Presented by Trellix ADVANCED RESEARCH CENTER

February 2026

TRELLIX® SECONDSIGHT THREAT HUNTING REPORT

Insights gleaned from Trellix SecondSight, expert threat hunters, and a global network of telemetry and intelligence

TABLE OF CONTENTS

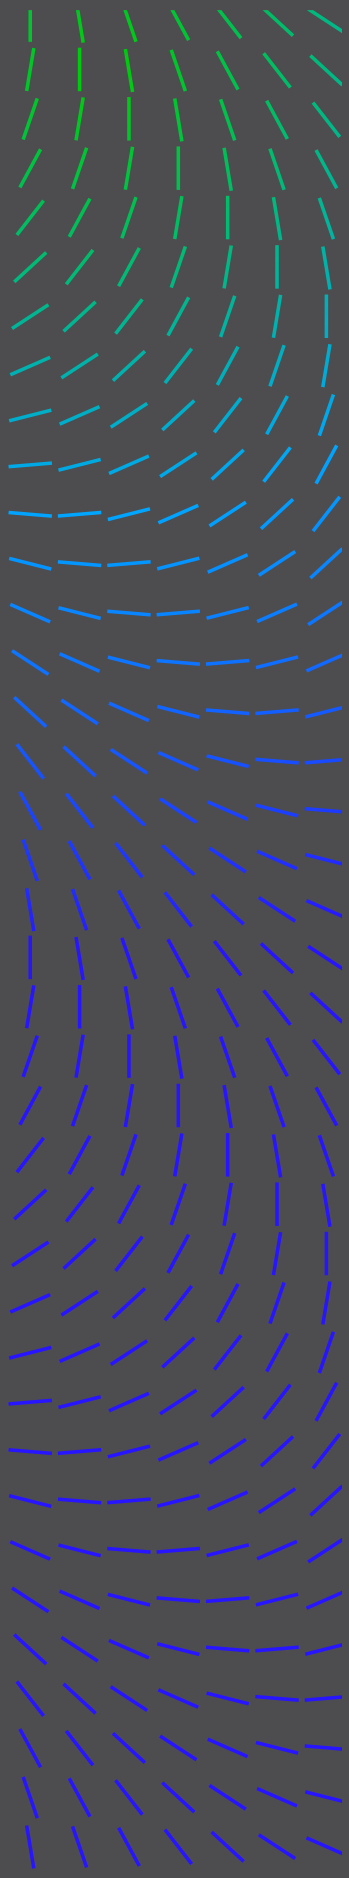

TRELLIX SECONDSIGHT THREAT HUNTING REPORT

Authored by the Trellix Advanced Research Center, this report (1) highlights threat hunting insights, intelligence, and guidance gleaned from multiple sources of critical data, including Trellix SecondSight, on the top five critical campaigns observed in 2025, and (2) develops expert, thorough case studies to inform and enable best practices in defending against these types of campaigns. This edition focuses on data and insights captured primarily between July 1, 2025 and December 31, 2025.

INTRODUCTION

Trellix SecondSight is a threat hunting service combining the telemetry of Trellix products with continual global human oversight, built from a simple belief: the adversary does not wait for perfect telemetry, and neither should we. This report reflects the work of our hunters, who operate alongside customers every day, combining Trellix intelligence with real-world telemetry to find what traditional detection often misses. We go into this level of detail because early signals matter. The first misused domain, the abnormal DLL load, and the unusual OAuth flow are the moments where a breach can still be stopped rather than investigated after the fact.

Our methodology is deliberately practical and adversary-focused. We start with campaigns and tradecraft, not products, mapping attacker behavior to customer environments to expose gaps between what organizations believe they can detect and what is actually happening on their networks. Each investigation in this report shows how Trellix connects weak signals to active operations, validates them through multiple data sources, and translates findings into concrete actions to reduce risk. This approach is driven by a fundamental philosophy of proactive hunting, where we prioritize understanding the attacker's intent over simply waiting for a known signature to match.

While security products wait for alerts to fire, threat hunters press the adversary, provide context, and help customers act faster to mitigate damage, working in tandem to empower security operations teams with proactive defense. The case studies in this report, from targeted espionage operations to OAuth abuse and zero-day exploitation, demonstrate this approach in practice and reinforce why proactive hunting remains one of the most effective defenses against modern threats. Of the thousands of threats observed in SecondSight over the last six months, we've highlighted these five due to their potential impact, techniques used, and importance for our customers.

John Fokker
Vice President, Threat Intelligence Strategy, Trellix

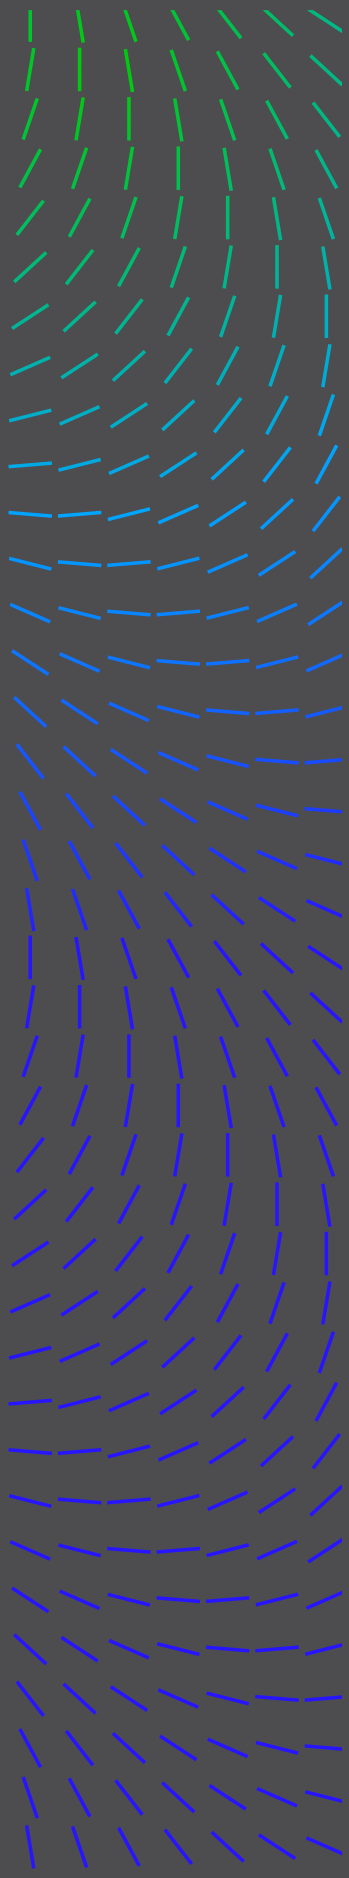

HUNTING METHODOLOGY AND FRAMEWORK

SecondSight hunting is driven by a structured prioritization framework designed to identify meaningful risk early, before attackers reach their objective. Our hunters do not start with alerts alone. We begin by continuously assessing active threat campaigns, adversary objectives, and tradecraft observed across the global landscape, then mapping that intelligence to customer environments to determine where exposure is most likely to exist. This approach allows us to focus our efforts where attacker intent, capability, and opportunity intersect.

Hunting decisions are informed by multiple signals rather than single indicators. Low-confidence alerts, anomalous behaviors, infrastructure reuse, and identity abuse are evaluated together and weighted based on relevance, credibility, and potential impact. Individually, these signals may appear benign or fall below traditional detection thresholds. When combined and placed in an adversarial context, they often reveal early-stage intrusion activity, campaign staging, or abuse of trusted services. Examples from the past six months include the correlation of subtle OAuth misuse with known espionage tradecraft, or the chaining of seemingly unrelated endpoint behaviors that exposed lateral movement before data access occurred.

This framework ensures consistency without sacrificing analyst judgment. Each hunt progresses through clear decision points: validating threat intel relevance, confirming behavioral alignment with known techniques, assessing environmental exposure, and determining whether proactive customer notification is warranted. The result is a disciplined, repeatable approach that prioritizes early disruption over retrospective investigation. This methodology reflects our belief that effective security is not defined by the number of alerts generated, but by the ability to connect intent, behavior, and context in time to prevent meaningful impact.

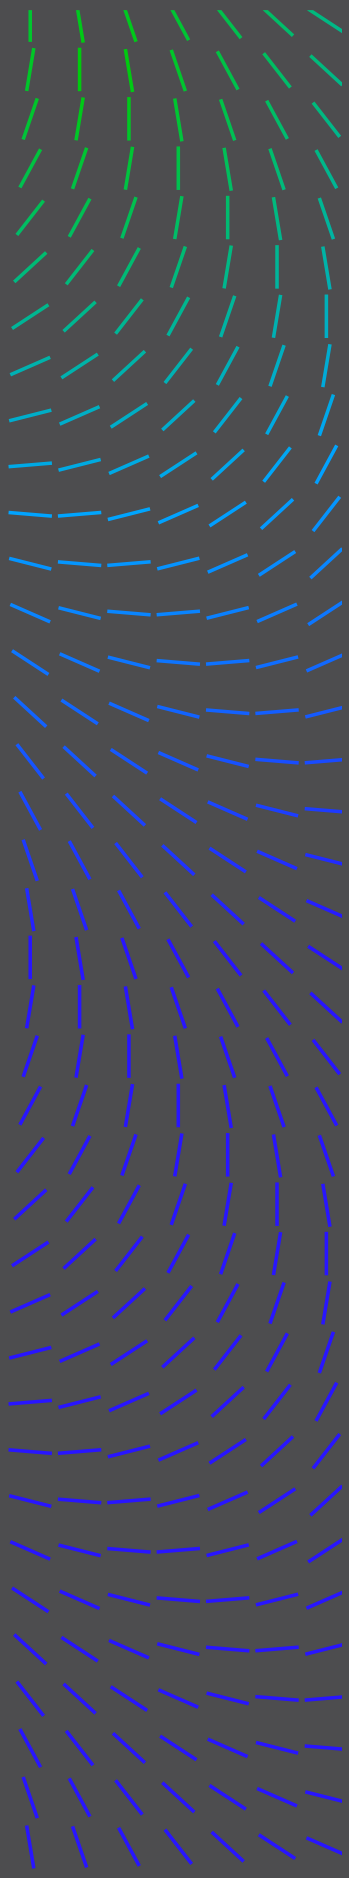

1 PERSISTENT WEB SERVER COMPROMISE: SHAREPOINT CVE-2025-53770 & COBALT STRIKE C2 INFRASTRUCTURE

Incident summary

In July 2025, threat actors exploited a Microsoft SharePoint vulnerability (CVE-2025-53770) targeting financial, healthcare, and life science organizations worldwide to establish a persistent web-based compromise.

On July 18, 2025, threat actors targeted a Canadian health services organization and attempted to execute Base64-encoded PowerShell commands through the w3wp.exe process. While this specific attempt failed to deploy persistent web shells, it signaled an active campaign targeting SharePoint LAYOUTS directories.

Simultaneously, between July 18 and 23, 2025, a U.S.-based organization was hit by a successful multi-stage intrusion. Adversaries utilized the same SharePoint vulnerability to install multiple web shells. The actors performed extensive reconnaissance using whoami and netstat, archived sensitive system data into ZIP files for exfiltration, and established long-term persistence by creating a rogue administrative account named sso-ishaansehgal.

On July 21, 2025, the campaign expanded to a Vietnamese financial institution, where an attacker exploited a public-facing PHP-CGI application on the Online-TTDT host. It resulted in the deployment of a Cobalt Strike beacon (artifact_x64.exe), which initiated persistent command-and-control (C2) communications to IP 84[.]247[.]151[.]254 via port 1233.

These operations, characterized by the abuse of legitimate web server processes and the deployment of advanced post-exploitation agents, highlight a high-risk environment geared toward long-term espionage and potential ransomware delivery.

TTP Progression

Target	Status	Initial access (T1190)	Execution & persistence	Command & control
Canadian Health Services Organization	Exploitation attempt	**SharePoint exploit:** Adversary targeted host via CVE-2025-53770.	**PowerShell execution:** Used w3wp.exe to run $b64$-encoded commands	**Blocked phase:** Managed to execute code but failed to deploy persistent web shells.
U.S. Company	Successful compromise	**SharePoint exploit:** Adversary successfully exploited host.	**Persistence (T1505.003):** Installed web shells docssended.aspx and test.aspx.	**Data staging:** Archived sensitive files into F.zip and a.zip for exfiltration.
Vietnamese financial institution	Full compromise	**PHP-CGI exploit:** Exploited public-facing app on host.	**Admin Persitence** Created rogue local admin account and enumeration commands.	**C2 (T1071):** Deployed Cobalt Strike beacon with persistent check-ins to 84.247.151.254.

Threat hunting process

The immediate, high-severity threat posed by the easily exploitable CVE-2025-53770 vulnerability prompted our team to proactively investigate its exploitation vectors to establish robust detection methods. This analysis revealed a consistent post-exploitation pattern across multiple threat actors: the spawning of PowerShell instances that execute Base64-encoded commands. This key insight guided our threat hunting to anomalous PowerShell executions originating from the w3wp.exe process, which was responsible for running the SharePoint server. By leveraging this specific detection opportunity against customer telemetry, we quickly identified successful exploitation attempts executing PowerShell commands to deploy webshells and other malicious tooling.

Our hunting methodology involved:

- Proactive vulnerability research to define post-exploitation behavioral patterns.
- Process tree anomaly hunting for unusual child processes spawned from w3wp.exe.
- Entropy-based script analysis to identify obfuscated or Base64-encoded commands.

The primary lesson from this campaign is the critical importance of vulnerability exploitation research.

Remediation next steps

Organizations should immediately patch **CVE-2025-53770** and rotate **ASP.NET machine keys** to invalidate potential token theft. Security teams must hunt for and remove the sso-ishaansehgal admin account and web shells such as docssended.aspx or test.aspx. To disrupt active C2, block communication to 84[.]247[.]151[.]254 and monitor for w3wp.exe or php-cgi.exe spawning command shells.

Threat hunting tips

The SharePoint persistent web server compromise offers several critical takeaways for threat-hunting teams.

- **Proactive research:** Use vulnerability research to better understand the initial actors vectors so defenders can anticipate post-exploitation behavior rather than react to alerts.
- **Process-tree anomalies:** Prioritize creating detections for any anomalous processes spawned from w3wp.exe or other public-facing service processes.
- **Untrusted execution:** Never trust processes running in the context of a web server to execute arbitrary system utilities.
- **High-entropy strings:** Specifically hunt for high-entropy strings in command lines, which often indicate encoded or obfuscated commands.

By focusing detection efforts on process tree anomalies, it's possible to turn knowledge about a vulnerability into a highly effective behavioral detection rule, protecting against both known and future exploits that share this common post-exploitation pattern.

2 SIDEWINDER: TARGETED ESPIONAGE AND DLL SIDELOADING ANALYSIS

Incident summary

On September 23, 2025, the India-affiliated threat actor SideWinder launched a targeted spear-phishing campaign against a European government institution with a presence in India. The operation utilized a diplomatic lure, "China Bangladesh Think Tank Forum," to lend an air of legitimacy.

The intrusion began when the attacker, spoofing a Pakistani diplomatic account (asresearch@mofa[.]gov[.]bd[.]pk-mail[.]org), sent an email containing a malicious PDF. This document leveraged social engineering to convince the victim to download a fake "Adobe Reader" application.

Once executed, the installer utilized a ClickOnce application to download legitimate MagTek software, which was then used to sideload a malicious component, DEVOBJ.dll. This multi-stage chain ultimately deployed StealerBot, a malware designed for persistent access and the exfiltration of sensitive data.

Such attacks serve as initial vectors for sophisticated, long-term compromises aimed at stealing confidential data, intellectual property, or classified national security information. Such breaches can severely impact diplomatic relations and national defense. This incident underscores SideWinder's evolving use of DLL sideloading to circumvent traditional defenses in high-value government sectors.

TTP Progression

Phase	Technique (ID)	Observed Adversary Behavior
Initial access	Phishing: Spear phishing Attachment (**T1566.001**)	Sent an email with a weaponized PDF titled "China bangladesh Think Tak Forum-2025.pdf".
Execution	User execution: Malicious file (**T1204.002**)	Victim prompted to click "Install" in the PDF to run a fake Adobe application.
Persistence	Boot or logon autostart execution: Registry run keys / startup folder (**T1547.001**)	Malware adds itself to Registry Run keys to maintain access after reboots.
Defense evasion	Hijack execution flow: DLL (**T1574.001**)	Employs **DLL sideloading** by using legitimate MagTek software to load malicious DEVOBJ.dll.
	System binary proxy execution (**T1218**)	Abuses legitimate binaries to mask the execution of StealerBot.
Discovery	Software discovery: Security software discovery (**T1518.001**)	Scans the system specifically for installed security products to aid in evasion.
Collection	Data from local system (**T1005**)	StealerBot automatically identifies and gathers sensitive data from the local machine.
Command & control	Application layer protocol: Web protocols (**T1071.001**)	Communicates with C2 servers using standard web protocols.
Exfiltration	Exfiltration over C2 channel (**T1041**)	Transfers stolen data out of the network via the established C2 channel.

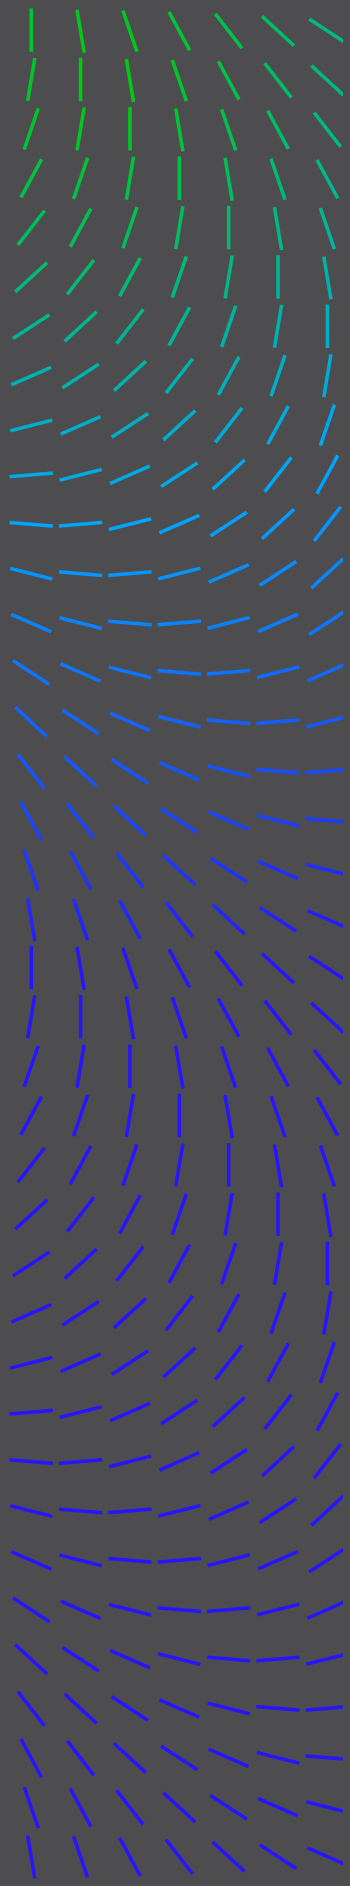

Threat hunting process

The threat hunting process focused on identifying professional email lures and analyzing telemetry that leveraged sophisticated techniques such as geo-fencing.

SideWinder typically targets diplomatic entities, and this campaign was no different. Initial discovery was challenging because the specific TTPs for this particular campaign were previously unknown.

Our investigation focused on identifying emails containing a PDF with an attached link. Key indicators were the professional tone of the email lure and the sender's attempt to mimic a Southeast Asian government authority email account, both of which are common tactics used by SideWinder.

Further analysis revealed a significant finding: the initial infection payload, the Click-Once application, could only be downloaded from a specific country, India in this case, a sophisticated geo-fencing tactic not commonly employed by less-skilled cybercriminals.

Leveraging the initial email, we were able to retrieve additional telemetry data. In some instances, this telemetry included domains and email accounts previously linked to SideWinder. The final analysis of the collected malware samples solidified our attribution efforts.

Our hunting methodology involved:

- Email metadata analysis to identify professional lures masquerading as government authorities.
- Infrastructure pivoting using initial email identifiers to uncover linked domains and accounts.
- Geographic telemetry correlation to identify targeted geo-fencing tactics.

Remediation next steps

To remediate similar spear-phishing campaigns, implement strict email filtering and software restriction policies to block unauthorized applications and the execution of ClickOnce installers. Immediately isolate any systems showing signs of infection and configure firewalls to block connections to the malicious domains filenest[.]live and pk-mail[.]org. Finally, conduct targeted user awareness training on recognizing spear-phishing attempts that masquerade as legitimate software updates.

Threat hunting tips

The SideWinder incident offers several critical takeaways for threat-hunting teams.

- **Geographic telemetry:** Focus on the sophisticated use of geo-fencing tactics in high-risk regions to deliver the initial ClickOnce payload. It demonstrates that adversaries are employing geographic targeting to evade global security monitoring.
- **Rapid infrastructure analysis:** Rather than focusing on stale threat feeds, prioritize immediate analysis and proactive blocking of new domains, given the short URL lifespan of adversary C2 infrastructure.

- **Atypical access vectors:** Pay attention to threat actors' willingness to adopt less-common initial access vectors, such as the ClickOnce application, which often bypasses traditional email attachment filters.
- **DLL anomalies:** Deploy enhanced endpoint detection and response (EDR) rules specifically targeting DLL execution and continuous monitoring for suspicious execution chains even when initiated by trusted binaries. For example, the successful DLL sideloading attack, leveraging a legitimate MagTek application.
- **Installer restrictions:** Implement application allowlisting to restrict the use of rarely used installer types, such as ClickOnce.

A key learning is the actor's willingness to adopt less-common initial access vectors, like ClickOnce, which often bypasses traditional email attachment filters.

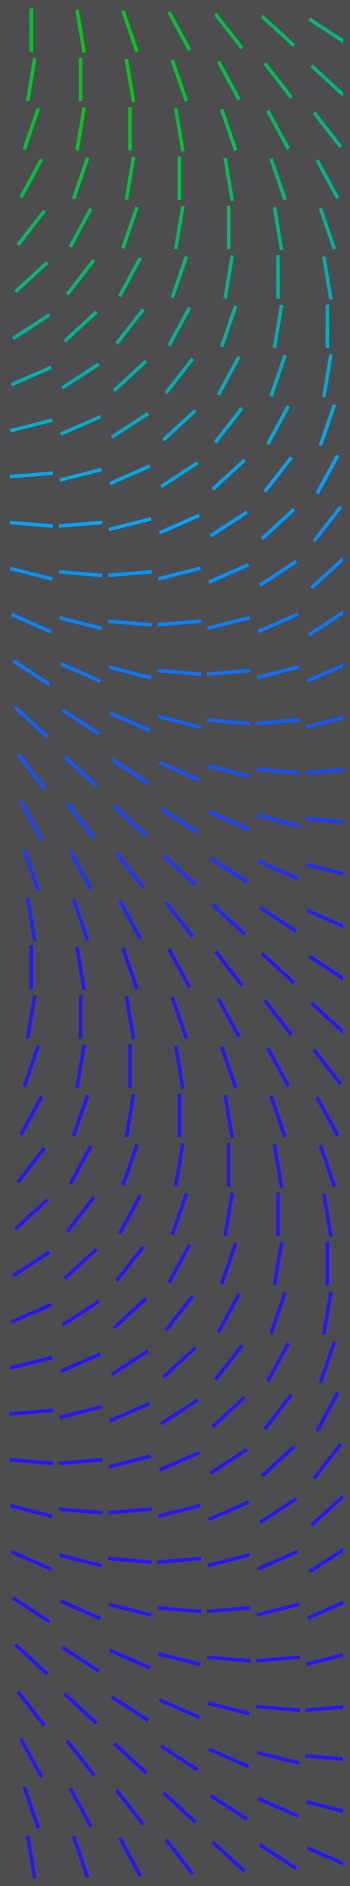

3 MUSTANG PANDA: EUROPEAN DIPLOMATIC ESPIONAGE CAMPAIGN

Incident summary

On September 26, 2025, an official from a European public administration received spear-phishing emails purporting to facilitate fake border crossings, kicking off a complex, multi-stage infection chain. Once the user clicked the malicious link, a ZIP archive was downloaded containing a malicious LNK file that exploited a command-line padding vulnerability to execute obfuscated PowerShell. The script abused the native Windows tar.exe to extract a legitimate Canon printer utility for DLL side-loading, ultimately deploying a PlugX backdoor through a novel EnumSystemGeoID callback hijacking technique.

Shortly after, on September 29 and 30, 2025, there was a similar wave of attacks against another European government institution, only a week after the SideWinder attack. These incidents leveraged lures related to NATO's Joint Analysis, Training, and Education Centre (JATEC) defense procurement workshops and the European Political Community (EPC) Summit in Copenhagen. The emails successfully bypassed security controls to reach recipient inboxes.

Attributed with high confidence to the PRC-nexus actor Mustang Panda, these operations established persistent backdoor access for long-term espionage and data exfiltration. The campaign highlights a strategic focus on European diplomatic personnel and NATO defense cooperation.

TTP Progression

Phase	Technique (ID)	Observed Adversary Behavior
Initial access	Spear-phishing (T1566.002)	Emails targeting EU Government officials using EU/NATO themes.
Delivery	HTML smuggling (T1027.006)	Azure Blob Storage URLs used to smuggle and download malicious ZIP archives.
Exploitation	LNK padding	LNK files exploiting **ZDI-CAN-25373** to hide PowerShell commands beyond display limits.
Extraction	Binary proxy execution (T1218)	PowerShell carves a TAR archive and extracts it using native **tar.exe**.
Execution	DLL side-loading (T1574.002)	Signed Canon utility (**cnmpaui.exe**) used to load malicious **cnmpaui.dll**.
Payload injection	Callback hijacking (T1106)	**EnumSystemGeoID** API misused to execute PlugX shellcode and evade EDR.
Persistence	Registry run keys (T1547.001)	Establishes survival via the **"Canon Printer"** key and masquerading in Public folders.
C2 communication	Web protocols (T1071.001)	PlugX beacons to Cloudflare-proxied domains using HTTPS with randomized parameters.

Threat hunting process

This investigation originated from open-source intelligence published by StrikeReady Labs on October 3, 2025, detailing Mustang Panda operations against Serbian government entities. The OSINT report revealed that the threat actor used Azure Blob Storage to host HTML smuggling payloads. We immediately pivoted to our email telemetry, querying for similar delivery infrastructure using the SQL LIKE pattern: %download%web.core.windows.net%. This uncovered additional targeting against European diplomatic personnel that had successfully bypassed initial security controls.

The breakthrough came when analyzing the HTML payloads on VirusTotal. The fake Cloudflare turnstile pages contained a distinctive JavaScript variable, SFAFAT_URL, combined with XOR obfuscation (key=23)—a unique artifact that enabled confident clustering of related samples. This encoding mechanism (decimal values like 127 99 99 103... translating to URLs) became a reliable signature for hunting additional campaign infrastructure.

```
try {
  turnstile.render('#cf-turnstile', {
    sitekey: TURNSTILE_SITE_KEY,
    callback: function (token) {
      if (SFAFAT_URL) location.assign(SFAFAT_URL);
    },
    'error-callback': function () { console.warn('Turnstile error'); },
    'timeout-callback': function () {
      try { turnstile.reset('#cf-turnstile'); } catch(e) {}
    }
  });
} catch (e) {
  console.error('Failed to render Turnstile:', e);
}
```

Figure 1: The SFAFAT_URL variable is a distinctive artifact within a benign Cloudflare Turnstile CAPTCHA check.

We expanded coverage by hunting for the distinctive TTPs: LNK files exploiting ZDI-CAN-25373, Canon printer DLL side-loading (cnmpaui.dll/cnmplog.dat), and EU/NATO-themed lures targeting diplomatic personnel. This multi-vector approach enabled proactive detection of related attacks before they were publicly disclosed. By exploiting UI character limits to hide malicious code, ZDI-CAN-25373 enables remote code execution; Microsoft's recent fix now prevents this by exposing the full command string.

Our hunting methodology involved:

- OSINT correlation using StrikeReady Labs' reporting on Azure Blob Storage delivery patterns.
- Proactive email telemetry searches using SQL LIKE patterns for specific storage URLs.
- Artifact signature clustering using the unique JavaScript variable SFAFAT_URL and XOR obfuscation.
- Multi-vector pivoting using Trellix IVX and FAUDE to identify additional affected organizations.

Remediation next steps
In response, organizations should conduct immediate hunts for DLL side-loading involving cnmpaui.exe in non-standard locations and to block identified C2 infrastructure like racineupci[.]org and cseconline[.]org. Focus recovery efforts on neutralizing persistence by removing malicious Canon Printer registry keys and implementing detection rules for HTML smuggling and suspicious Azure Blob Storage URLs.

External threat research often reveals infrastructure patterns applicable to your environment.

Threat hunting tips
The Mustang Panda European espionage campaign offers several critical takeaways for threat-hunting teams.

- **OSINT integration:** Leverage OSINT for proactive threat hunting. Often external research reveals infrastructure patterns applicable to your environment.
- **Multi-source telemetry:** Don't limit hunting to a single telemetry source. Expand hunting beyond email logs to include web and endpoint solutions to identify the full scope of a campaign.
- **Geopolitical context:** Context matters for attribution and prioritization. Monitor for themed lures (e.g., EU summits, NATO meetings) that coincide with major international conferences and cross-reference against known APT targeting preferences to improve threat hunting and threat attribution.
- **Unique JavaScript variables:** Focus on unique artifacts, such as the SFAFAT_URL rather than common IOCs, transient IPs or domain names for clustering related campaigns, for technical pivoting.
- **Cloud infrastructure:** Query for specific HTML smuggling patterns and Azure Blob Storage URLs (e.g., %download%web.core.windows.net%). In this instance, it revealed additional affected customers who had accessed the malicious infrastructure.

For attribution and prioritization, context matters. Monitor for thematically relevant lures and cross-reference against known APT targeting preferences.

4 THE "LONG GAME": ANALYSIS OF KIMSUKY'S TRUST-BUILDING SPEAR-PHISHING CAMPAIGN

Incident summary

In late 2025, the North Korean APT group **Kimsuky** executed a sophisticated, multi-stage social engineering campaign targeting the HR department of a South Korean organization. The operation was characterized by a high degree of patience and localized precision, designed to bypass traditional email filters through a "trust-building" methodology.

The intrusion began on September 28 when the attacker, impersonating an applicant named "Park Sung-hwan," sent a benign PDF transcript to an HR representative to build trust. After a strategic 10-day delay, the threat actor sent a follow-up email to a different HR contact, spoofed to appear as an internal HR thread containing a malicious ZIP archive. Inside the archive was a weaponized **LNK file** masquerading as a DOCX document via a double file extension.

Upon execution, the LNK file triggered an obfuscated **PowerShell** chain. The script was designed to conduct registry-based geofence checks and download secondary payloads—such as majority.docx and entiment.ps1—from attacker-controlled **GitHub** repositories. To maintain persistence, the malware attempted to create a scheduled task to execute every 30 minutes.

Trellix Email Security successfully mitigated by quarantining the second-stage payload, preventing the establishment of a persistent backdoor and potential data exfiltration. This incident underscores Kimsuky's continued reliance on long-form social engineering and the abuse of legitimate cloud services to target South Korean commercial interests.

TTP Progression

Phase	Technique (ID)	Observed Adversary Behavior
Initial access	**Spear-phishing attachment (T1566.001)**	Stage 1: Attacker impersonated a job applicant to deliver a clean PDF Stage 2: Sent a follow-up email containing a malicious ZIP.
Execution	**Human-triggered (T1204.002)**	Required the recipient to open the ZIP and double-click the weaponized LNK file.
Execution	**PowerShell (T1059.001)**	LNK file executed cmd.exe to launch heavily obfuscated PowerShell commands.
Persistence	**Scheduled task (T1053.005)**	Attempted to create a task named "Majority Company" to run entiment.ps1 every 30 minutes.
Defense Evasion	**Obfuscation (T1027) & masquerading (T1036.007)**	Used Base64 encoding for scripts and double file extensions (.docx.lnk) to hide the true file type.
Command & control	**Web service: GitHub (T1102)**	Abused GitHub repositories to host and download secondary payloads using hardcoded API tokens.

Threat hunting process

This discovery emerged from our team's research into the DPRK-linked GitHub C2 espionage campaign documented in "The Coordinated Embassy Hunt." While analyzing diplomatic targeting patterns, we pivoted to hunt for similar TTPs targeting South Korean commercial entities, using password-protected ZIP archives and double-file extension lures (.docx.lnk, .pdf.lnk).

The finding came when hunting Trellix Email Security telemetry for emails containing both Korean-language themes and suspicious attachment chains. Our internal methodology for handling password-protected archives flagged this APT campaign. We systematically searched for common Kimsuky indicators: GitHub-hosted PowerShell payloads and Base64-encoded execution chains. The "miffiya1@naver[.]com" sender pattern, cross-referenced with our GitHub C2 research, revealed overlapping TTPs, including hardcoded API tokens, obfuscated PowerShell droppers, and the XenoRAT delivery, consistently observed across multiple North Korean campaigns.

Our hunting methodology involved:

- Thematic lure analysis searching for Korean-language themes and password-protected ZIPs.
- File detection focusing on .lnk files with double extensions like .docx.lnk.
- Cloud service misuse monitoring for unauthorized connections to raw.githubusercontent[.]com.
- Cross-referencing the "miffiyal@naver[.]com" sender pattern with known GitHub C2 research.
- Analysis of GitHub-hosted PowerShell payloads, hardcoded API tokens, and obfuscated execution chains.

Remediation next steps

Block the sender miffiyal@naver[.]com and conduct hunts for .lnk files nested within ZIP archives. Security teams must identify and remove the "Majority Company" scheduled task and delete malicious scripts—specifically, entiment.ps1, addition.ps1, and system_first.ps1—from %TEMP% and %APPDATA% directories. To disrupt active C2, monitor for unauthorized connections to raw.githubusercontent[.]com and block the execution of PowerShell commands originating from masqueraded documents with double extensions like .docx.lnk.

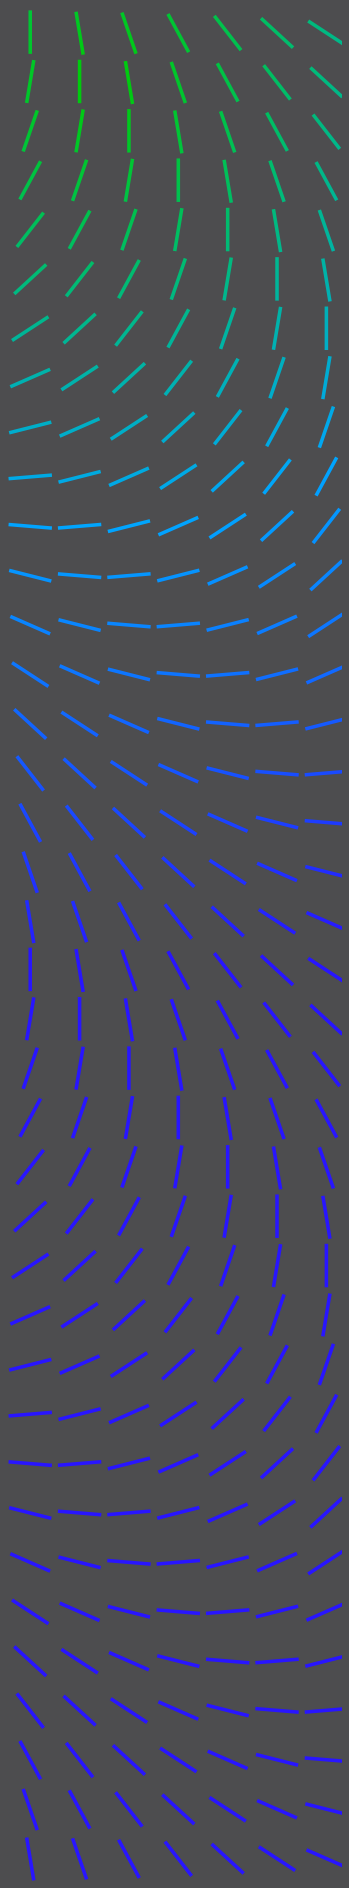

Threat hunting tips

To proactively identify similar activity within your environment, analysts should focus on the following hunting parameters:

- **File masquerading:** Prioritize hunting for .lnk files within password-protected ZIPs, especially with double extensions (.docx.lnk, .pdf.lnk). These consistently evade automated sandboxing due to password protection and user interaction requirements.
- **GitHub C2 detection:** Monitor for raw.githubusercontent.com connections from unexpected processes, particularly PowerShell executing Base64-encoded commands. Implement detections for unusual repository access patterns.
- **Behavioral hunting:** Search for scheduled tasks with Korean company names, PowerShell execution from %TEMP%/%APPDATA%.

Prioritize hunting for .lnk files within password-protected ZIPs, with double extensions as they evade automated sandboxing due to password protection and user interaction requirements.

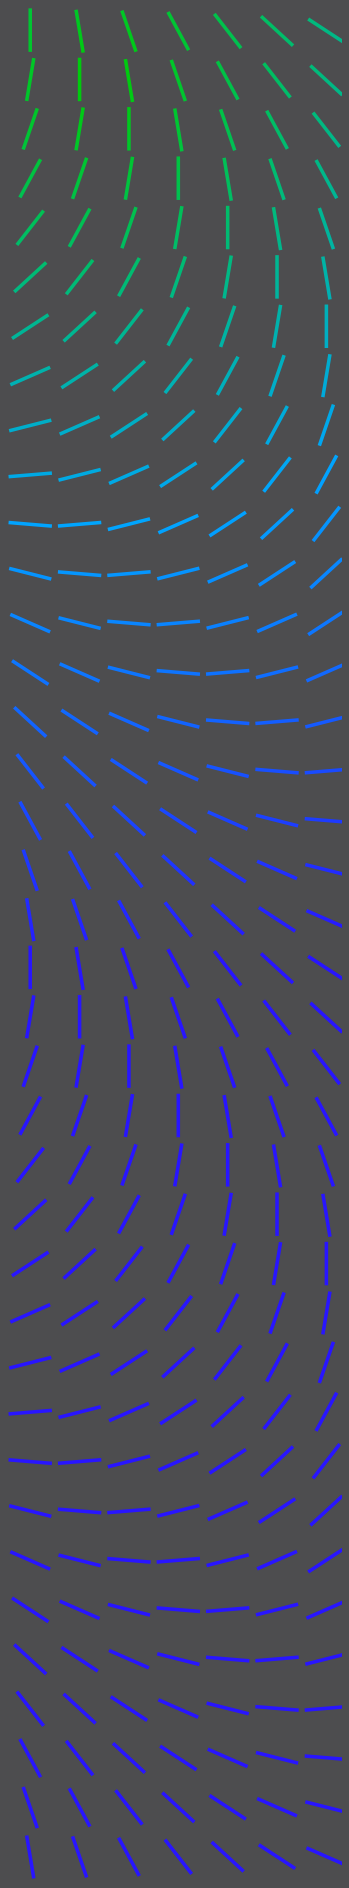

5 UTA0355: WESTERN POLICY AND DIPLOMATIC SPEAR-PHISHING CAMPAIGN

Incident summary

On November 14, 2025, the Russian threat actor UTA0355 launched a sophisticated spear-phishing campaign against the U.S. organization. The operation utilized a Belgrade Security Conference lure to target policy professionals with a fraudulent registration domain, bsc2025[.]org.

The intrusion began when the attacker, impersonating a researcher via a Gmail account, sent emails designed to trick recipients into an OAuth Device Code authentication workflow. While early attempts were blocked, a follow-up email on November 18 bypassed defenses and reached a recipient. The attack chain relied on harvesting OAuth tokens, which allowed the actor to bypass multi-factor authentication (MFA) and gain persistent access to Microsoft 365 accounts.

Once access was established, the attacker registered rogue devices in Microsoft Entra ID to mask data exfiltration and maintain stealth. This incident underscores UTA0355's focus on Western think tanks and their transition toward OAuth abuse to circumvent traditional perimeter security. Such compromises pose a severe risk of intelligence gathering and lateral movement within sensitive national security networks.

TTP Progression

Phase	Technique	Observed adversary behavior
Initial access	Phishing: spearphishing link (T1566.002)	Targeted emails were sent to policy professionals containing a link to a fake conference registration portal.
Persistence	Account manipulation: additional cloud credentials (T1098.001)	After obtaining access, attackers registered new devices in Microsoft Entra ID to maintain persistent access to the tenant.
Defense evasion	Impersonation (T1656)	The campaign posed as a legitimate researcher from the Belgrade Centre for Security Policy.
Credential access	Steal application access token (T1528)	The attack used OAuth Device Code workflows to trick users into granting permissions to their Microsoft 365 accounts.
Credential access	Input capture: web portal capture (T1056.003)	A fraudulent landing page was used to harvest corporate credentials and email addresses.
Discovery	Account discovery: cloud account (T1087.004)	The threat actor identified and targeted specific high-value cloud accounts within the organization.
Collection	Email collection: remote email collection (T1114.002)	Attackers gained unauthorized access to collect sensitive data from email, OneDrive, and Teams.
Command and control	Proxy: external proxy (T1090.002)	Residential proxy networks, such as Comcast IP addresses, were used to mask the attacker's true geographic origin.

Threat hunting process

The investigation began with Volexity's "Dangerous Invitations" threat intelligence reporting on UTA0355's campaign spoofing European security events for OAuth phishing. This OSINT provided critical context on Russian threat actors impersonating conferences like the Belgrade Security Conference and Brussels Indo-Pacific Dialogue.

Thereafter, we proactively hunted our email telemetry for the bsc2025[.]org domain and related infrastructure patterns. The finding came when pivoting from the malicious domain to identify the sender persona "Ivana Ranković <ivanaarankovic@gmail[.]com>" impersonating Belgrade Centre for Security Policy researchers.

Our hunting methodology involved:

- OSINT correlation using Volexity's campaign intelligence.
- Proactive domain searches for conference impersonation patterns.
- Email metadata analysis to identify sender personas and delivery patterns.
- Correlation of spam scores to detect evasion tactics: Trellix Email Security provided existing detection coverage with spam scores of 8.608 for initial emails.

Remediation next steps

To remediate similar spear-phishing campaigns, immediately revoke all OAuth tokens and active sessions for compromised accounts. Audit and remove unauthorized OAuth application permissions and rogue devices registered in Microsoft Entra ID. Organizations should block the device code authentication flow via Conditional Access policies and restrict user consent for unverified apps. Finally, conduct targeted training on rapport-building social engineering and OAuth phishing tactics to prevent future exploitation.

Threat hunting tips

The RU spear-phishing campaign offers several critical takeaways for threat-hunting teams.

- **Leverage proactive threat intelligence:** Subscribe to high-quality research, such as Volexity's reporting.
- **Cross-reference threat intel:** Compare public threat intelligence reports with your telemetry using campaign patterns, infrastructure IOCs, and targeting profiles.
- **Identify conference impersonation patterns:** Monitor newly registered domains (<90 days) that contain security/policy event keywords, especially during conference seasons. UTA0355's bsc2025[.]org was registered 29 days before the attacks.

Cross-reference public threat intel against your telemetry using campaign patterns, infrastructure IOCs, and targeting profiles.

THE TRELLIX SECONDSIGHT ADVANTAGE

Trellix built SecondSight on the principle that we cannot wait for perfect telemetry while an adversary is active. While security products are excellent at surfacing data, sophisticated attackers often hide in the noise of legitimate administrative activity. SecondSight bridges this gap by augmenting your SOC with elite human hunters who provide a "second set of eyes" over your environment.

CONCLUSION

As the real-world incidents in this report reveal, modern adversaries do not rely on a single playbook. They are agile, patient, and increasingly adept at blending into legitimate environments. From the SideWinder group's use of geo-fenced ClickOnce installers to UTA0355's abuse of OAuth workflows, the threat landscape is defined by "weak signals" often bypassing traditional defenses.

Key takeaways

- **Adversaries prioritize trust and context:** Modern campaigns, such as those by Kimsuky, utilize "long-game" social engineering to build trust before delivering a payload, making initial detection through automated filters significantly more difficult.
- **Legitimate tools are becoming primary vectors:** Threat actors consistently abuse native Windows binaries and legitimate software (e.g., MagTek or Canon utilities) to sideload malware and mask their execution, turning trusted applications into vehicles for compromise.
- **Proactive visibility is non-negotiable:** Relying on stale threat feeds is no longer sufficient. Effective defense requires immediate analysis of new infrastructure and hunting for process-tree anomalies, such as unusual child processes spawned from web servers like w3wp.exe.

To keep pace with modern adversaries and their evolving tactics, organizations should prioritize a proactive threat intelligence security strategy. Operating alongside our customers, Trellix combines global intelligence with real-world telemetry to expose the gaps between perceived and actual network security. With Trellix SecondSight, we apply human intuition to raw product data to understand the "intent" behind the signal, turning threats into concrete, actionable steps to reduce organizational risk.

Amidst an ever-evolving threat landscape, Trellix remains steadfast in our commitment to empowering our customers, partners, and communities with actionable threat intelligence to strengthen resilience.

Nanhi Singh
President and Chief Customer Officer, Trellix

Learn more about how Trellix SecondSight turns your telemetry into decisive defensive action at Trellix.com/SecondSight

Contributors

Ale Houspanossian
Duy-Phuc Pham
Ernesto Fernández Provecho
Heather Mackey
Ilya Kolmanovich
Jenn Jackson
John Fokker
John Wells
Megan Haley
Nanhi Singh
Ryan Delany
Alex Lanstein

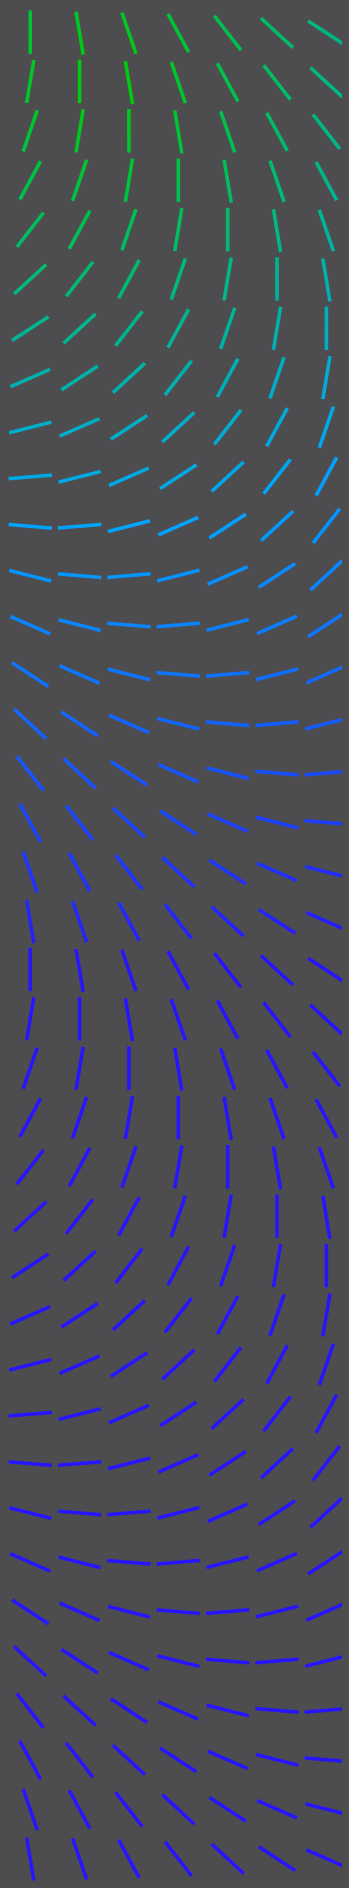

Trellix