TLP:CLEAR

ADGM THREAT INTELLIGENCE NEWSLETTER

PREPARED BY ADGM CTI

CATEGORY	ACTIONABLE
AUDIENCE	ADGM FSRA ENTITIES
DATE	5/3/2026
OVERALL THREAT SCORE	GUARDED
TARGET SECTOR	FINANCIAL SERVICES
TARGET REGION	UAE, MENA & GLOBAL
ATTRIBUTION	MULTIPLE
TLP	CLEAR

WEEKLY SUMMARY REPORT – 5 March 2026

11	4	0	0
Campaigns	**Vulnerability**	**Cyber Breach**	**Threat Actors**
Threat Campaigns of Potential Relevance to Finance Sector	Actively Exploited & Critical Vulnerabilities	Major Compromises and Breaches	Threat actor activities in the UAE & Middle East impacting Finance Sector

Summary

This week's cybersecurity newsletter highlights state linked espionage, credential phishing, supply chain abuse in developer ecosystems, and stealthy remote access tooling. Key campaigns include MuddyWater spear phishing with macro documents, GRIDTIDE use of cloud API based C2, Ruby Jumper operations across air gapped networks using removable media, and multistage Agent Tesla infections. We also track a malicious Go module that deploys the Rekoobe backdoor, a fake Zoom update that installs Teramind for covert monitoring, and the Moonrise RAT with low detection. Developer workflows face risk from trojanized Next.js repositories and a malicious npm package named "ambar-src". For the financial sector, the most material exposure is unauthorized access and data exfiltration driven by stolen credentials and persistent footholds. Act now by patching actively exploited and critical issues in Cisco Catalyst SD WAN, ServiceNow AI Platform, Google Chrome, and Trend Micro Apex One. Reinforce controls with EDR and MFA, tighten developer trust boundaries, restrict removable media, and monitor API and cloud service usage as well as Ivanti Connect Secure and other edge appliances.

ADGM THREAT INTELLIGENCE SUMMARY

[Critical Vulnerabilities in Trend Micro Apex One Could Allow Remote Code Execution] [Vulnerability] [Medium]

Name	Threat Severity Rating	TLP	Attribution	Originating Source
MuddyWater APT Campaign Targets Middle Eastern Sectors with Spear-Phishing Techniques	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers at Genians have identified cyber operations attributed to the MuddyWater APT group, targeting critical sectors in the Middle East. The group employs spear-phishing emails with malicious Microsoft Office documents to gain initial access, leveraging social engineering to manipulate users into enabling macros that execute malicious payloads.

These tactics pose significant risks to organizations in the financial services sector, particularly in investment banking, wealth management, and fintech. As MuddyWater continues to use legitimate remote management tools and advanced social engineering, financial institutions must remain vigilant against potential long-term infiltration and intelligence gathering by state-sponsored actors.

Technical Details

- The MuddyWater APT group is linked to Iran, indicating state-sponsored motivations behind their attacks.
- Spear-phishing emails are the primary vector for initial access, often containing malicious Microsoft Word documents that prompt users to enable macros.
- Attackers exploit social engineering techniques to encourage users to enable macros, especially in environments with outdated software.
- Post-compromise tactics include PowerShell-based script execution, DLL side-loading, and the abuse of legitimate remote management tools for persistence.
- Recent attacks have shown the use of Rust-based malware, indicating an evolution in the group's capabilities and sophistication.
- The group has targeted sectors such as telecommunications, energy, and government, reflecting a strategic focus on long-term intelligence collection.
- Document-based threats that utilize OLE and embedded object features remain prevalent in their attack methodology.
- EDR (Endpoint Detection and Response) is emphasized as a critical countermeasure to detect and respond to these sophisticated attacks.
- The group has demonstrated a pattern of reusing attack infrastructure and techniques across multiple campaigns.
- Traditional perimeter security measures are insufficient against the sophisticated tactics employed by this APT group.

Recommendations

- Financial institutions should implement robust EDR solutions to enhance visibility into endpoint behaviors and detect anomalous activities.
- Regularly update and patch software to mitigate vulnerabilities that could be exploited by macro-based attacks.
- Conduct employee training on recognizing phishing attempts and the dangers of enabling macros in documents from unknown sources.
- Employ multi-factor authentication (MFA) to add an additional layer of security against unauthorized access.
- Monitor network traffic for unusual patterns that may indicate the presence of remote management tools being abused.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Cyber Espionage Campaign GRIDTIDE Targeting Critical Sectors	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Google Threat Intelligence Group (GTIG) and Mandiant have successfully disrupted the GRIDTIDE global cyber espionage campaign, which was targeting telecommunications and government organizations across 42 countries. The threat actor, identified as UNC2814, utilized API calls to legitimate SaaS applications to mask their command-and-control (C2) communications, enhancing the stealth of their operations.

This campaign may impact organizations globally, particularly those within sectors that handle sensitive information. The use of cloud-hosted products for malicious purposes highlights the evolving tactics of threat actors, and organizations should remain vigilant against similar techniques that could lead to unauthorized access and data exfiltration.

Technical Details

- The GRIDTIDE campaign was executed by UNC2814, a suspected PRC-nexus cyber espionage group.
- Attackers employed API calls to Google Sheets as a C2 channel, disguising malicious traffic as legitimate.
- The GRIDTIDE backdoor allows the execution of arbitrary commands, file uploads, and downloads, leveraging cloud infrastructure for stealth.
- Initial detection was triggered by suspicious activity on a CentOS server, revealing a malicious binary named "xapt" that escalated privileges to root.

- The threat actor created a service for the malware in the systemd directory to ensure persistence within the compromised environment.
- A SoftEther VPN Bridge was deployed to establish an encrypted outbound connection, indicating long-term infrastructure use since 2018.
- The campaign targeted personally identifiable information (PII), including names, phone numbers, and national IDs, consistent with cyber espionage objectives.
- Historical data suggests that similar campaigns have led to the exfiltration of sensitive communications and call data records.
- GRIDTIDE employs AES-128 encryption for its configurations, requiring a cryptographic key for execution.
- The backdoor sanitizes its Google Sheet by deleting previous data to avoid interference with current operations.

Recommendations

- Implement strict access controls and monitor API usage to detect any unauthorized access attempts.
- Regularly audit cloud services and associated accounts for unusual activity or configurations.
- Employ endpoint detection and response (EDR) solutions to identify and mitigate suspicious processes and privilege escalations.
- Educate employees on the risks associated with cloud-based tools and the importance of reporting suspicious activities.
- Develop an incident response plan that includes procedures for addressing potential cyber espionage activities.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Punchbowl Phishing Campaign Exploits Digital Invitations to Steal Credentials	**MEDIUM**	**CLEAR**	**Campaign**	**Open Source**

Executive Summary

Researchers at Cofense Phishing Defense Center have identified a phishing campaign utilizing digital invitation platforms such as Punchbowl and Paperless Post. This campaign leverages the familiarity of these platforms to trick recipients into logging in, ultimately redirecting them to a phishing site that mimics well-known brands like Microsoft and Google to extract credentials.

The implications of this campaign may impact organizations within the financial services sector, as stolen credentials can lead to account access, identity theft, and further compromises. Financial institutions

should be aware of this evolving threat and consider the potential risks associated with credential theft through seemingly innocuous digital invitations.

Technical Details

- The phishing campaign utilizes digital invitations that appear legitimate, prompting users to log in to view event details.
- Upon interaction, users are redirected to a phishing site featuring familiar brand login options, increasing the likelihood of credential submission.
- The phishing page is designed to return fake error messages, encouraging users to try different credentials even if they are correct.
- Credentials entered on the phishing site are exfiltrated to a domain controlled by the threat actor.
- Threat actors often register new domains for phishing sites, allowing them to control DNS records and evade detection by security tools.
- New domains lack historical data, making them less likely to be flagged by reputation-based security measures.
- Stolen credentials can be sold on the dark web or used for credential stuffing attacks, particularly in cases where users reuse passwords.
- The campaign poses risks of privilege escalation and business email compromise, especially affecting corporate email accounts.
- Identity theft, fraud, and extortion are potential outcomes of credential theft from this campaign.
- Threat actors may include stolen credentials in botnets for conducting further attacks.

Recommendations

- Recipients should verify the authenticity of digital invitations by contacting the host directly using verified contact information.
- Be vigilant when redirected to login screens after attempting to RSVP; inspect the page and address bar for anomalies.
- Report any irrelevant or suspicious invitations to the appropriate security team.
- If credentials are entered on a malicious site, reset passwords immediately and monitor accounts for suspicious activities.
- Enable two-factor or multi-factor authentication on accounts to enhance security against credential theft.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
APT37's Ruby Jumper Campaign Enhances Capabilities for Air-Gapped Networks	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

A campaign linked to APT37, also known as ScarCruft, was discovered, utilizing Windows shortcut files to initiate attacks. This campaign, termed Ruby Jumper, employs a range of newly identified tools to facilitate surveillance and data exfiltration, particularly targeting air-gapped networks through removable media.

The implications of this campaign may impact organizations within the financial services sector, particularly those that handle sensitive data or operate in environments with strict network segmentation. Institutions should be aware of the evolving tactics employed by APT37, as the use of removable media for command and data transfer poses unique challenges for maintaining security in air-gapped systems.

Technical Details

- APT37's Ruby Jumper campaign uses malicious Windows shortcut (LNK) files to launch an attack via PowerShell commands.
- The initial implant, RESTLEAF, utilizes Zoho WorkDrive for command and control (C2) communications to fetch additional payloads.
- SNAKEDROPPER, a next-stage loader, installs the Ruby runtime and drops additional malware, including THUMBSBD and VIRUSTASK.
- THUMBSBD acts as a backdoor, using removable media to relay commands and facilitate data transfer between connected and air-gapped systems.
- VIRUSTASK is designed to weaponize removable media, replacing legitimate files with malicious LNK shortcuts to spread malware.
- FOOTWINE, a backdoor delivered later in the attack chain, offers surveillance capabilities such as keylogging and audio/video capturing.
- The campaign employs a two-stage shellcode execution process, making detection more challenging.
- APT37 leverages cloud services for C2 communications, indicating a shift in their operational tactics.
- The malware utilizes obfuscation techniques, including XOR encryption, to conceal payloads and evade detection.
- The campaign's architecture allows for bidirectional command delivery and data exfiltration across air-gapped networks.

Recommendations

- Implement strict controls on the use of removable media within organizational networks to mitigate potential threats.
- Enhance endpoint monitoring to detect unusual activities associated with removable media connections.

- Educate employees on the risks of opening unknown files, particularly LNK files, to prevent initial infection vectors.
- Utilize advanced threat detection solutions capable of identifying obfuscated payloads and suspicious PowerShell activity.
- Regularly update and patch systems to reduce vulnerabilities that could be exploited by such campaigns.

For a detailed breakdown of tactics, techniques, and procedures (TTPs), refer to the Annexure – APT37's Ruby Jumper Campaign

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Multi-Stage Campaign Utilizing Agent Tesla Targets Windows Users	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Fortinet has identified a multi-stage campaign leveraging Agent Tesla, a persistent threat that enables data harvesting through sophisticated phishing techniques. The attack begins with a deceptive email containing a compressed RAR file that, when executed, initiates a complex infection chain designed to evade detection and extract sensitive information.

The implications of this campaign may impact organizations in the financial services sector, particularly those utilizing Windows systems. As the malware employs advanced evasion techniques and targets sensitive data, financial institutions should be aware of the potential risks associated with phishing attacks and the importance of robust security measures to mitigate such threats.

Technical Details

- The campaign initiates with a phishing email containing a RAR attachment designed to bypass email filters.
- An obfuscated JScript loader (.jse) is executed from the RAR file, which serves as the initial downloader.
- The loader fetches an encrypted PowerShell script from an external file-hosting service to continue the infection process.
- The PowerShell script uses in-memory decryption to avoid leaving traces on the disk, enhancing stealth.
- The malware employs process hollowing to inject malicious code into a legitimate Windows process, making detection difficult.

- Defensive checks are performed to ensure the environment is not a virtual machine or sandbox, leading to disruption in analysis.
- Agent Tesla systematically extracts sensitive data, including browser cookies and contacts, for exfiltration.
- Stolen data is sent to the attacker's command-and-control mail server via SMTP, indicating a large-scale operation.
- The campaign showcases the adaptability of Agent Tesla, allowing low-skilled actors to execute complex attacks.

Recommendations

- Implement advanced email filtering solutions to detect and block phishing emails and malicious attachments.
- Educate employees on recognizing phishing attempts and the importance of verifying unexpected emails.
- Utilize endpoint detection and response (EDR) solutions to monitor for process hollowing and memory-based attacks.
- Regularly update and patch systems to protect against vulnerabilities that malware may exploit.
- Conduct simulated phishing exercises to enhance user awareness and resilience against real-world threats.

For a detailed breakdown of tactics, techniques, and procedures (TTPs), refer to the Annexure – Multi-Stage Campaign Utilizing Agent Tesla

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Malicious Go Module Impersonates Legitimate Library to Deploy Rekoobe Backdoor	**MEDIUM**	**CLEAR**	**Campaign**	**Open Source**

Executive Summary

Researchers at Socket's Threat Research Team have uncovered a malicious Go module, which impersonates the legitimate codebase. This module captures passwords entered through a modified ReadPassword function, exfiltrates them to a threat actor-controlled endpoint, and executes a remote shell stager that installs a Rekoobe backdoor on Linux systems.

The implications of this campaign may impact organizations in the financial services sector, particularly those utilizing Go-based applications. The use of a trusted library for malicious purposes highlights the need for vigilance in dependency management and the potential risks associated with supply chain attacks. Financial

institutions should be aware of the evolving tactics employed by threat actors to exploit high-value libraries and frameworks.

Technical Details

- The malicious module mimics the legitimate golang[.]org/x/crypto library, targeting its high trust within the Go ecosystem.
- It captures passwords through a modified ReadPassword function, storing them locally before exfiltration.
- The module reaches out to a GitHub-hosted resource for follow-on instructions, facilitating further malicious actions.
- A shell script is executed to append a threat actor controlled SSH key to the authorized_keys file, ensuring persistent access.
- The script alters iptables default policies to ACCEPT, weakening the host's firewall defenses.
- Payloads disguised with the ".mp5" extension are downloaded and executed, reducing detection risk during manual review.
- The campaign utilizes a multi-stage Linux dropper chain, with several network hops following the initial password capture.
- The Rekoobe backdoor, confirmed as one of the payloads, is known for its use in espionage-oriented operations.
- The threat actor maintains a GitHub account with multiple repositories, indicating ongoing operational relevance.
- The campaign's low-effort, high-impact nature suggests potential for recurrence, targeting other credential management libraries.

Recommendations

- Treat Go module roots as critical supply chain boundaries and review dependency changes as security sensitive.
- Implement endpoint and CI detections for behaviors associated with the malicious module, such as unauthorized writes to sensitive paths.
- Block suspicious utility additions that enable outbound network access or shell execution in dependency management.
- Utilize security tooling to scan for and flag suspicious module introductions before merging code changes.
- Deploy firewall rules to block known malicious packages and transitive dependencies before they are fetched.

For a detailed breakdown of tactics, techniques, and procedures (TTPs), refer to the Annexure – Malicious Go Module Impersonates Legitimate Library

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
RESURGE Malware Campaign Targets Critical Infrastructure via Ivanti Connect Secure Exploits	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

The Cybersecurity and Infrastructure Security Agency (CISA) has provided an updated analysis of the RESURGE malware, which exploits Ivanti CVE-2025-0282 to gain initial access to critical infrastructure devices. This malware remains dormant until activated by a remote actor, utilizing advanced techniques for command and control, including forged TLS certificates and encryption methods.

The implications of RESURGE are potentially relevant to organizations in the financial services sector, particularly those utilizing Ivanti Connect Secure devices. The malware's sophisticated capabilities for stealth and command execution may impact operational integrity and security, necessitating heightened vigilance and proactive defense measures.

Technical Details

- RESURGE exploits Ivanti CVE-2025-0282 to gain initial access to devices, marking the beginning of the attack chain.
- The malware hooks into the native Ivanti web server process, monitoring TLS packets to differentiate between legitimate and malicious connections.
- It employs a CRC32 fingerprint hashing scheme to authenticate incoming TLS HELLO packets, enhancing its stealth.
- If a connection is deemed malicious, RESURGE responds with a forged TLS SERVER HELLO packet, establishing a covert communication channel.
- The malware utilizes advanced cryptographic methods, including Elliptical Curve Cryptography (ECC), for secure communications.
- Forged TLS certificates are generated to facilitate covert interactions, serving as network indicators for detection.
- RESURGE mimics legitimate TLS/SSH traffic to evade detection, complicating remediation efforts.
- The malware can modify critical files and processes, further enhancing its control over compromised devices.
- It remains latent on devices, waiting for a remote actor to initiate a connection, which poses a significant threat to critical infrastructure.
- The similarities to SPAWNCHIMERA Malware in command-and-control capabilities underscore the evolving nature of such threats.

Recommendations

- Implement robust monitoring and detection mechanisms for Ivanti Connect Secure devices to identify potential RESURGE activity.
- Regularly update and patch systems to mitigate vulnerabilities associated with CVE-2025-0282.
- Utilize the IOCs and detection signatures provided by CISA to enhance threat detection capabilities.
- Conduct regular security assessments and penetration testing to identify and remediate potential vulnerabilities.
- Educate employees on the risks associated with phishing and other social engineering tactics that could facilitate malware deployment.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Fake Zoom Meeting Update Campaign Installs Surveillance Software on Victims' Machines	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

A campaign has been identified where a fake Zoom meeting website trick users into downloading surveillance software onto their Windows machines. This malicious operation masquerades as a legitimate Zoom video call, automatically prompting users with an “Update Available” message that leads to the silent installation of a Teramind monitoring tool without user consent.

The misuse of Teramind's legitimate software poses potential risks to individuals and organizations in the financial services sector. Cybercriminals are leveraging this trusted tool to conduct unauthorized surveillance, which may impact privacy and security. Organizations in the financial sector should be aware of this campaign as it highlights the growing trend of attackers abusing legitimate software for malicious purposes.

Technical Details

- The campaign operates through a fake Zoom meeting website, which simulates a Zoom waiting room.
- Upon loading, the site sends a notification to the attackers indicating a visitor's arrival.
- The site features scripted fake participants to create an illusion of a live meeting, enhancing the deception.
- After 10 seconds, an “Update Available” pop-up appears, prompting an automatic download of a malicious installer.
- The downloaded file is designed to mimic legitimate software installations.
- The installer is configured to connect to an attacker-controlled Teramind server, allowing unauthorized surveillance.

- The installation process is designed to be stealthy, avoiding typical user prompts and detection by antivirus software.
- The installer employs techniques to detect sandbox environments, altering its behavior to evade analysis.
- Once installed, the monitoring agent runs in the background, collecting sensitive information without user knowledge.
- The campaign illustrates a significant risk as it utilizes a legitimate software product for malicious purposes, complicating detection efforts.

Recommendations

- Financial institutions should educate employees on verifying meeting links and using official applications directly.
- Implement strict policies regarding software installations and updates, ensuring they are sourced from trusted vendors.
- Regularly monitor systems for unauthorized installations and unusual network activity.
- Encourage the use of endpoint protection solutions that can identify and block suspicious software behavior.
- Establish a response plan for potential compromises, including immediate reporting and remediation steps.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Moonrise RAT: A Low-Detection Remote Access Threat and Operational Exposure	**MEDIUM**	**CLEAR**	**Campaign**	**Open Source**

Executive Summary

ANY.RUN experts have identified a new Go-based remote access trojan (RAT) named Moonrise, which operates without early static detection and can execute credential theft, remote command execution, and user monitoring. This RAT establishes active command and control (C2) communication before triggering vendor alerts, posing a significant risk to organizations that may not detect its presence in a timely manner.

The implications of Moonrise for the financial services sector are noteworthy. Organizations should be aware that the stealthy nature of this RAT may lead to prolonged dwell times, increasing the risk of data loss, operational disruption, and potential financial impact. The ability of Moonrise to facilitate remote control and data extraction could expose sensitive information and disrupt critical business operations.

Technical Details

- Moonrise operates without early static detection, establishing C2 communication before vendor alerts are triggered.
- The RAT supports credential theft, remote command execution, persistence, and user monitoring, enabling full control of infected endpoints.
- Silent C2 activity increases business exposure, extending dwell time and raising risks of data loss and operational disruption.
- Credential theft and clipboard monitoring can expose passwords and sensitive data copied between systems.
- Remote command execution allows operators to run scripts and manipulate business applications.
- File upload and execution capabilities create paths for additional payloads, including stealers or ransomware.
- Extensive user interaction monitoring can reveal sensitive activities within finance workflows and internal communications.
- Persistence and privilege-related functions make removal of the RAT more challenging.
- The lifecycle management functions of Moonrise facilitate ongoing control and disruption of operations.
- Early detection relies on behavior-based analysis rather than static reputation checks.

Recommendations

- Implement behavior-based detection mechanisms to identify unusual activity indicative of RAT presence.
- Enhance monitoring capabilities to catch early signals of potential compromise.
- Conduct regular security training for staff to recognize signs of credential theft and suspicious activities.
- Establish a rapid triage process to confirm potential threats when static checks fail.
- Review and update incident response plans to include scenarios involving stealthy RATs like Moonrise.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Threat Actors Abusing Next.js Repositories to Target Developer Workflows	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Microsoft Defender Experts have identified a coordinated developer-targeting campaign that utilizes malicious repositories disguised as legitimate "Next.js" projects. The campaign employs job-themed lures to blend into routine developer workflows, increasing the likelihood of code execution through various entry points. The malicious repositories lead to the execution of attacker-controlled JavaScript, facilitating command-and-control operations and data exfiltration.

This campaign may impact organizations in the financial services sector, particularly those with development teams that interact with external code repositories. Developers should be aware of the risks associated with running untrusted projects, as these compromises can expose sensitive assets and credentials, potentially leading to broader security incidents.

Technical Details

- The campaign leverages malicious "Next.js" repositories, disguised as legitimate projects, to execute JavaScript on developer systems.
- Multiple execution paths are employed, including Visual Studio Code workspace automation and build-time execution during application development.
- Malicious logic is embedded in seemingly legitimate application assets, such as modified JavaScript libraries, to act as loaders.
- The campaign initiates with a lightweight registration stage that establishes host identity and facilitates command-and-control communication.
- A dynamic remote code execution mechanism is utilized, allowing attackers to execute JavaScript returned from their servers in memory.
- The campaign's execution paths are designed to trigger during normal developer activities, increasing the chances of successful exploitation.
- Staged uploads and directory browsing are observed, indicating potential data exfiltration activities.
- The use of base64-encoded endpoints and environment variable exfiltration enhances the attack's stealth and effectiveness.
- The campaign's infrastructure includes multiple command-and-control IP addresses and domains for tasking and data exfiltration.
- The overall design supports operator-driven discovery and follow-on payload delivery, posing a significant risk to developer environments.

Recommendations

- Implement strict trust boundaries for developer workflows, ensuring that untrusted repositories are not executed without thorough review.

- Utilize endpoint telemetry to monitor for unusual Node.js execution patterns and unexpected outbound connections.
- Enforce strong authentication and conditional access controls to protect sensitive assets on developer systems.
- Regularly review and audit workspace automation files in Visual Studio Code to prevent automatic execution of untrusted code.
- Centralize monitoring and hunting activities to detect and respond to suspicious behaviors associated with this campaign.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
New Malicious npm Package Targeting Developers and Raising Supply Chain Risks	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Tenable Research has identified a malicious npm package named "ambar-src" that was downloaded approximately 50,000 times before its removal. This package employs various detection-evasion techniques and executes harmful code during the installation process, compromising systems without the need for explicit invocation by developers. The package was designed to mimic legitimate software, making it particularly dangerous for developers across multiple operating systems, including Windows, Linux, and macOS.

The implications of this campaign may impact organizations in the financial services sector, particularly those utilizing open-source software in their development processes. The use of malicious npm packages like "ambar-src" highlights the risks associated with supply chain vulnerabilities, and financial institutions should be aware of the potential for similar attacks that could lead to significant operational disruptions and data breaches.

Technical Details

- The "ambar-src" package utilized npm's preinstall script hook, allowing it to execute malicious code during installation without user interaction.
- The package was designed to evade detection through various techniques, including hex-encoding malicious command strings.
- Upon installation, it executes OS-specific one-liner shell commands to fetch and run additional malware from remote servers.
- For Windows systems, the payload downloaded is named "msinit[.]exe", which contains encrypted shellcode.

- Linux systems receive a bash script that fetches an ELF binary, saving it as "osa" for execution.
- macOS systems execute a JavaScript payload fetched using "osascript", which is a built-in utility for executing scripts.
- The initial infection vector involved typo squatting, mimicking a popular package to lure unsuspecting developers.
- The campaign demonstrates a sophisticated level of supply chain attack, with the malicious package gaining significant traction before detection.
- All versions of "ambar-src" are deemed malicious, as it lacks any legitimate use cases.
- The malware communicates with known domains to exfiltrate data, leveraging common web services to mask its activity.

Recommendations

- Organizations should conduct thorough audits of their development environments to identify the presence of the "ambar-src" package.
- Implement strict controls and monitoring for npm package installations to prevent unauthorized software from being executed.
- Educate developers on the risks associated with installing packages from public registries and the importance of verifying package integrity.
- Utilize security tools that can detect and alert on the installation of known malicious packages in development environments.
- Establish incident response protocols to quickly address any potential compromises resulting from malicious package installations.

For the indicators of compromise (IOCs), refer to the attached CSV sheet.

Reference to the Source

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Critical Authentication Bypass Vulnerability in Cisco Catalyst SD-WAN Actively Exploited	**HIGH**	**CLEAR**	**Vulnerability**	**CSC**

Executive Summary

Cisco Catalyst SD-WAN Controller and Manager are affected by a critical authentication bypass vulnerability, CVE-2026-20127, which allows unauthenticated remote attackers to gain administrative access. This vulnerability is being actively exploited, enabling attackers to perform software downgrade attacks and escalate privileges using CVE-2022-20775, ultimately leading to root-level command execution.

The implications of this vulnerability may impact organizations in the financial services sector that utilize Cisco's SD-WAN solutions. Financial institutions should be aware of the potential for unauthorized access

and subsequent manipulation of network infrastructure, which could affect the integrity and confidentiality of sensitive financial data.

Technical Details

- CVE-2026-20127 is a critical authentication bypass vulnerability with a CVSS score of 10.0, allowing remote, unauthenticated access.
- Attackers can exploit this vulnerability to gain administrative access to Cisco Catalyst SD-WAN systems without needing credentials.
- Once access is obtained, attackers can downgrade the software to an older, vulnerable version, facilitating further exploitation.
- CVE-2022-20775, a privilege escalation vulnerability, is then exploited to execute arbitrary commands as the root user.
- This chained attack sequence allows attackers to manipulate network routing and intercept traffic.
- The exploitation status of CVE-2026-20127 indicates it is actively being exploited in the wild.
- Attackers can establish persistence and move laterally within the network post-exploitation.

Recommendations

- Upgrade all affected SD-WAN systems to fixed versions to mitigate the vulnerability.
- Block unauthorized management-plane access from the internet to prevent exploitation.
- Audit control-plane peering logs for anomalies to detect suspicious activity.
- Disable unused administrative accounts to reduce potential attack vectors.
- Rotate all SD-WAN administrative credentials to enhance security.

Detailed Vulnerability Details and Affected Products can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Multiple High-Severity Vulnerabilities Discovered in Google Chrome	**HIGH**	**CLEAR**	**Vulnerability**	**CSC**

Executive Summary

Google has released a Stable Channel update addressing three high-severity vulnerabilities in the desktop versions of Chrome. These vulnerabilities involve out-of-bounds memory operations and an inappropriate implementation issue, which could lead to memory corruption, browser crashes, or arbitrary code execution under specific conditions.

Organizations in the financial services sector should be aware of these vulnerabilities, as they may impact the security of web applications and online banking services. Ensuring that browsers are updated to the latest versions is crucial to mitigate potential risks associated with these flaws.

Technical Details

- Three high-severity vulnerabilities have been identified in Google Chrome, affecting desktop versions.
- CVE-2026-3061 involves an out-of-bounds read in the Media component, which could lead to memory corruption.
- CVE-2026-3062 includes both out-of-bounds read and write vulnerabilities in the Tint component, posing risks of browser crashes.
- CVE-2026-3063 pertains to an inappropriate implementation issue in DevTools, which could allow for arbitrary code execution.
- Exploitation of these vulnerabilities may occur under specific conditions, emphasizing the need for immediate updates.
- The vulnerabilities have been addressed in the latest Stable Channel update for Chrome.
- Fixed versions include 145.0.7632.116 and 145.0.7632.117 for Windows and Mac, and 144.0.7559.116 for Linux.

Recommendations

- Update Google Chrome to the latest version immediately to mitigate vulnerabilities.

Detailed Vulnerability Details and Affected Products can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Critical Vulnerability in ServiceNow AI Platform Enables Remote Code Execution	**HIGH**	**CLEAR**	**Vulnerability**	**CSC**

Executive Summary

A critical vulnerability, identified as CVE-2026-0542, has been discovered in the ServiceNow AI Platform, which allows unauthenticated remote code execution within the platform's AI Sandbox. This vulnerability poses a significant risk as it could enable attackers to execute arbitrary code without authentication, potentially compromising the integrity of the affected systems.

Organizations in the financial services sector should be aware of this vulnerability, as it may impact their operations if they utilize the ServiceNow AI Platform. The critical nature of this vulnerability highlights the importance of applying security patches promptly to mitigate potential risks associated with unauthorized access and exploitation.

Technical Details

- CVE-2026-0542 is classified as a critical vulnerability with a CVSSv4 score of 9.2.
- The vulnerability affects the ServiceNow AI Platform specifically within its Sandbox environment.

- It allows for unauthenticated remote code execution (RCE), which can lead to arbitrary code execution.
- Attackers can exploit this vulnerability without needing any form of authentication.

Recommendations

- Verify the instance version of the ServiceNow AI Platform and ensure it is up to date.
- Apply the latest security patches as soon as they are released by ServiceNow.
- If hosted, confirm that the January 2026 security update has been applied to the instance.

Detailed Vulnerability Details and Affected Products can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Critical Vulnerabilities in Trend Micro Apex One Could Allow Remote Code Execution	**MEDIUM**	**CLEAR**	**Vulnerability**	**Open Source**

Executive Summary

Trend Micro has identified critical vulnerabilities in its Apex One management console that could allow remote attackers to execute malicious code on affected installations. Specifically, CVE-2025-71210 and CVE-2025-71211, both rated with a CVSS score of 9.8, enable directory traversal attacks, while additional vulnerabilities could lead to local privilege escalation.

Organizations in the financial services sector should be aware that these vulnerabilities may impact their security posture, particularly if the management console is exposed to the internet. It is recommended that customers apply the latest patches and review their security configurations to mitigate potential risks.

Technical Details

- CVE-2025-71210 and CVE-2025-71211 are critical vulnerabilities that allow remote code execution via directory traversal in the Apex One management console.
- Both vulnerabilities have a CVSS score of 9.8, indicating a critical severity level.
- Exploitation requires that an attacker has access to the management console, particularly if its IP address is exposed externally.
- CVE-2025-71212, CVE-2025-71213, and several other vulnerabilities allow local privilege escalation, with CVSS scores ranging from 7.2 to 7.8.
- Local privilege escalation vulnerabilities require an attacker to execute low-privileged code on the target system before exploitation.
- The vulnerabilities were reported through responsible disclosure by researchers via the Zero Day Initiative.

Recommendations

- Apply the latest Critical Patch (CP) for Trend Micro Apex One and Apex One (Mac) immediately.

Detailed Vulnerability Details and Affected Products can be found *here*.

back to top

Appendix A - Tactics, Techniques & Procedures (TTPs)

APT37's Ruby Jumper Campaign Enhances Capabilities for Air-Gapped Networks

TACTIC	TECHNIQUE
Execution	T1204.001 User Execution: Malicious Link
Execution	T1059.001 Command and Scripting Interpreter: PowerShell
Persistence	T1053.005 Scheduled Task/Job: Scheduled Task
Persistence	T1574 Hijack Execution Flow
Defense Evasion	T1027 Obfuscated Files or Information
Defense Evasion	T1055 Process Injection
Defense Evasion	T1620 Reflective Code Loading
Defense Evasion	T1036.005 Masquerading: Match Legitimate Name or Location
Defense Evasion	T1564.001 Hide Artifacts: Hidden Files and Directories
Discovery	T1082 System Information Discovery
Discovery	T1057 Process Discovery
Discovery	T1083 File and Directory Discovery
Command and Control	T1132.002 Data Encoding: Non-Standard Encoding
Command and Control	T1092 Communication Through Removable Media
Exfiltration	T1052.001 Exfiltration Over Physical Medium: Exfiltration over USB
Exfiltration	T1567.002 Exfiltration Over Web Service: Exfiltration to Cloud Storage
Collection	T1056.001 Input Capture: Keylogging
Collection	T1113 Screen Capture
Collection	T1123 Audio Capture
Collection	T1125 Video Capture

Multi-Stage Campaign Utilizing Agent Tesla Targets Windows Users

TACTIC	TECHNIQUE
Initial Access	T1566.001 Phishing: Spearphishing Attachment
Execution	T1059.001 Command and Scripting Interpreter: PowerShell
Execution	T1059.007 Command and Scripting Interpreter: JavaScript
Defense Evasion	T1055.012 Process Injection: Process Hollowing
Defense Evasion	T1620 Reflective Code Loading
Defense Evasion	T1497.001 Virtualization/Sandbox Evasion: System Checks
Credential Access	T1539 Steal Web Session Cookie
Credential Access	T1555.003 Credentials from Web Browsers
Collection	T1005 Data from Local System
Exfiltration	T1048.003 Exfiltration Over Alternative Protocol: SMTP

Malicious Go Module Impersonates Legitimate Library to Deploy Rekoobe Backdoor

TACTIC	TECHNIQUE
Initial Access	T1195.002 Supply Chain Compromise: Compromise Software Supply Chain
Execution	T1204.005 User Execution: Malicious Library
Defense Evasion	T1036 Masquerading
Defense Evasion	T1036.008 Masquerading: Masquerade File Type
Defense Evasion	T1656 Impersonation
Collection	T1056 Input Capture
Command and Control	T1071.001 Application Layer Protocol: Web Protocols
Command and Control	T1102.001 Web Service: Dead Drop Resolver
Command and Control	T1105 Ingress Tool Transfer
Execution	T1059.004 Command and Scripting Interpreter: Unix Shell
Persistence	T1098.004 Account Manipulation: SSH Authorized Keys

Defense Evasion	T1562.004 Impair Defenses: Disable or Modify System Firewall
Defense Evasion	T1070.004 Indicator Removal: File Deletion

Appendix B – Threat Severity Ratings & Definitions

In this newsletter, each campaign, breach, and vulnerability has been assigned a severity level—High, Medium, or Low—based on a contextual assessment of its impact, relevance to the financial sector, regional exposure, and exploit status.

A **High** severity rating is assigned to threats that involve widely used technologies, nation-state actors, or actively exploited zero-day vulnerabilities. These incidents pose a significant risk due to their potential for widespread disruption, data compromise, and strategic targeting of core infrastructure.

A **Medium** severity rating reflects threats that are technically impactful but have limited regional exposure or affect less prominent entities. These may include ransomware attacks on mid-sized firms, blockchain-based malware delivery mechanisms, or smart contract exploits in DeFi platforms.

A **Low** severity rating is applied to threats that are peripherally related to the financial sector, have no confirmed impact in the region, or target non-critical systems. These include credential stuffing attacks on platforms outside the MENA region or supply chain breaches that have not yet affected local infrastructure.

Threat Score Ratings & Definitions

1. **Severe**: Widespread, high-impact attacks ongoing. Critical infrastructure affected by destructive malware or multiple zero-day exploits.
2. **High**: Confirmed breaches or coordinated campaigns in progress. High-impact threat actor activity and zero-day exploitation observed.
3. **Elevated**: Active exploitation of known vulnerabilities or targeted campaigns detected. Some breaches or zero-days may surface.
4. **Guarded**: Slight increase in malicious activity. Low-level campaigns or threat actor reconnaissance may be underway.
5. **Normal**: No significant cyber activity beyond baseline. No major breaches, campaigns, or zero-day activity observed.

Appendix C – Traffic Light Protocol (TLP) Definitions and Usage

TLP	When should it be used?	How should it be shared?
TLP:Red	Sources may use TLP:RED when information cannot be effectively acted upon without significant risk for the privacy, reputation, or operations of the organizations involved. For the	Recipients may not share TLP:RED information with any parties outside of the specific exchange, meeting, or conversation in which it was originally disclosed. In the context of a meeting, for example, TLP:RED information is

	eyes and ears of individual recipients only, no further.	limited to those present at the meeting. In most circumstances, TLP:RED should be exchanged verbally or in person.
TLP:Amber+Strict	Sources may use TLP:AMBER+STRICT when information requires support to be effectively acted upon, yet carries risk to privacy, reputation, or operations if shared outside of the organization.	Recipients may share TLP:AMBER+STRICT information only with members of their own organization on a need-to-know basis to protect their organization and prevent further harm.
TLP:Amber	Sources may use TLP:AMBER when information requires support to be effectively acted upon, yet carries risk to privacy, reputation, or operations if shared outside of the organizations involved. Note that TLP:AMBER+STRICT should be used to restrict sharing to the recipient organization only.	Recipients may share TLP:AMBER information with members of their own organization and its clients on a need-to-know basis to protect their organization and its clients and prevent further harm.
TLP:Green	Sources may use TLP:GREEN when information is useful to increase awareness within their wider community.	Recipients may share TLP:GREEN information with peers and partner organizations within their community, but not via publicly accessible channels. Unless otherwise specified, TLP:GREEN information may not be shared outside of the cybersecurity or cyber defense community.
TLP:Clear	Sources may use TLP:CLEAR when information carries minimal or no foreseeable risk of misuse, in accordance with applicable rules and procedures for public release.	Recipients may share this information without restriction. Information is subject to standard copyright rules.

Appendix D - Acronyms & Technical Terms

Term / Acronym	Meaning / Description
.mp5 payloads	Files using a misleading extension to reduce detection during manual review.
AES-128	Encryption noted in GRIDTIDE configuration that requires a specific key to run.
Agent Tesla	Phishing delivered stealer that uses in memory decryption and process hollowing and sends data via SMTP.
AI Sandbox	ServiceNow component where the remote code execution issue exists.
Air gapped network	Isolated network with no internet connection that attackers reach by passing data via removable media.
ambar-src	Malicious npm package that auto executes during install using a preinstall hook and fetches OS specific payloads.
ANY.RUN	Research team that analyzed Moonrise operations and documented its exposure.

API	Application Programming Interface used by attackers to blend malicious traffic with normal cloud use.
APT	Advanced Persistent Threat. Long running targeted intrusion focused on persistence and stealth.
APT37	Threat actor also referenced in the Ruby Jumper campaign targeting air gapped environments.
authorized_keys	File on Linux where attacker SSH keys were added to keep access.
Backdoor	Malware capability that allows remote command execution and file transfer once installed.
Base64 encoded endpoints	Obfuscation used to hide command and exfiltration destinations.
Business Email Compromise (BEC)	Outcome where stolen mail access is abused for fraud and approvals.
C2	Command and Control. Infrastructure used to manage compromised systems and move data.
CentOS	Linux system where suspicious GRIDTIDE activity was observed with a binary that escalated to root.
Cisco Catalyst SD-WAN	Controller and Manager affected by a critical authentication bypass that leads to admin and then root command execution.
Cofense Phishing Defense Center	Researchers who identified the invitation themed phishing campaign.
Control-plane peering logs	Records to audit for anomalies after suspected SD-WAN exploitation.
CRC32	Fingerprint hashing used by RESURGE to check incoming TLS traffic.
Credential theft	Theft of usernames and passwords that can lead to account access and fraud.
Critical Patch (CP)	Vendor update recommended to remediate Apex One issues.
CSC	UAE Cyber Security Council
CVE-2022-20775	Privilege escalation used after the Cisco auth bypass to run commands as root.
CVE-2026-0542	Critical ServiceNow AI Platform vulnerability that enables unauthenticated remote code execution.
CVE-2026-20127	Critical authentication bypass in Cisco Catalyst SD-WAN that allows unauthenticated admin access.
DevTools	Chrome component noted with an inappropriate implementation issue that could allow code execution.
Directory traversal	Technique that allowed remote code execution in the Apex One console.
DLL side loading	Running malicious code by placing crafted DLLs where legitimate programs will load them.
Dropper	Multistage installer chain that fetches and runs additional payloads.
Dynamic remote code execution	Mechanism that fetches JavaScript from attacker servers and runs it in memory.
ECC	Elliptic Curve Cryptography referenced for secure communications in RESURGE.
EDR	Endpoint Detection and Response. Monitors endpoint behaviour and helps detect suspicious activity.
ELF	Linux binary format fetched and saved as osa by the malicious installer.
Environment variable exfiltration	Theft of secrets stored in local environment variables.
Exfiltration	Unauthorized transfer of data out of a network.
Fake Zoom update site	Website that simulates a meeting and pushes an update prompt that installs surveillance software.
FOOTWINE	Backdoor that supports surveillance like keylogging and audio or video capture.
Genians	Researchers who reported on MuddyWater operations.
golang.org/x/crypto	Trusted Go library name that was impersonated by the malicious module.

Google Chrome Stable Channel	Desktop browser builds updated to fix high severity issues including memory errors and a DevTools flaw.
Google Sheets (as C2)	Spreadsheet service used as a channel to send commands and receive data from infected hosts.
GRIDTIDE	Global espionage campaign that uses API calls to cloud apps like Google Sheets to hide command and control.
In memory decryption	Decrypting payloads in RAM to avoid leaving traces on disk.
iptables	Host firewall rules that were weakened by setting defaults to ACCEPT.
Ivanti Connect Secure	VPN device targeted for initial access in the RESURGE activity.
JScript (.jse)	Obfuscated loader executed from the archive to fetch encrypted PowerShell.
LNK (shortcut)	Windows shortcut files that are weaponized to start the attack chain.
Local Privilege Escalation (LPE)	Vulnerabilities that let a low privilege user gain higher privileges on a system.
Macros (Office)	Embedded document scripts that attackers ask users to enable which then execute malicious code.
Malicious Go module	Module that mimics golang.org/x/crypto, captures passwords via a modified ReadPassword, and deploys Rekoobe.
Management plane	Administrative interface of network systems that should be blocked from the internet.
Mandiant	Team that participated in the GRIDTIDE disruption.
MFA	Multi Factor Authentication. Adds an extra verification step to block unauthorized access.
Microsoft Defender Experts	Team that identified the coordinated developer targeting campaign.
Moonrise RAT	Go based remote access trojan that establishes C2 early and supports credential theft and remote execution.
msinit.exe	Windows payload name used by the malicious npm package.
MuddyWater	Threat actor using spear phishing with malicious Office documents, social engineering, PowerShell, and DLL side loading to gain and keep access.
Next.js malicious repositories	Trojanized projects that execute attacker controlled JavaScript during normal developer workflows.
Node.js execution patterns	Telemetry focus to detect unusual runtime behaviour and outbound connections.
npm preinstall hook	Script step that allowed the ambar src package to run code automatically on install.
OLE / Embedded objects	Document features that are abused in macro-based attack chains.
osascript	Built in macOS utility invoked to execute JavaScript payloads.
Out of bounds read/write	Memory errors that can cause crashes or code execution as noted in the Chrome entry.
Phishing	Deceptive messages that trick users into revealing credentials or running malware.
PII	Personally Identifiable Information such as names, phone numbers, and national IDs targeted in espionage.
PowerShell	Windows scripting used by attackers for execution and persistence after initial compromise.
PRC-nexus	Descriptor in the newsletter that links the actor's activity to PRC nexus in an espionage context.
Process hollowing	Injecting code into a legitimate process to evade detection.
Punchbowl / Paperless Post	Digital invitation platforms used as lures that redirect to fake brand login pages.
RAR archive	Compressed attachment used to deliver a JScript loader in the Agent Tesla chain.
ReadPassword	Function altered by the malicious module to capture credentials typed by users.
Rekoobe	Backdoor payload confirmed as part of the Go module attack.
Removable media	USB or similar media used to move commands and data between connected and isolated systems.
RESTLEAF	Initial implant that uses Zoho WorkDrive to fetch more payloads.

RESURGE	Malware that exploits Ivanti CVE 2025 0282 and uses forged TLS certificates and ECC for covert control.
Root level command execution	Highest privilege command execution observed in the Cisco exploit chain.
Ruby Jumper	Campaign that uses LNK shortcuts, loaders, and removable media to bridge air gapped networks.
SaaS	Cloud hosted products abused to mask attacker communications.
Sandbox evasion	Behaviours in the installer that change execution to avoid analysis environments.
ScarCruft	Also known as APT37 in the newsletter context.
ServiceNow AI Platform	Platform affected by an unauthenticated remote code execution issue in the AI Sandbox.
Sheet sanitization	Deleting older rows in the Google Sheet so current operations are not affected.
SMTP exfiltration	Sending stolen data to attacker mail servers over SMTP.
SNAKEDROPPER	Loader that installs the Ruby runtime and drops later stage components.
Social engineering	Manipulating people to bypass security and gain access or information.
Socket's Threat Research Team	Researchers who uncovered the malicious Go module.
SoftEther VPN Bridge	Component deployed to maintain encrypted outbound connections for long term use.
Software downgrade attack	Reverting to an older version after access to enable further exploitation.
Spear phishing	Targeted phishing emails that entice users to open macro enabled documents and trigger payloads.
Supply chain attack	Compromising third party components or packages to reach the primary target.
systemd service	Startup entry created to keep malware persistent across reboots.
Teramind	Legitimate monitoring tool that was connected to attacker servers for unauthorized surveillance.
THUMBSBD	Backdoor that uses removable media to move commands and data in and out of air gapped systems.
TLS HELLO / SERVER HELLO	TLS handshake messages that RESURGE inspects and forges to set up covert channels.
Trend Micro Apex One	Endpoint security management console with critical directory traversal and other privilege issues.
Typo squatting	Using a name similar to a popular package to lure installs.
UNC2814	Actor identified as operating the GRIDTIDE campaign using cloud based C2 methods.
VIRUSTASK	Component that replaces files on removable media with malicious LNK shortcuts.
Visual Studio Code workspace automation	Paths that can auto run tasks and were abused for execution.
XOR	Obfuscation noted to hide payloads within the Ruby Jumper activity.
Zero Day Initiative (ZDI)	Channel noted for responsible disclosure of several Apex One issues.
Zoho WorkDrive	Cloud storage used for command retrieval in the Ruby Jumper toolset.