windows. While the nature of this anomaly remained unknown, we began detecting some suspicious installer packages backdoored with Metasploit stagers. We were unable to attribute these packages to any threat actor until the middle of 2019 when we found a host that served these installers among FinSpy Mobile implants for Android. Over the course of our investigation, we found out that the backdoored installers are nothing more than first stage implants that are used to download and deploy further payloads before the actual FinSpy Trojan. Apart from the Trojanized installers, we also observed infections involving usage of a UEFI or MBR bootkit. While the MBR infection has been known since at least 2014, details on the UEFI bootkit were publicly revealed for the first time in our report.

Towards the end of Q3, we identified a previously unknown payload with advanced capabilities, delivered using two infection chains to various government organizations and telecoms companies in the Middle East. The payload makes use of a Windows kernel-mode rootkit to facilitate some of its activities and is capable of being persistently deployed through an MBR or a UEFI bootkit. Interestingly enough, some of the components observed in this attack have been formerly staged in memory by Slingshot agent on multiple occasions, whereby Slingshot is a post-exploitation framework that we covered in several cases in the past (not to be confused with the Slingshot APT). It is mainly known for being a proprietary commercial penetration testing toolkit officially designed for red team engagements. However, it's not the first time that attackers appear to have taken advantage of it. One of our previous reports from 2019 covering FruityArmor's activity showed that the threat group used the framework to target organizations across multiple industries in the Middle East, possibly by leveraging an unknown exploit in a messenger app as an infection vector. In a recent private intelligence report, we provided a drill-down analysis of the newly discovered malicious toolkit that we observed in tandem with Slingshot and how it was leveraged in

clusters of activity in the wild. Most notably, we outlined some of the advanced features that are evident in the malware as well as its utilization in a particular long-standing activity against a high-profile diplomatic target in the Middle East.
