securelist.com /apt-trends-report-q2-2024/113275/

APT trends report Q2 2024

GReAT ⋮⋮ 8/13/2024

For over six years now, Kaspersky's Global Research and Analysis Team (GReAT) has been sharing quarterly updates on advanced persistent threats (APTs). These summaries draw on our threat intelligence research, offering a representative overview of what we've published and discussed in more detail in our private APT reports. They're designed to highlight the key events and findings that we think people should know about.

In this latest installment, we focus on activities that we observed during Q2 2024.

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

Most notable findings

In March, a backdoor was discovered in XZ, a compression utility integrated into many popular distributions of Linux. The backdoored library liblzma is used by the OpenSSH server process sshd. OpenSSH is patched to use systemd features on a number of systemd-based distributions, including Ubuntu, Debian and RedHat/Fedora Linux, and therefore depends on this library (Arch Linux and Gentoo are not affected). The code was inserted in February and March 2024, mostly by Jia Cheong Tan – probably a fictitious identity. The likely goal of the attack was to introduce exclusive remote code execution capabilities into the sshd process by targeting the XZ build process, and then to push the backdoored code to major Linux distributions as a part of a large-scale supply-chain attack. The attackers used social engineering to gain prolonged access to the source/development environment, and extended that access by faking human interactions in plain sight to build credibility for introducing the malicious code.

There are two levels at which the backdoor in the liblzma library was introduced. The source code of the build infrastructure that generated the final packages was tweaked slightly (by adding an extra file, build-to-host.m4) to extract the next stage script that was hidden in a test case file (bad-3-corrupt_lzma2.xz). The script then extracted a malicious binary component from another test case file (good-large_compressed.lzma) that was linked with the legitimate library during the compilation process to be shipped to Linux repositories. Some of the big vendors ended up shipping the malicious component in beta and experimental builds without realizing it. The compromise of XZ Utils was given the identifier CVE-2024-3094 and a maximum severity score of 10.

The attackers' initial goal was to successfully hook one of the functions related to RSA key manipulation. In our analysis of the hook process, we focused on the backdoor's behavior inside OpenSSH, specifically OpenSSH portable version 9.7p1 (the most recent version). Our analysis revealed a number of interesting details about the functionality of the backdoor.

- The attacker set an anti-replay feature to make sure the backdoor communication couldn't be captured or hijacked.
- The author hid the public key for backdoor decryption in the x86 code using a custom steganography technique.
- The backdoor hooks the logging function to hide its logs of unauthorized connections to the SSH server.
- The backdoor hooks the password authentication function, which allows the attacker to use any username/password to log in to the infected server without any further checks. It also does the same for public key authentication.
- The backdoor has remote code execution capabilities, which means the attacker can run any system command on the infected server.

You can read our analysis here, here and here.

Chinese-speaking activity

In an earlier report on ToddyCat, we described various tools used to collect and exfiltrate files of interest to this APT threat actor. One of these tools was PcExter, which was initially only used to exfiltrate data previously collected with the help of other tools, such as FileScan. However, we recently found a new version, PcExter 2.0, which has been completely redesigned and rewritten in .NET to be able to collect the data itself, as well as use an improved file search mechanism. We found several versions of this tool, together with a set of special loaders.

In 2021, we published a private report describing the technical details of QSC, a framework that was discovered while investigating an attack on the telecoms industry in South Asia. While our research did not reveal how the framework was deployed, or the threat group behind it, we continued to monitor our telemetry for further detections of the QSC framework. In October 2023, we saw multiple detections of QSC framework files in the West Asia region targeting an ISP. Our investigation revealed that the target machines had already been infected with Quarian Backdoor version 3 (aka Turian) since 2022, and the same attackers used this access to deploy the QSC framework starting from October 10, 2023. In addition to the QSC framework, the attackers also deployed a new backdoor written in Golang, which we named "GoClient": we saw the first deployment of this GoClient backdoor on October 17, 2023. After analyzing all the artifacts from this campaign, we assess with medium confidence that the CloudComputating threat actor is behind the deployment of the QSC framework and the GoClient backdoor.

Early in 2023, the activities of GOFFEE were discovered when this threat actor used a modified version of a monitored malicious IIS module called Owowa. Since then, GOFFEE has stopped using Owowa, as well as a PowerShell RCE implant VisualTaskel; however, it has continued to conduct intrusions leveraging PowerTaskel, the threat actor's previous HTA-based infection chain, and has added a new loader, disguised as a legitimate document and distributed via email, to its arsenal.

We recently found a new remote access tool (RAT) with a low detection rate called SalmonQT that was uploaded from a computer in China to a public multi-scanner platform. What caught our attention was that the sample used GitHub's REST API to accept instructions and upload data, thereby acting as a C2 (command and control) server. At first glance, it appeared that the path to the GitHub repository had been

deleted, but on closer inspection, the repository was set to private and the REST API could only be accessed using the correct token. The C2 server was active from early January 2024 up to the completion of our report at the end of June this year. We attribute this newly discovered RAT with low confidence to the threat actor CNC. CNC (aka APT-C-48) is highly focused on Chinese entities.

Middle East

Gaza Cybergang has been active since at least 2012, targeting the Middle East and North Africa. When we first started tracking the group, its attacks were relatively basic in nature, often relying on publicly available malware families such as QuasarRAT. Nevertheless, the group exhibited a particular TTP that we can still see today – going after only a few targets per campaign. At the start of this year we detected several cases involving Gaza Cybergang in which the threat actor adjusted its TTPs slightly. Instead of using tabcal.exe as a vehicle to sideload its initial access downloader IronWind, the group switched to setup_wm.exe, another legitimate Windows Media Utility file. The lures were also changed to a more generic theme, rather than focusing on a specific geopolitical situation.

Southeast Asia and Korean Peninsula

We discovered Mysterious Elephant in 2023 while investigating attacks using a set of malware families previously associated with other known threat actors, such as SideWinder and Confucius. As we analyzed the infrastructure, we realized that the attacks were not in fact delivered by any of the previously known actors, but by a new threat actor that we dubbed Mysterious Elephant. The threat actor has remained active since then and has launched several attacks since our initial report. We have discovered a wealth of new malware families developed and used by Mysterious Elephant in its recent attacks, as well as recently created infrastructure and updated tools – mostly backdoors and loaders to minimize detection in the early stages of attacks. In our report, we describe the latest attacks delivered by this threat actor and analyze the newly discovered malware samples and associated infrastructure.

Hacktivism

With the start of the Russian-Ukrainian conflict in February 2022, hundreds of different hacktivist groups have emerged on both sides. One such group is -=Twelve=-. This group announced itself in the information sphere by claiming to have hacked various government and industrial enterprises of the Russian Federation. Some of the targets were published on the group's official channel on its own platform, while others remained in the shadows. While there are several reports on the internet about the Twelve group from various CTI (Cyber Threat Intelligence) vendors that attempt to describe the group's activities, we have not seen any that detail the tools and techniques used in the attacks. Our report on Twelve provides a detailed overview of the TTPs used by the group, as well as the connections to its infrastructure.

In February, the Institute of Geography and Statistics of Albania (INSTAT) was attacked. The attack was the work of Homeland Justice – a self-described hacktivist group, but suspected of being a state sponsored group – that has been relentlessly attacking Albanian targets, particularly in the government sector, for over three years. The attackers were able to obtain more than 100TB of data, as well as disrupt the official websites and email services of organizations and wipe database servers and backups.

One of the main reasons for the attacks is the presence of a Mujahedeen-e-Khalq (MEK) refugee camp on Albanian territory: Homeland Justice considers this group to be a terrorist organization and believes that specific sectors of the Albanian government and certain companies provide them with support and funding. The threat actor conducts ongoing cyber operations aimed at conveying its anti-MEK political message. They are attempting to garner support among the Albanian people for the government to abandon the MEK – their actions are framed within what are known as psychological operations (PsyOps) campaigns.

We have analyzed the group's campaign history, which spans almost three years of cyberattacks aimed at exerting long-term pressure on the Albanian government and populace. In our report, we cover its main campaigns, ranging from sophisticated operations involving collaboration with allied groups with the same aims, to opportunistic attacks. We also describe the main techniques employed by the group, which range from exploiting internet-facing servers for initial access, lateral movement activities, expanding the attack surface, to using custom wiping malware and ransomware in the final disruptive phase of the cyber operations. Additionally, we examine the group's persuasion mechanisms, such as amplifying messaging through social networks and news media, sharing stolen data to gain notoriety and advocate for change, and the continual threat of future attacks to induce a state of permanent vigilance among its targets.

Other interesting discoveries

We discovered a new modular malware framework, which we dubbed "Aniseed Vodka", on a system in East Africa: the system was infected in 2018. The framework consists of a main module, a JSON-formatted configuration file, and a set of plug-ins. The framework is highly configurable, allowing its operator both to specify operating parameters for plug-ins and to schedule plug-in tasks (such as screen capture, webcam capture, and data exfiltration) at specific intervals. The framework employs anti-detection and anti-forensics techniques, enabling it to operate covertly. It uses non-traditional communication channels to evade network detection, using Google Chat as a C2 channel, Gmail to send alerts and Google Drive as an exfiltration channel. The framework we presented in our report is, as far as we know, not publicly known. We have not been able to tie this framework to an existing threat actor.

Our previous report on DinodasRAT showed a wealth of overlaps in features between the Linux backdoor version and its Windows counterpart, as well as additional Linux-specific functionalities such as persistence through systemd or SystemV. In recent months, we were able to collect more relevant samples, giving us a deeper insight into the Linux variant. There are indications that it has been used in campaigns dating back to 2021. Previously identified as XDealer, an ongoing APT campaign using the Windows version of this threat was disclosed by ESET and named "Operation Jacana". DinodasRAT was also used in a recent APT campaign, which included both its Windows and Linux versions, as described by Trend Micro. In our latest report on the Linux variant of DinodasRAT, we focus on the network communication with the C2 and the operations performed by the malware on the infected machine, beyond establishing persistence and awaiting C2 commands.

In May 2024, we discovered a new APT targeting Russian government entities. The CloudSorcerer malware is a sophisticated cyber-espionage tool used for stealth monitoring, data collection and exfiltration via Microsoft, Yandex and Dropbox cloud infrastructures. The malware uses cloud resources for its C2 servers, accessing them through APIs using authentication tokens. Additionally, CloudSorcerer uses GitHub as its initial C2 server. CloudSorcerer's modus operandi of is reminiscent of the

CloudWizard APT, which we reported on in 2023. However, the malware code is completely different. We believe that CloudSorcerer is a new threat actor that has adopted a similar method of interacting with public cloud services.

In April, we discovered a previously unknown campaign targeting organizations in Russia, including the government sector, using the Telemos backdoor. The malware is delivered via spear-phishing emails as a ZIP file containing one of two types of dropper – a PE64 executable with an .SCR extension or a Windows Script File with a .WSF extension. These drop and execute a PowerShell-based script with backdoor functionality. We found several malicious samples associated with these attacks and were able to restore the original source code. The main purpose of this threat is espionage – collecting data from browsers such as login credentials, cookies and browsing history, as well as collecting files of interest from available drives on the affected system. The operation cannot be tied to a known threat actor at this point.

Final thoughts

While some threat actors' TTPs remain the same, such as a heavy reliance on social engineering to gain entry to a target organization or compromising an individual's device, others have updated their toolsets and broadened the scope of their activities. Our regular quarterly reports are designed to highlight the most significant developments related to APT groups.

Here are the key trends we saw in Q2 2024:

- The key highlight this quarter was the backdooring of the XZ compression utility integrated into many popular Linux distributions – in particular, the use of social engineering to gain persistent access to the development environment.
- This quarter we saw APT campaigns focused on Europe, the Americas, Asia, the Middle East and Africa targeting a range of sectors including government, military, telecoms and judicial systems.
- The purpose of most APT activities is cyber-espionage, although some campaigns are driven by financial gain.
- Hacktivist attacks have also been a feature of the threat landscape this quarter. Not all of these attacks are focused on areas of open conflict, as illustrated by the attacks on entities in Albania by the Homeland Justice group.

As always, we would like to point out that our reports are the product of our insight into the threat landscape. However, it is important to remember that while we strive for continuous improvement, there is always the possibility that there are other sophisticated attacks that may go unnoticed.

Disclaimer: when referring to APT groups as Russian-speaking, Chinese-speaking or other-language-speaking, we refer to various artifacts used by the groups (such as malware debugging strings, comments found in scripts, etc.) containing words in these languages, based on the information that we obtained directly or that is otherwise publicly known and widely reported. The use of certain languages does not necessarily indicate a specific geographic relation, but rather points to the languages that the developers behind these APT artifacts use.