APT trends report Q3 2023

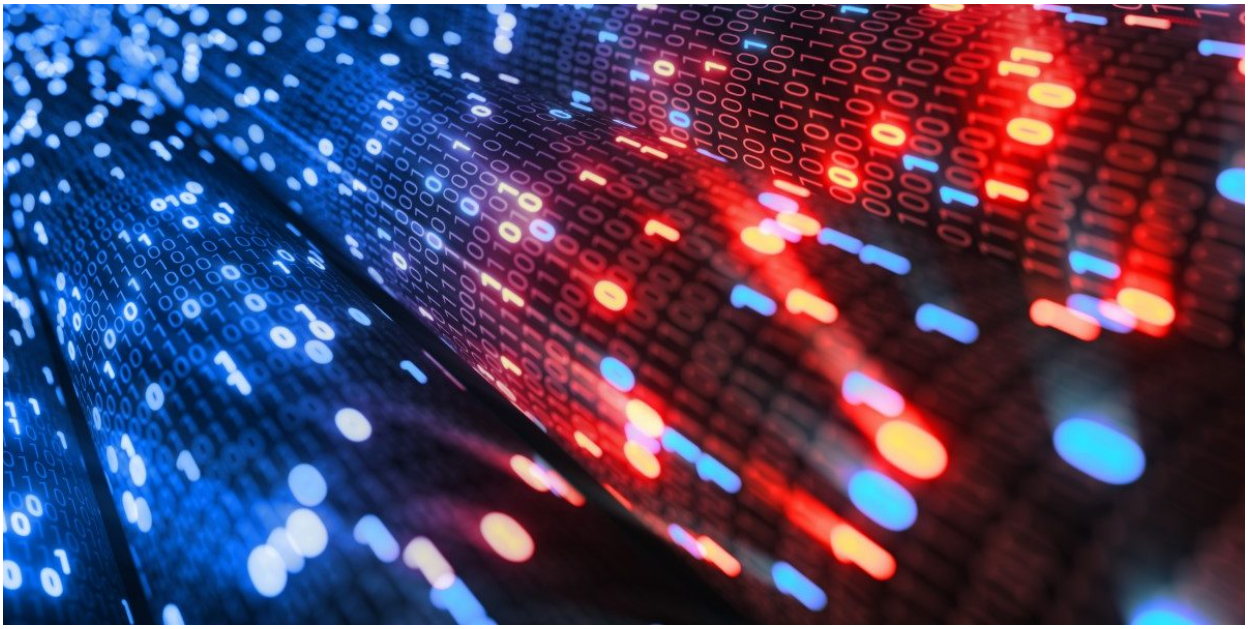

Authors

- Expert GReAT

For more than six years, the Global Research and Analysis Team (GReAT) at Kaspersky has been publishing quarterly summaries of advanced persistent threat (APT) activity. These summaries are based on our threat intelligence research; and they provide a representative snapshot of what we have published and discussed in greater detail in our private APT reports. They are designed 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 Q3 2023.

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 early 2023, we discovered an ongoing attack targeting government entities in the APAC region by compromising a specific type of a secure USB drive, which provides hardware encryption. Such secure USB drives are used by the government organisations of the country to securely store and transfer data physically between computer systems. The USB drive contains a protected partition which can only be accessed via custom software bundled on an unencrypted part of the USB and a passphrase known to the user.

Further investigation revealed a long-running campaign consisting of various malicious modules, used to execute commands and collect files and information from compromised machines and pass them on to further machines using the same or other secure USB drives as a carrier. They are also capable of executing other malicious files on the infected systems.

The attack comprises sophisticated tools and techniques, including virtualization-based software obfuscation for malware components, low-level communication with the USB drive using direct SCSI commands, self-replication through connected secure USB drives to propagate to other air-gapped systems and injection of code into a legitimate access management program on the USB drive which acts as a loader for the malware on a new machine.

The attacks were extremely targeted and had a quite limited number of victims. Our investigation revealed a high level of sophistication in the malicious tools used in the deployment of the attacks. We believe these attacks have been carried out by a highly-skilled and resourceful threat actor interested in espionage activities in sensitive and protected government networks. It is therefore very important to build a deep understanding of the TTPs of this threat actor and to watch out for future attacks.

BlindEagle has targeted both government entities and individuals in South America. While the threat actor's primary focus is espionage, it has also shown an interest in stealing financial data. One notable characteristic of BlindEagle is the threat actor's propensity to cycle through various open-source remote access Trojans (RATs) such as AsyncRAT, Lime-RAT and BitRAT, and use them as the final payload to achieve its objectives. A recent development in BlindEagle's modus operandi involves a shift in its choice of final payload. Previously, our reports highlighted the group's transition from the Quasar RAT to njRAT, showcasing its dynamic approach. In its most recent wave of attacks, BlindEagle has once again adapted, embracing yet another open-source RAT, Agent Tesla. This strategic shift signals its intent to intensify its surveillance capabilities and expand its range of targets.

BlindEagle subsequently added the Remcos RAT to its toolset in attacks on government entities, private companies and individuals in Colombia. The initial attack vector was a phishing email disguised as an email from a government entity or service. This email contained a link leading to a password-protected archive hosted on Google Drive, which represented the first stage of the infection – a .NET binary that was obfuscated and trying to pass itself off as an OpenVPN binary, when in fact it was a malware loader. The final payload was a loaded malware implant, the Remcos RAT – the fourth type of remote administration tool adopted by this threat actor within a few months of operation.

In a stunning display of grit, given the group's lack of resources and development talent, this threat actor cycles through open-source RATs to supply the implants it uses. However, we continue to observe these types of groups and their successful intrusion sets in the less-reported parts of the world. This ongoing trend is one that, over time, will transform these lesser actors into better resourced, more capable offensive adversaries.

Russian-speaking activity

Beginning in late 2022, a new and unknown APT group launched attacks against multiple entities in Russia. The group targets its victims by sending spear-phishing emails with Microsoft Office documents attached. This initiates a multi-level infection scheme leading to the installation of a new Trojan, which is

primarily designed to exfiltrate files from the victim's machine and gain control by executing arbitrary commands. The first observed wave of attacks commenced in October 2022, followed by another in April 2023, with the aim of compromising dozens of victims, including government entities, military contractors, universities and hospitals. We are currently unable to link the activity to any known group, so we dubbed both waves of attacks "BadRory". The name is a reference to text in one of the malicious files taken from the classic novel "The Adventures of Roderick Random", which tells the life story of the protagonist, Roderick "Rory" Random.

Chinese-speaking activity

One of the more unique components of HoneyMyte's malware set includes cookie stealers for web-based productivity and email services from a wide range of browsers. The stolen cookies can be used later to remotely access victims' email accounts. These implants, which haven't been seen in any other APTs, were recently observed in Vietnam, starting in June of this year. The new cookie-stealer variant adds support for "Coc coc", a localized browser built specifically for the Vietnamese market. These are new variants of the same hacking tools that we first reported in August 2019, and later detailed again in June 2021.

In late 2021, we first documented "Owowa", a malicious IIS backdoor module that had been deployed primarily in Asia since 2020. At that time, we showed that Owowa may have been developed by a Chinese-speaking individual. We continued to monitor the possible deployment of Owowa and discovered that, starting in May 2022, an updated variant was uniquely used against targets in Russia. We were later able to link the deployment of Owowa with an email-based intrusion chain that mimicked known CloudAtlas activity. We refer to the malicious campaign of attacks that leverage both Owowa and the email-based intrusion chain against common targets in Russia as "GOFFEE". The campaign is still ongoing.

In January, we reported on an in-memory implant named "TargetPlug": operations using this implant focused on the gaming sector, primarily in South Korea. In April of this year, we saw a fresh surge of attacks involving related variants of TargetPlug. Notably, these attacks expanded their scope to include entities within the software and entertainment domains located in Spain and Mexico. These newer iterations of TargetPlug exhibit various changes in their loader components. The malware architects removed a distinctive string that previously served as a telltale compromise marker within the loader, and introduced a string hashing algorithm derived from omniORB, an open-source Common Object Request Broker Architecture (CORBA) implementation.

Spanish-speaking activity

See above, "The most remarkable findings".

Middle East

Dark Caracal, a highly skilled threat group operating with nation-state level capabilities, has been conducting cyber-espionage campaigns since at least 2012. The group's campaigns target governments, military entities, utilities, financial institutions, manufacturing companies and defense contractors worldwide. This threat actor is known for stealing valuable data, including intellectual property and

personally identifiable information, impacting thousands of victims. This group has been referenced as a “cyber mercenary threat group” due to the variety of targets and the apparent targeting of multiple governments in its campaigns. Since 2021, this group’s activity has been reported to be focused on Spanish-speaking countries, mainly in Latin America. While tracking Dark Caracal’s activity, we discovered an ongoing campaign targeting public and private sector entities in multiple Spanish-speaking countries.

StrongyPity (aka PROMETHIUM) is a Turkish-speaking threat actor known to have been active since at least 2012. We first reported it in 2016 following a series of attacks against users in Italy and Belgium, where it used watering-hole attacks to deliver malicious versions of WinRAR and TrueCrypt. In 2020, we discovered a new version of the StrongPity implant that we named StrongPity4, which appeared to be more advanced than previous variants and was detected with a small number of victims in Egypt, Syria and Turkey. We have continued to monitor the attacker’s activities over the past few years. Another Chinese vendor reported a similar attack against entities in China and partially described a similar implant and various modules deployed by the attacker on specific victim machines. These new modules were not publicly available, but we were able to detect them on targets in the Middle East and North Africa. In a recent report, we described these additional modules, which are used to search for relevant files on victim machines that can then be exfiltrated, as well as log keystrokes and take screenshots. We also discovered new variants of loaders used to start the main StrongPity implants, and observed an expansion of attacker activity to new countries: Algeria, Lebanon, Armenia and Iran.

In a recent investigation, we discovered malicious PowerShell scripts that indicate the attackers behind these scripts are active inside multiple high-profile private and public organizations in the Levant region. Our telemetry shows that this threat actor’s latest endeavor – Operation BlackCrescent – is focused on entities in specific countries, namely: Palestine, Syria, Lebanon and Iran. The artefacts we collected in this investigation indicate that the operation may have been active for several years, dating back to at least 2020 or earlier. While we couldn’t tie the majority of the victims to specific industries or the activities to a specific threat actor, one of the victims is a telecom equipment supplier in the Middle East. We assess with low confidence that this operation has ties to the Lyceum threat group. Lyceum is believed to be a Farsi-speaking threat group that has been active since 2018 and may be behind this novel PowerShell tool set. This threat group’s tradecraft, in turn, has some notable similarities to the prolific OilRig and the infamous DNSpionage.

BellaCiao is .NET-based malware publicly associated with the threat actor Charming Kitten (aka Newsbeef and APT35). In May, we published a detailed research report investigating hitherto unreported activity by this threat actor. More recently, however, we have encountered new iterations of this malware, exposing additional command-and-control (C2) addresses and some minor changes in its methodology. By consolidating and analyzing all discovered samples we were able to shed some light on the malware’s development and evolution through its PDB paths.

We observed MuddyWater leveraging Ligolo, an open-source reverse tunneling project hosted on GitHub, in its malicious activities. Microsoft unveiled a sample called “vpnui.exe” that was identified as a Ligolo tool. However, there is not much information about the specifics of how Ligolo works, so we tried to shed some light on it, focusing on its customized variants and the threat actor’s efforts to conceal its operations. Our investigation uncovered two distinct files that can definitively be categorized as customized Ligolo tools. These variants go beyond Ligolo’s standard functionality and attempt to emulate

VPN solutions from Cisco and Palo Alto. This emulation goes beyond appearance, as these customized Ligolo tools incorporate metadata reminiscent of authentic VPN services, effectively masking their true nature. The primary objective of this customization is to evade detection and maintain a covert presence within targeted systems, making it difficult for security tools and security staff to identify them as malicious.

Southeast Asia and Korean Peninsula

We discovered an ongoing Lazarus campaign targeting the defense industry and nuclear engineers. Lazarus uses Trojanized apps, especially backdoored VNC apps, to access enterprise systems. The threat actor tricks job seekers on social media into opening malicious apps for fake job interviews. To avoid detection by behavior-based security solutions, this backdoored application operates discreetly, only activating when the user selects a server from the drop-down menu of the Trojanized VNC client. The application launches additional payloads into memory and retrieves further malicious code. Through our telemetry, we have observed the installation of an additional payload on the victim's system. Once the compromised VNC client is executed by the victim, it triggers the creation of further malware known as "LPEClient", which has been previously used by the Lazarus group on multiple occasions. It also employs sophisticated C2 communication methods and disables behavior monitoring by unhooking user-mode syscalls. We also saw the use of an updated version of COPPERHEDGE as an additional backdoor, with a complex infection chain. In addition, we observed the presence of a malware variant specifically designed to transfer targeted files to a remote server. The purpose of this particular malware is to exfiltrate specific files chosen by the Lazarus group and send them to their designated remote server. Through our telemetry, we have confirmed numerous instances of compromised companies. The majority of these affected companies are directly involved in defense manufacturing, including radar systems, unmanned aerial vehicles (UAVs), military vehicles, ships, weaponry and maritime companies. Furthermore, our observations led us to identify the username associated with the initial infection. Subsequently, through conversations with the victim, we verified that this indeed corresponded to the real name of the victim. This individual is a nuclear engineer based in Hungary who received the malicious file after being in contact with a suspicious account via Telegram and WhatsApp.

During our assessment of Origami Elephant, we encountered an initial infection sample that resembled the infection chains of Vtyrei and RTY, leading to its classification under Origami Elephant. However, the first and final payloads differed from their typical malware. Upon closer examination, we noticed similarities between RTY's infection chain and the downloader, which we subsequently named CSVtyrei. CSVtyrei bore a striking resemblance to Vtyrei, a downloader used to deploy RTY. Our thorough investigation uncovered a novel .NET=based backdoor called Firebird, with a main loader and at least three plugins. All samples demonstrated robust protection through ConfuserEx, resulting in an exceptionally low detection rate. With a limited impact, we identified only a handful of victims in Pakistan and Afghanistan. Some code within the examples appeared non-functional, hinting at ongoing development efforts.

While tracking the activities of the ScarCruft group, we stumbled upon a novel infection chain. It was initiated by a cleverly crafted macro-embedded Word document that posed as a commercial invoice written in Russian. Once permission is obtained to execute the embedded macro, a highly sophisticated infection chain is executed. The initial infection vector injects a relatively large first-stage shellcode,

including additional packages needed for installation, into a legitimate Windows process. Once the victim's status is verified, the shellcode proceeds to install the Windows Resource Kit, using it to establish a persistent Windows service. In addition, if Python is not present on the victim's system, the shellcode drops the Python package to enable further infection. It then deploys the malicious Python script and the next-stage payload to the victim. This process goes through two more shellcode stages before ultimately executing the Windows executable payload without any disk involvement, deploying the well-known RokRat or BlueLight malware. The final payload uses only cloud services such as OneDrive, GoogleDrive, PCloud and BackBlaze for its C2 operations and data exfiltration. While investigating this case, we discovered that the attacker had compromised their own host for testing purposes, giving us a valuable insight into their testing environment. Moreover, our telemetry strongly suggests that the intended target of this attack was an individual or entity associated with a trading company linked to Russia and North Korea.

Final thoughts

While the TTPs of some threat actors remain consistent over time, relying heavily on social engineering as a means of gaining a foothold in a target organization or compromising an individual's device, others have refreshed their toolsets and expanded the scope of their activities. Our regular quarterly reviews are intended to highlight the most significant developments among APT groups.

Here are the main trends that we've seen in Q3 2023:

- Key highlights this quarter include the attack on government entities in the APAC region by compromising a specific type of a secure USB drive and BlindEagle's Latin American activities – the latter underlining the fact that not all "successful" APT incidents require technical sophistication.
- We've become accustomed to seeing established threat actors enhance their toolsets over time. This quarter, for example, ScarCruft's novel multi-stage infection chain, successive RATs used by BlindEagle and MuddyWater's emulation of VPN applications.
- We have also seen a campaign from a newly discovered threat actor, BadRory.
- APT campaigns continue to be very geographically dispersed. This quarter, we've seen actors focusing their attacks on Europe, South America, the Middle East and various parts of Asia.
- We have seen attacks targeting a wide variety of sectors, including government, military, defence, gaming, software, entertainment, utilities, finance and manufacturing.
- Geo-politics remains a key driver of APT development, and cyber-espionage remains a prime goal of APT campaigns.

As always, we would note that our reports are the product of our visibility into the threat landscape. However, it should 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.

Disclaimer: when referring to APT groups as Russian-speaking, Chinese-speaking or "other-speaking" languages, we refer to various artefacts used by the groups (such as malware debugging strings, comments found in scripts, etc.) containing words in these languages, based on the information 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 artefacts use.