TLP:CLEAR

ADGM THREAT INTELLIGENCE NEWSLETTER

PREPARED BY ADGM CTI

CATEGORY	ACTIONABLE
AUDIENCE	ADGM FSRA ENTITIES
DATE	25/6/2026
OVERALL THREAT SCORE	ELEVATED
TARGET SECTOR	FINANCIAL SERVICES
TARGET REGION	MENA & GLOBAL
ATTRIBUTION	MULTIPLE
TLP	CLEAR

WEEKLY SUMMARY REPORT – 25 June 2026

10	7	0	0
Campaigns	Vulnerability	Cyber Breach	Threat Actors
Threat Campaigns of Potential Relevance to Financial Sector	Actively Exploited & Critical Vulnerabilities	Major Compromises and Breaches	Threat actor activities in the Middle East impacting Financial Sector

Summary

This week's cybersecurity activity showed a broad rise in financially motivated threats, including phishing, malvertising, supply chain compromise, ransomware activity, and attacks on outdated edge devices and enterprise platforms. The campaigns emphasized stealth and scale, using tactics such as credential theft, abuse of trusted services, package poisoning, EDR-disabling tools, covert reconnaissance infrastructure, and evasive malware delivery chains. From a financial sector perspective, these activities may affect remote access, administrative systems, collaboration tools, developer environments, and platforms holding sensitive customer or operational data. The observed vulnerabilities also exposed key technologies to unauthenticated access, privilege escalation, and silent data theft, making rapid patching, reduced internet exposure, stronger MFA, tighter access reviews, and closer monitoring of suspicious installations, uploads, prompts, and persistence activity immediate priorities.

ADGM THREAT INTELLIGENCE SUMMARY

Phantom Stealer Phishing Campaign Targets Banking Organizations [Campaign] [High]

Rokarolla : Android Banker with Complete Device Takeover Capabilities [Campaign] [High]

Prinz Eugen Uses Go-Based Encryption and Anti-Forensic Ransomware Tactics [Campaign] [High]

FortiBleed Credential Harvesting Campaign Targets FortiGate Firewalls Globally [Campaign] [High]

DragonForce Ransomware Group Abuses Microsoft Teams for Covert Command and Control [Campaign] [High]

Gentlemen Ransomware Group Uses Advanced EDR Killer Framework [Campaign] [High]

Sapphire Sleet Targets Mastra npm with Hidden Postinstall Payload [Campaign] [Medium]

AryStinger Botnet Compromises Legacy Routers for Global Proxy Attacks [Campaign] [Medium]

OxLoader Operational Evolution Supporting CASTLESTEALER Infostealer Campaigns [Campaign] [Medium]

EtherRAT Delivered Through Malicious Web Infrastructure and Phishing Pages [Campaign] [Medium]

Cisco SD-WAN Manager File Upload Vulnerability Exploited in the Wild [Vulnerability] [High]

ManageEngine SSO Flaw Could Enable Account Takeover in Integrated Products [Vulnerability] [High]

SearchLeak Chain Exposes Microsoft 365 Copilot Enterprise Search Data [Vulnerability] [High]

Atlassian Security Updates Address Multiple Third-Party Dependency Flaws [Vulnerability] [High]

F5 Out-of-Band Advisory Addresses Critical NGINX Vulnerabilities [Vulnerability] [High]

Oracle Patch Update Fixes Critical Remote Vulnerabilities Across Enterprise Products [Vulnerability] [High]

pgAdmin 4 Patches Three Critical Vulnerabilities [Vulnerability] [Medium]

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Phantom Stealer Phishing Campaign Targets Banking Organizations	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers at Fortra has identified an active phishing campaign targeting high-capital organizations, particularly in the banking sector, through phishing emails carrying RAR (Roshal Archive) files with malicious batch files disguised as business documents. Once executed, the chain uses obfuscated batch and PowerShell stages to decrypt and launch Phantom Stealer entirely in memory.

This campaign uses browser credential theft, screenshot capture, and multi-channel exfiltration, which may impact financial institutions if a compromised employee endpoint provides access to sensitive systems or customer information. Organizations in the financial sector should be aware that the stealer is offered as a maintained MaaS (Malware as a Service) product and may be reused in repeated or parallel campaigns.

Technical Details

- The campaign begins with phishing emails carrying a RAR archive that contains a malicious BAT (Batch) file disguised as a quote request or other business-related document.
- Once executed, the BAT file uses heavy variable noise and command obfuscation, then relaunches itself in a minimized window to reduce user visibility and hinder analysis.
- The script copies a hidden version of itself into 'AppData', applies hidden and system attributes, and creates persistence through a 'RunOnce' registry entry and a fake COM registration.
- A large Base64 blob (Binary Large Object) is decoded into a PowerShell script, which runs with hidden window settings and bypasses execution policy to continue the infection chain without visible prompts.
- The PowerShell stage applies a multi-layer decryption routine consisting of Base64 decoding, two XOR operations, and AES-256-CBC (Cipher Block Chaining) decryption to recover the next-stage payload.
- The decrypted payload is 'DonutLoader' shellcode, which injects Phantom Stealer directly into explorer.exe and keeps the malware entirely in memory, leaving no portable executable on disk.
- Phantom Stealer launches browsers including Chrome, Firefox, and Edge in isolated mode to harvest stored credentials, cookies, autofill data, and session tokens from compromised systems.
- The malware also captures screenshots, clipboard data, keystrokes, and financial or cryptocurrency-related information, then exfiltrates collected data through multiple parallel channels (Telegram, Discord, FTP and SMTP) for redundancy.

Recommendations

- Deploy behavior-based endpoint detection to identify hidden scripting, process injection, and fully in-memory execution.
- Quarantine unsolicited archive attachments containing batch or executable content from unverified external senders.

- Monitor for suspicious process chains involving command shell, hidden PowerShell, and browser launches without user activity.
- Audit browser-stored credentials on potentially exposed endpoints and rotate sensitive credentials where compromise is suspected.
- Enforce multi-factor authentication and reinforce phishing awareness around document-themed attachment lures.

For a detailed breakdown of tactics, techniques, and procedures (TTPs), refer to the Annexure – Phantom Stealer Phishing Campaign Targets Banking Organizations

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
Rokarolla : Android Banker with Complete Device Takeover Capabilities	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers has identified Rokarolla, an Android banking trojan distributed through malicious websites that impersonate legitimate applications such as TikTok or Google Chrome to trick users into installing a dropper and second-stage payload. The malware then abuses Accessibility Services and additional permissions to gain broad control over the infected device and target banking and cryptocurrency applications.

This campaign could affect financial institutions and payment users by enabling theft of banking credentials, lock screen secrets, SMS data, and other sensitive information from compromised mobile devices. Organizations in the financial sector should be aware that the malware combines fraud-focused overlays, device control, and stealth features to support undetected financial abuse.

Technical Details

- The infection starts when a victim is lured to a malicious website and installs a dropper that impersonates a legitimate Android security or update component. The dropper then installs a secondary payload containing the core banking trojan.
- After installation, the malware abuses Accessibility Services and requests additional permissions, including access to SMS messages and notifications. This gives it the visibility and control needed to monitor screens and automate actions on the device.
- Rokarolla communicates with its command-and-control infrastructure over HTTPS and first sends basic device telemetry to identify the victim. It then receives commands and updated configuration data to maintain operations.

- The malware supports multiple fallback domains and can dynamically switch its active command-and-control endpoint. This improves resilience if one communication path is blocked or becomes unavailable.
- Rokarolla uses fake overlays that closely imitate the Android lock screen to capture PINs, patterns, and passwords entered by the victim. This allows attackers to continue operating even when the device is locked.
- It also retrieves targeted banking and cryptocurrency application mappings from the server and downloads fake HTML login pages. When a targeted app is opened, the malware places a fraudulent overlay on top of the real app to steal credentials and payment data.
- Beyond banking theft, the malware can exfiltrate SMS messages, intercept or block calls, log keystrokes, capture visible screen content, and manipulate clipboard data. It also uses repeated screenshots for snapshot-based surveillance.
- To reduce detection, the malware hides its app icon, suppresses audio and vibration, attempts to disable Google Play Protect, and keeps the device screen active to prevent its fraudulent workflows from being interrupted.

Recommendations

- Restrict sideloading and closely monitor mobile devices for unauthorized app installations originating from websites or non-trusted sources.
- Enforce mobile security controls that detect abuse of Accessibility Services, overlay-based phishing behavior, and suspicious permission requests.
- Review and limit SMS, notification, and call-handling permissions on managed devices, especially for applications that imitate system components.
- Educate users to avoid installing updates or security tools from websites and to verify unexpected prompts requesting accessibility or device control privileges.
- Increase monitoring for mobile banking fraud patterns involving stolen OTPs, blocked calls, and suspicious activity originating from enrolled Android devices.

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
Prinz Eugen Uses Go-Based Encryption and Anti-Forensic Ransomware Tactics	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified ‘Prinz Eugen’ a new Go-based ransomware family that was observed following suspected compromise through RDP credentials, after which the actor downloaded an encryptor and

launched it against selected folders. The malware recursively encrypts files, prioritizes recently modified data, and can delete originals after verifying the encrypted output.

This activity could affect organizations in the financial sector where recently modified files, cloud-synced folders, and shared data repositories are critical to daily operations. Organizations in the financial sector should be aware that the actor also used out-of-band extortion, remote management tooling, and anti-forensic cleanup to reduce visible evidence during the attack.

Technical Details

- The intrusion likely began with compromised RDP (Remote Desktop Protocol) credentials, giving the actor interactive access to the environment before ransomware execution.
- After gaining access, the actor downloaded the ransomware executable through Chrome and launched it manually against multiple folders, showing a hands-on deployment approach.
- The encryptor is written in Go and processes each supplied directory through a fully recursive walk, with no depth limit in the analyzed sample.
- Files are selected by most recent modification time first, which means the ransomware targets active and newly updated data before older content.
- For each file, the malware creates an encrypted temporary copy, renames it with the final ransomware extension, and can remove the original when the delete option is enabled.
- Encryption uses 'ChaCha20-Poly1305' with a custom file header, per-file random IVs, SHA-256 integrity hashing, and a multi-stage key derivation process.
- The sample does not contain ransom-note functionality, indicating the operator relies on direct or external communication channels for extortion.
- Before exiting, the malware wipes its hardcoded encryption key from memory, triggers garbage collection, and deletes itself to reduce forensic recovery opportunities.
- In the same environment, the actor also used RemotePC to launch PowerShell stagers and created an additional admin account, suggesting persistence and follow-on access.

Recommendations

- Review and harden RDP access, with close monitoring for unusual or unauthorized interactive logons.
- Restrict and monitor remote management tools and PowerShell activity, especially when used to stage or launch additional payloads.
- Hunt for unexpected local admin account creation and investigate systems where new privileged users were added without approved change activity.
- Deploy behavioral anti-ransomware controls that can detect recursive encryption, file replacement, and self-deletion activity.
- Prioritize backup validation and recovery testing for recently modified and shared data repositories that may be encrypted first.

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
FortiBleed Credential Harvesting Campaign Targets FortiGate Firewalls Globally	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified an active campaign targeting FortiGate firewalls and SSL VPN (Secure Sockets Layer Virtual Private Network) gateways by reusing credentials from previous incidents, applying brute-force methods against weak passwords, and abusing devices without multi-factor authentication. Once access is obtained, compromised devices are used to monitor SSL VPN traffic and harvest additional credentials, creating a self-sustaining credential collection cycle.

This activity may impact organizations in the financial sector that rely on FortiGate appliances for perimeter security and remote access, especially where internet-facing management and legacy password settings remain in place. Organizations in the financial sector should be aware that this campaign is not described as a new vulnerability, but as ongoing credential compromise that could affect firewall administration, VPN access, and potentially wider internal environment.

Technical Details

- The campaign targets internet-facing FortiGate firewalls and SSL VPN gateways using credentials gathered from earlier incidents and infostealer logs, combined with brute-force attempts against weak or unchanged passwords.
- Fortinet stated the activity is not tied to a new vulnerability or recent advisory and linked the compromise to reused credentials and poor password hygiene on devices without MFA (Multi-Factor Authentication).
- SOCRadar described the operation as highly automated and active since at least February 2026, with more than 86,644 compromised devices identified across 194 countries.
- After initial access, attackers use the compromised device as a listening point and monitor SSL VPN traffic to collect more working credentials from passing sessions.
- The campaign feeds newly harvested usernames and passwords back into automated scanning, allowing the intrusion set to expand without manual intervention.
- The exposed attacker data showed broad targeting across sectors, including banks, government entities, telecoms, hospitals, universities, and multinational enterprises.
- Commonly exposed accounts included generic administrator and built-in system accounts, suggesting many devices still used default-style identities or had not rotated credentials after earlier compromise.
- Some activity may also include unapproved configuration changes, unexpected VPN users, password resets, and signs of lateral movement into internal networks, particularly where AD or LDAP integration is enabled.

Recommendations

- Reset all administrator and VPN credentials immediately, especially on internet-facing devices.
- Enforce MFA on all administrator and remote-access accounts without exception.

- Review firewall and VPN configurations against a known good baseline and remove any unrecognized accounts or unauthorized changes.
- Restrict management access using trusted hosts, local-in policies, or by removing internet administration entirely where possible.
- Review logs for unexpected administrator access, unusual VPN usage, suspicious password resets, and possible lateral movement into internal systems.

For a detailed breakdown of tactics, techniques, and procedures (TTPs), refer to the Annexure – FortiBleed Credential Harvesting Campaign Targets FortiGate Firewalls Globally

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
DragonForce Ransomware Group Abuses Microsoft Teams for Covert Command and Control	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified a DragonForce ransomware campaign in which attackers gained access to a victim network, likely through an SQL or MSSQL server weakness or purchased access, then deployed malicious tooling through a ZIP archive containing a sideloaded DLL. The intrusion used a custom Go-based backdoor malware (Backdoor[.]Turn), which abused Microsoft Teams relay infrastructure to hide command-and-control traffic behind legitimate outbound connections.

This campaign could affect organizations in the financial sector by reducing visibility into command-and-control activity, while also supporting credential theft, reconnaissance, and lateral movement inside enterprise environments. Organizations in the financial sector should be aware that the attackers combined stealthy backdoor communications with driver-based defense evasion and ransomware deployment to prolong access and hinder detection.

Technical Details

- The attack appears to begin with access to the victim environment through an SQL or MSSQL server issue, or possibly access obtained from a third party, followed by the download of a malicious ZIP archive.
- The archive contained a legitimate VirtualBox or DbgView executable alongside a malicious DLL, allowing the attackers to use DLL sideloading and run their code through a trusted signed process.
- After execution, the sideloaded DLL downloaded additional code used for securing access, reconnaissance, and evasion, helping the attackers establish a stronger foothold in the environment.

- To maintain resilience, the attackers changed system settings, added users or groups, and modified firewall rules to keep remote access and command-and-control communications available.
- The group also used multiple vulnerable drivers for defense evasion, including a custom technique involving Huawei's HWAuidoOs2Ec[.]sys, along with other vulnerable signed drivers to terminate security processes.
- Backdoor[.]Turn (Malware) was injected into a windows executable for stealth and installed after the ransomware deployment, suggesting it may support persistence, later re-entry, or access resale.
- The backdoor requests an anonymous Teams visitor token, uses a legitimate Microsoft TURN relay for connection setup, and then establishes a QUIC session to the attacker's real command-and-control server.
- Backdoor[.]Turn supports command execution, process creation, network scanning, LDAP or AD discovery, credential-based lateral movement, and browser credential theft from compromised endpoints.

Recommendations

- Monitor for suspicious DLL sideloading involving trusted applications and investigate unexpected use of tooling associated with process injection or archive-based delivery.
- Restrict and review the use of remote administration paths, firewall changes, and unauthorized user or group creation on critical systems.
- Hunt for unusual driver loading activity and assess whether signed but vulnerable drivers are being abused to disable security controls.
- Increase monitoring of Microsoft Teams-related outbound traffic for anomalies, especially where it is followed by unusual process or network activity on internal hosts.
- Strengthen controls around exposed database services, credential hygiene, and privileged access to reduce the chance of intrusion and lateral movement.

For a detailed breakdown of tactics, techniques, and procedures (TTPs), refer to the Annexure - DragonForce Ransomware Group Abuses Microsoft Teams for Covert Command and Control

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
Gentlemen Ransomware Group Uses Advanced EDR Killer Framework	HIGH	CLEAR	Campaign	Open Source

Executive Summary

Researchers at ESET have analyzed the Gentlemen ransomware group's EDR-killing tools made available to affiliates by the operators, helping them weaken or disable endpoint security early in an intrusion. The group uses its own GentleKiller framework along with other acquired tools, all prepared with similar evasion methods and quickly updated to use newly disclosed BYOVD (Bring Your Own Vulnerable Driver) techniques. This helps affiliates bypass defenses more easily before moving on to credential theft, data exfiltration, and ransomware deployment.

This activity could affect organizations in the financial sector by lowering the barrier for affiliates to disable endpoint protections before data theft or ransomware deployment. Organizations in the financial sector should be aware that the group can quickly adapt newly disclosed BYOVD techniques and integrate them into a consistent affiliate-ready toolkit.

Technical Details

- Gentlemen started as a ransomware service in late 2025. It became very active in early 2026 and is still active. The group has been seen in the region and provides ransomware and other support tools to its partners.
- Unlike many ransomware groups, Gentlemen operators actively maintain and distribute a portfolio of EDR killers rather than leaving affiliates to source their own tools.
- The group's in-house framework, named GentleKiller, has at least eight known variants, each using a different vulnerable or malicious driver while keeping a shared internal structure.
- GentleKiller variants share common traits such as repeated process termination loops, code obfuscation, and broad targeting of security-related processes across many products.
- Observed samples were often staged in a directory named GentlemenCollection, where GentleKiller and other EDR killers appeared together during unrelated intrusions.
- Gentlemen also integrates third-party or leaked tools such as HexKiller, ThrottleBlood, and HavocKiller into its broader defense-evasion toolkit.
- The group applies a unified evasion strategy across these tools by using binary protection, filenames resembling legitimate cybersecurity software, and copied icons, signatures, and version information.
- ESET also linked a Rust-based credential stealer, OxideHarvest, to one of Gentlemen's affiliates, showing that credential theft can accompany EDR disruption in the group's operations.

Recommendations

- Monitor for vulnerable or malicious driver abuse, especially where newly disclosed BYOVD proof-of-concepts are rapidly operationalized in real environments.

- Investigate binaries that imitate legitimate software through suspicious filenames, copied icons, or fabricated version information.
- Hunt for repeated attempts to terminate security processes and review endpoint telemetry for unusual process-killing loops.
- Review environments for suspicious staging locations and tool bundles associated with affiliate operations, including credential-stealing utilities.
- Strengthen endpoint and driver control policies to detect or block unsigned, impersonated, or abnormal driver-loading activity.

For a detailed breakdown of tactics, techniques, and procedures (TTPs), refer to the Annexure - Gentlemen Ransomware Group Uses Advanced EDR Killer Framework

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
Sapphire Sleet Targets Mastra npm with Hidden Postinstall Payload	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Researchers at Microsoft has identified a large-scale npm supply chain attack affecting more than 140 packages across the 'mastra' and '@mastra' scopes after the takeover of the 'ehindero' npm maintainer account. The attacker injected a malicious typosquatted dependency 'easy-day-js' into poisoned package versions so that running npm install or npm update would trigger a 'postinstall' hook and execute the payload automatically.

This campaign may impact organizations in the financial sector that rely on developer workstations or CI/CD pipelines to build internal applications, especially where package installation can expose credentials, tokens, or build environments. Organizations in the financial sector should be aware that Microsoft assesses this activity with high confidence to Sapphire Sleet, which primarily targets the financial sector.

Technical Details

- The attack began with the compromise of the 'ehindero' npm maintainer account, which had published rights across the Mastra ecosystem and could push new package versions.
- The attacker created 'easy-day-js', a typosquat of the legitimate 'dayjs' library, and used a staged approach by first publishing a clean version and then a weaponized version.
- More than 140 Mastra packages were then republished with 'easy-day-js' added as a new dependency, and the poisoned versions were tagged as the latest releases.

- When developers or CI/CD systems ran 'npm install' or 'npm update', the dependency resolved to the weaponized version and executed a malicious 'postinstall' hook automatically.
- The postinstall script used obfuscation, disabled TLS certificate verification, dropped tracking markers, contacted attacker-controlled infrastructure, and fetched a second-stage payload.
- The second-stage payload was saved as a randomly named JavaScript file and launched as a detached hidden Node.js process for continued execution.
- On Windows, the payload also performed host reconnaissance, loaded a '.NET DLL' directly into memory through reflection, and used fileless execution to avoid leaving disk artifacts.
- The implant established persistence across Windows, macOS, and Linux, collected browser history, wallet extension information, host details, and process data, and later enabled delivery of a PowerShell backdoor and service-level persistence.

Recommendations

- Review dependency trees for direct or transitive use of the affected Mastra package versions in developer and CI/CD environments.
- Check for 'easy-day-js' in 'node_modules' and 'package-lock[.]json' files across impacted systems.
- Use known-good package versions and run npm 'Install --ignore-scripts' where appropriate to prevent automatic execution of postinstall hooks.
- Investigate systems for unexpected marker files, unusual JavaScript artifacts, and suspicious 'postinstall' activity.
- Rotate credentials, tokens, and API keys that may have been exposed on systems where the compromised packages were installed.

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
AryStinger Botnet Compromises Legacy Routers for Global Proxy Attacks	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Researchers has identified an attack campaign targeting legacy router devices based on RTL819X-series chips by exploiting long-disclosed vulnerabilities to deploy AryStinger malware. The campaign also expanded to NAS (Network-Attached Storage) devices, where a more feature-rich variant was observed delivering scanning, command execution, and tunneling capabilities.

This campaign may impact organizations in the financial sector where outdated edge devices or NAS systems remain exposed, as infected systems can be used for reconnaissance, relay activity, and follow-on intrusion

support. Organizations in the financial sector should be aware that the malware is designed to support pre-intrusion footprinting, persistence, and concealed attacker access.

Technical Details

- The campaign was observed exploiting CVE-2013-3307 and CVE-2016-5681 to compromise old Linksys and D-Link routers, and later CVE-2025-11837 to target NAS devices.
- AryStinger exists in two main versions: a C-based RTL819X variant for older routers and a Go-based Standard variant for NAS devices with broader functionality.
- After infection, the malware authenticates to its command-and-control server, sends device fingerprint data, and receives an assigned Executor ID used in later communications.
- The malware communicates over HTTP or HTTPS and sends data in a compact binary format called Protobuf. This data is protected with basic XOR encryption, and the Standard version also compresses it using Gzip.
- The router-focused variant downloads and deploys 'dropbear' and 'dropbearkey', opens a local SSH service on a specified port, and adjusts firewall rules to maintain remote access.
- AryStinger receives scanning tasks from the command-and-control server and distributes them across infected nodes, allowing large reconnaissance jobs to run in parallel.
- Observed tasking includes domain scanning, IP scanning, HTTP probing, DNS scanning, source-level payload execution in Go, Java, and Python, and system command execution.
- The Standard variant also supports intranet reconnaissance through integrated tools and can return collected scan results to the command-and-control server for follow-on activity.

Recommendations

- Identify and replace or isolate outdated routers and NAS devices that no longer receive security updates.
- Review exposed edge devices for unexpected SSH services, suspicious processes, and unauthorized binaries.
- Monitor for unusual outbound communication with suspicious command-and-control infrastructure and abnormal scanning behavior.
- Inspect network devices and attached systems for signs of persistence, remote management tooling, and unauthorized task execution.
- Prioritize hardening and asset reviews for legacy internet-facing infrastructure that may support reconnaissance or relay activity.

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
OxLoader Operational Evolution Supporting CASTLESTEALER Infostealer Campaigns	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Researchers has identified an active malvertising campaign in which victims searching for Node[.]js are redirected through a fake landing page and a Storj-hosted batch script that downloads and launches OXLOADER. The loader then unpacks itself through multiple decryption stages and delivers the CASTLESTEALER infostealer entirely in memory.

This activity could affect organizations in the financial sector if developer or employee endpoints are exposed through sponsored search results and unauthorized software downloads. Organizations in the financial sector should be aware that the loader's low detection rates, sandbox evasion, and stealthy staging may allow credential theft activity to proceed with limited early visibility.

Technical Details

- The campaign used malicious Google Ads impersonating Node[.]js, directing victims to a fake landing page and then to a Storj-hosted batch script that initiated the infection chain.
- The batch script displayed a fake software installation wizard, downloaded the next-stage executable from Storj, and launched it with elevated privileges through a UAC prompt.
- OXLOADER masqueraded as legitimate software, including API Monitor and a Node.js-themed installer, helping the payload blend in with expected user activity.
- The loader began execution during CRT initialization, before normal user code, and used self-modifying decryption stubs to unpack its next stages at runtime.
- OXLOADER applied several obfuscation methods, including control-flow flattening, opaque predicates, mixed Boolean-Arithmetic, and fragmented function layout to hinder static analysis.
- The malware performed environment checks before continuing, including CPU, RAM, display refresh rate, geographic region, and Russian-language exclusions, to avoid sandbox and analyst systems.
- After passing these checks, it copied a legitimate Windows DLL, created a new executable section, and used the Windows '.reloc' section to stage malicious shellcode.
- The final stage used DonutLoader to execute CASTLESTEALER in memory, allowing the infostealer to run without relying on a traditional disk-based payload chain.

Recommendations

- Block or closely monitor sponsored search-driven software downloads, especially when they imitate developer tools or popular platforms.
- Restrict execution of newly downloaded scripts and installers from user-writable or temporary locations.
- Monitor for unusual PowerShell activity, suspicious elevated process launches, and copied system DLLs loaded from unexpected paths.

- Hunt for behaviors associated with in-memory '.NET' execution, reflective code loading, and self-modifying loader activity.
- Reinforce user awareness around fake software update pages and verify all software downloads through trusted vendor-controlled sources.

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
EtherRAT Delivered Through Malicious Web Infrastructure and Phishing Pages	MEDIUM	CLEAR	Campaign	Open Source

Executive Summary

Researchers have identified an active campaign distributing EtherRAT through websites hosting open directories, MSI installers, and PowerShell loaders, with some infection chains also beginning from phishing emails carrying document lures. The malware deploys a Node[.]js-based RAT that can execute arbitrary code from its command-and-control server and uses the Ethereum blockchain to retrieve active infrastructure.

This activity may impact organizations in the financial sector if endpoints interact with malicious installers, phishing pages, or follow-on remote access content hosted across the same infrastructure. Organizations in the financial sector should be aware that the campaign combines malware delivery, phishing, and covert remote-control mechanisms across a broad and flexible malicious distribution network.

Technical Details

- The campaign was first identified through an open directory serving MSI installers and PowerShell scripts from an /install path, with versioned files such as v1 through v10 used to distribute EtherRAT.
- The MSI-based chain contains a batch launcher, a first-stage JavaScript loader, and an encrypted EtherRAT payload. The batch file extracts components into a random "%LOCALAPPDATA%" path and re-executes itself through conhost[.]exe in headless mode.
- If Node[.]js is not already present, the malware downloads it from the official website, extracts it locally, and uses it to execute the loader and final payload. This helps the malware run consistently on systems without a prior Node[.]js installation.
- The JavaScript loader decrypts embedded code with a custom routine, copies node.exe to a renamed file, adds a persistence-related registry key, and then decrypts and launches the final EtherRAT stage through standard input.
- EtherRAT can execute arbitrary JavaScript from the command-and-control server, manipulate files and registry data, exfiltrate information, and re-obfuscate itself so that each execution generates a different file hash.

- Instead of relying only on fixed infrastructure, the malware uses Ethereum JSON-RPC 'eth_call' requests to retrieve the active command-and-control server, making takedown more difficult.
- Related infrastructure also hosted phishing pages, malicious documents, remote desktop software, and URL-cloaking content, indicating a broader ecosystem used for multiple delivery chains.

Recommendations

- Review email-driven document lures that redirect users to external links, especially when they lead to download pages or credential prompts.
- Investigate endpoints for unexpected Node[.]js downloads, renamed local Node executables, and conhost[.]exe running in headless mode.
- Hunt for persistence through suspicious registry entries linked to script or Node.js execution from user-local directories.
- Monitor for connections to blockchain RPC services followed by unusual, randomized API-style polling patterns from the same host.
- Prioritize review of systems that accessed installer pages, open directories, phishing endpoints, or suspicious remote software downloads associated with the 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
Cisco SD-WAN Manager File Upload Vulnerability Exploited in the Wild	HIGH	CLEAR	Vulnerability	CSC

Executive Summary

Cisco has disclosed CVE-2026-20262, a vulnerability in Cisco Catalyst SD-WAN Manager that is being actively exploited in the wild. The flaw affects the web-based management interface and allows an authenticated attacker, including a low-privileged user, to send crafted HTTP requests to vulnerable file upload functions and create or overwrite files on the underlying operating system.

This activity could affect organizations in the financial sector that rely on SD-WAN Manager for centralized network administration, particularly where the platform is internet-exposed or broadly accessible to users. Organizations in the financial sector should be aware that Cisco has stated the issue may be leveraged to gain root-level privileges, which could lead to full compromise of the management platform.

Technical Details

- CVE-2026-20262 affects Cisco Catalyst SD-WAN Manager due to insufficient validation of user-supplied input during file upload operations in the web-based management interface.

- An authenticated remote attacker can exploit the flaw by sending crafted HTTP requests to vulnerable API endpoints exposed by the management platform.
- The vulnerability does not require high privileges, as Cisco stated that valid credentials, including those associated with low-privileged accounts, are sufficient to trigger the issue.
- Successful exploitation allows the attacker to create or overwrite arbitrary files on the underlying operating system hosting the SD-WAN Manager instance.
- Cisco indicated that the uploaded or overwritten files may then be used to obtain root-level privileges on the affected system.
- The issue has been confirmed as under limited exploitation in the wild as of June 2026, increasing the urgency of remediation.
- Affected deployments include on-premises installations as well as Cisco SD-WAN Cloud-Pro, Cisco SD-WAN Cloud, and Cisco SD-WAN for Government environments.

Recommendations

- Upgrade Cisco Catalyst SD-WAN Manager to a fixed release immediately.
- Restrict internet exposure of SD-WAN Manager instances wherever possible.
- Review user accounts and remove unnecessary or inactive accounts.
- Monitor logs for signs of suspicious file uploads and unauthorized access.
- Investigate for the presence of "suspicious[.]war" and related malicious activity.
- Investigate systems for indicators of compromise and engage Cisco TAC if compromise is suspected.

Vulnerability and affected product details can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
ManageEngine SSO Flaw Could Enable Account Takeover in Integrated Products	**HIGH**	**CLEAR**	**Vulnerability**	**CSC**

Executive Summary

ManageEngine has disclosed CVE-2026-11374, a high-severity vulnerability affecting ADSelfService Plus, RecoveryManager Plus, M365 Manager Plus, and ADAudit Plus when they are deployed as integrated components within AD360. The issue stems from predictable SSO ticket generation, which may allow an unauthenticated remote attacker to predict valid authentication tickets, obtain user identity and role details, and take over targeted account.

This vulnerability could affect organizations in the financial sector that use AD360-integrated ManageEngine products for identity, audit, recovery, or Microsoft 365 administration. Organizations in the financial sector

should be aware that successful exploitation may provide access to privileged user sessions without requiring prior authentication.

Technical Details

- CVE-2026-11374 is a high-severity vulnerability tied to predictable SSO (Single Sign-On) authentication ticket generation in AD360-integrated ManageEngine products.
- The flaw affects deployments where ADSelfService Plus, RecoveryManager Plus, M365 Manager Plus, or ADAudit Plus are integrated with ManageEngine AD360 using SSO.
- The attack vector is network-based and does not require authentication, allowing a remote attacker to target exposed vulnerable integrations directly.
- An attacker may be able to predict valid SSO tickets generated by the affected environment and use them to impersonate legitimate users.
- Successful prediction of these tickets may also expose user identity and role information, which can help the attacker select valuable accounts.
- The reported impact includes account takeover of targeted users within the affected integrated products.
- Affected versions include ADSelfService Plus 6528 and earlier, RecoveryManager Plus 6320 and earlier, M365 Manager Plus 4816 and earlier, and ADAudit Plus 8702 and earlier.

Recommendations

- Upgrade all affected AD360-integrated ManageEngine products to the vendor's fixed versions immediately.

Vulnerability and affected product details can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
SearchLeak Chain Exposes Microsoft 365 Copilot Enterprise Search Data	HIGH	CLEAR	Vulnerability	Open Source

Executive Summary

Microsoft has remediated CVE-2026-42824, a critical vulnerability chain affecting Microsoft 365 Copilot Enterprise Search that can turn a trusted search link into a silent data exfiltration path. The attack begins when a victim clicks a crafted Microsoft 365 Copilot Search URL whose q parameter is interpreted as prompt instructions, causing Copilot to search accessible emails, calendar items, SharePoint content, and OneDrive data.

This vulnerability could affect organizations in the financial sector where users rely on Copilot Enterprise Search to access sensitive business content across Microsoft 365. Organizations in the financial sector

should be aware that the chain enables data exposure with only one click on a trusted domain link, without requiring plugins, special permissions, or a second user action.

Technical Details

- The chain, named SearchLeak, combines three weaknesses: Parameter-to-Prompt Injection, an HTML rendering race condition, and a Bing-based SSRF (Server-Side Request Forgery) path.
- The issue starts with the q parameter in Microsoft 365 Copilot Enterprise Search, which is passed directly to Copilot and treated as executable prompt instructions instead of only a search query.
- An attacker can use this behavior to instruct Copilot to search the victim's accessible Microsoft 365 content and place extracted data into an HTML image tag.
- During response generation, Copilot streams output before its safety wrapper is fully applied, allowing raw HTML such as an 'img' tag to be rendered temporarily in the browser.
- The browser processes the image request before the response is later wrapped in code formatting, creating a race condition that allows exfiltration to begin before sanitization completes.
- Direct exfiltration to an attacker-controlled domain is limited by Content Security Policy, but bing.com is allowlisted for image loading.
- The chain abuses Bing's image search endpoint, which performs a server-side fetch of a user-supplied image URL, allowing Bing infrastructure to request an attacker-controlled path containing stolen data.
- Because Copilot Enterprise Search operates with the victim user's graph permissions, the exposed data can include emails, meeting content, SharePoint documents, OneDrive files, and other indexed organizational content.

Recommendations

- Apply Microsoft's remediation for CVE-2026-42824 without delay.
- Monitor for suspicious Copilot Search URLs containing long or encoded 'q' parameter content, especially where HTML tags or image-based instructions appear.

Vulnerability and affected product details can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Atlassian Security Updates Address Multiple Third-Party Dependency Flaws	**HIGH**	**CLEAR**	**Vulnerability**	**CSC**

Executive Summary

Atlassian has released its June 2026 security updates to address multiple vulnerabilities affecting Bamboo, Bitbucket, Confluence, Crowd, Fisheye/Crucible, Jira Software, and Jira Service Management. The issues originate primarily from third-party dependencies and may be exploited over the network to trigger conditions

such as SSRF, injection, authentication bypass, HTTP request smuggling, information disclosure, remote code execution, and denial of service.

These vulnerabilities could affect organizations in the financial sector that rely on Atlassian platforms for development, collaboration, and service management workflows. Organizations in the financial sector should be aware that several of the disclosed issues carry critical severity ratings and impact core enterprise products, particularly Data Center deployments.

Technical Details

- Atlassian's June 2026 updates address multiple vulnerabilities across Bamboo, Bitbucket, Confluence, Crowd, Fisheye/Crucible, Jira Software, and Jira Service Management. Most of the issues are linked to third-party software components embedded within these products.
- The most severe issues include critical vulnerabilities in axios, Apache Tomcat, Spring Security, and Netty dependencies. These flaws can introduce SSRF, prototype pollution, injection, broken authentication, improper authorization, and HTTP request smuggling conditions.
- Jira Software and Jira Service Management Data Center are affected by several critical axios-related issues, including CVE-2026-42043, CVE-2026-40175, and CVE-2026-42264. These may expose systems to SSRF, prototype pollution, and injection risks.
- Confluence and Jira Software Data Center are affected by critical Tomcat dependency issues such as CVE-2026-41293 and CVE-2026-43512. These issues may enable injection or authentication bypass through vulnerable application components.
- Crowd Data Center is affected by critical Netty and Spring Security-related issues, including CVE-2026-42581, CVE-2026-42584, and CVE-2026-22732. These introduce request smuggling and business logic weakness exposure.
- Bamboo Data Center includes a high-severity ActiveMQ-related remote code execution issue, CVE-2026-41044, along with several axios-related SSRF, information disclosure, and denial-of-service issues.
- Additional high-severity issues affect multiple Atlassian products through minimatch, react-router, Netty, Tomcat, and other dependencies, increasing the breadth of exposure across enterprise deployments.
- Atlassian has released fixed versions for supported product branches, including updated LTS releases for Bamboo, Bitbucket, Confluence, Jira, and Jira Service Management, as well as a recommended update for Crowd.

Recommendations

- Upgrade affected Atlassian products to the fixed or latest supported versions released by Atlassian.

Vulnerability and affected product details can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
F5 Out-of-Band Advisory Addresses Critical NGINX Vulnerabilities	HIGH	CLEAR	Vulnerability	CSC

Executive Summary

F5 has released an out-of-band advisory for multiple critical vulnerabilities affecting NGINX Open Source, NGINX Plus, NGINX Gateway Fabric, NGINX Instance Manager, NGINX Ingress Controller, and App Protect modules. The issues affect HTTP/3 (QUIC), HTTP/2, and gRPC processing, and may allow remote, unauthenticated attackers to trigger denial-of-service conditions or, under certain configurations, achieve remote code execution.

These vulnerabilities could affect organizations in the financial sector that rely on internet-facing NGINX deployments for application delivery, API access, or edge protection. Organizations in the financial sector should be aware that some affected products do not yet have fixes available, increasing the importance of immediate mitigation and exposure reduction.

Technical Details

- F5's advisory covers multiple vulnerabilities across several NGINX products, including core open source, commercial, gateway, ingress, management, and protection-related components.
- The reported flaws impact protocol handling in HTTP/3 with QUIC, HTTP/2, and gRPC, which broadens exposure across common web and application traffic paths.
- The vulnerabilities may be exploited remotely without authentication, which increases risk for exposed services and externally reachable deployments.
- Reported impact includes denial-of-service across affected products, with the advisory also stating that remote code execution may be possible under certain configurations.
- CVE-2026-42530 affects NGINX Open Source, NGINX Instance Manager, NGINX Gateway Fabric, and NGINX Ingress Controller, with fixes available for Open Source and Gateway Fabric.
- CVE-2026-42055 affects NGINX Plus, NGINX Open Source, NGINX Instance Manager, F5 WAF for NGINX, NGINX App Protect WAF, F5 DoS for NGINX, NGINX App Protect DoS, NGINX Gateway Fabric, and NGINX Ingress Controller. Fixes are available only for NGINX Plus and NGINX Open Source.
- CVE-2026-11311 and CVE-2026-50107 affect NGINX Gateway Fabric, and both are fixed in version 2.6.4.
- Several affected products, including Instance Manager, Ingress Controller, and multiple protection modules, currently have no fixes listed in the advisory.

Recommendations

- Upgrade all affected NGINX products to the latest fixed versions immediately where patches are available.
- Disable HTTP/3 and QUIC where they are not operationally required until remediation is completed.
- Limit HTTP/2 and gRPC exposure, especially on internet-facing services.

Vulnerability and affected product details can be found *here*, *here*, *here and here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
Oracle Patch Update Fixes Critical Remote Vulnerabilities Across Enterprise Products	**HIGH**	**CLEAR**	**Vulnerability**	**CSC**

Executive Summary

Oracle has released its June 2026 Critical Security Patch Update addressing 245 vulnerabilities across multiple product families, including Fusion Middleware, WebLogic Server, Enterprise Manager, E-Business Suite, MySQL, JD Edwards, PeopleSoft, Solaris, and Coherence. Several of the flaws are remotely exploitable without authentication and may allow attackers to execute code, gain unauthorized access, disclose sensitive information, modify data, or fully compromise affected systems.

This update could affect organizations in the financial sector that rely on Oracle platforms for middleware, application hosting, identity services, or business operations. Organizations in the financial sector should be aware that multiple vulnerabilities carry critical severity ratings, including CVSS 10.0, and impact widely deployed enterprise components.

Technical Details

- Oracle's June 2026 Critical Security Patch Update addresses 245 vulnerabilities across a broad set of enterprise products, increasing the need for coordinated remediation across multiple teams.
- Several of the disclosed issues are remotely exploitable without authentication, which means attackers may be able to target exposed systems directly over the network.
- Oracle Coherence is affected by critical vulnerabilities including CVE-2026-35308 and CVE-2026-35307, both rated CVSS 10.0 and described as potentially enabling complete compromise.
- Oracle WebCenter Enterprise Capture includes CVE-2026-46778 and CVE-2026-46781, both critical issues that may allow unauthenticated attackers to gain full control of vulnerable instances.
- Oracle WebCenter Portal is affected by CVE-2026-46803 and CVE-2026-46846, critical remote vulnerabilities that may result in unauthorized access and broader system compromise.
- Oracle WebCenter Sites includes critical CVE-2026-46798 and CVE-2026-46800, which could allow unauthenticated remote attackers to compromise affected deployments.
- Oracle WebLogic Server is affected by CVE-2026-35301 and CVE-2026-35292, both rated CVSS 10.0 and described as remotely exploitable flaws that may lead to complete server compromise.
- Additional CVSS 9.8 vulnerabilities affect products including Oracle Enterprise Manager Base Platform, Oracle Unified Directory, Oracle Virtual Directory, Oracle WebCenter Content, WebCenter Content: Imaging, JD Edwards EnterpriseOne Tools, MySQL Router, and Oracle Coherence.

Recommendations

- Apply Oracle's June 2026 security updates across all affected product families as a priority.

Vulnerability and affected product details can be found *here*.

back to top

Name	Threat Severity Rating	TLP	Attribution	Originating Source
pgAdmin 4 Patches Three Critical Vulnerabilities	MEDIUM	CLEAR	Vulnerability	CSC

Executive Summary

pgAdmin has patched three critical vulnerabilities in pgAdmin 4 that affect the SQL Editor, AI Assistant, and HTML rendering paths. The issues include unauthenticated pickle deserialization in connection-related routes, AI Assistant prompt injection that can lead to SQL execution and read-only transaction bypass, and stored XSS through untrusted error and plan-node text.

These vulnerabilities may affect organizations in the financial sector that use pgAdmin for PostgreSQL administration, especially where administrative interfaces are broadly accessible or exposed to shared environments. Organizations in the financial sector should be aware that the issues could enable remote code execution, unauthorized SQL activity, authentication-related abuse, and client-side compromise through stored content.

Technical Details

- CVE-2026-12046 is a critical vulnerability with a CVSS v4 score of 9.5 affecting the SQL Editor close and update_connection routes. It involves unauthenticated pickle deserialization and may allow remote code execution.
- This issue is tied to insecure handling of serialized data during SQL Editor workflow operations. Because the flaw is unauthenticated, an attacker may not need valid credentials to trigger the vulnerable behavior.
- CVE-2026-12045 is a critical vulnerability with a CVSS v4 score of 9.4 affecting the AI Assistant feature. It allows prompt injection and SQL injection and can also bypass read-only transaction restrictions.
- This means attacker-controlled input may influence AI Assistant behavior in a way that results in unintended SQL execution.
- CVE-2026-12048 is a critical vulnerability with a CVSS v4 score of 9.3 involving stored XSS. The issue stems from untrusted error messages and plan-node text being rendered through html-react-parser.
- If malicious content is stored and later rendered in the interface, it may execute in the browser of another user who views the affected content.

Recommendations

- Organizations using pgAdmin should upgrade to fixed version immediately, restrict administrative access, secure deployment configurations.

Vulnerability and affected product details can be found *here*, *here* and *here*.

back to top

Appendix A - Tactics, Techniques & Procedures (TTPs)

Phantom Stealer Phishing Campaign Targets Banking Organizations

Tactic	Technique	ID
Execution	Command and Scripting Interpreter: Windows Command Shell	T1059.003
Execution	Command and Scripting Interpreter: PowerShell	T1059.001
Defense Evasion	Obfuscated Files or Information: Command Obfuscation	T1027.010
Defense Evasion	Obfuscated Files or Information: Base64 Encoding	T1027
Defense Evasion	Process Injection: Shellcode Injection	T1055.004
Defense Evasion	Modify Registry	T1112
Defense Evasion	Hide Artifacts: Hidden Files and Directories	T1564.001
Defense Evasion	Execution Guardrails: Environmental Keying	T1480
Persistence	Boot or Logon Autostart Execution: Registry Run Keys	T1547.001
Persistence	Hijack Execution Flow: COM Hijacking	T1546.015
Discovery	System Information Discovery	T1082
Discovery	System Network Configuration Discovery (External IP)	T1016.001
Collection	Credentials from Web Browsers	T1555.003
Collection	Screen Capture	T1113
Exfiltration	Exfiltration Over C2 Channel (Telegram/Discord/FTP)	T1041
Exfiltration	Exfiltration Over Alternative Protocol: SMTP	T1048.003

FortiBleed Credential Harvesting Campaign Targets FortiGate Firewalls Globally

Tactics	Techniques	Observed Activity
Reconnaissance	T1595.001 Active Scanning: Scanning IP Blocks	The actor conducted large-scale scanning to discover exposed remote-access, firewall, and database services across broad internet ranges.
Reconnaissance	T1596.005 Search Open Technical Databases	The actor used open technical databases to enumerate and enrich FortiGate targets with service, certificate, hostname, and organization-related metadata.
Initial Access	T1190 Exploit Public-Facing Application	The actor targeted exposed FortiGate and SSL-VPN portals through brute-force and credential-stuffing activity to gain access.
Credential Access	T1110.001 Brute Force: Password Guessing	The actor used automated password-guessing activity against FortiGate, database, and storage services using curated credential sets.
Credential Access	T1110.002 Brute Force: Password Cracking	The actor cracked harvested authentication material using distributed password-cracking infrastructure and automation.
Credential Access	T1110.003 Brute Force: Password Spraying	The actor performed password-spraying activity against domain accounts to validate credentials and expand access.
Credential Access	T1110.004 Brute Force: Credential Stuffing	The actor reused previously collected or purchased credentials against SSL-VPN portals and other exposed services.
Initial Access	T1078.001 Valid Accounts: Default Accounts	The actor used weak, default, or reused administrative credentials to access FortiGate and other exposed systems.

Initial Access	T1133 External Remote Services	The actor used compromised remote-access services to enter victim environments, including access to a defense contractor network.
Execution	T1059.004 Command and Scripting Interpreter: Unix Shell	The actor executed FortiOS command-line functionality over remote administrative access to enable packet capture activity.
Credential Access	T1040 Network Sniffing	The actor abused FortiGate diagnostic packet-capture capability to passively collect authentication traffic across multiple protocols.
Credential Access	T1557 Adversary-in-the-Middle	Compromised FortiGate devices acted as traffic interception points, allowing the actor to observe authentication flows traversing the gateway.
Credential Access	T1558.003 Steal or Forge Kerberos Tickets: Kerberoasting	The actor captured Kerberos service-ticket material from network traffic and cracked it offline for credential recovery.
Credential Access	T1558.004 Steal or Forge Kerberos Tickets: AS-REP Roasting	The actor captured AS-REP authentication material and used offline cracking to recover reusable credentials.
Discovery	T1087.002 Account Discovery: Domain Account	The actor enumerated domain users, computer accounts, machine accounts, and email-related identity data.
Credential Access	T1539 Steal Web Session Cookie	The actor reused active web session cookies to maintain authenticated access to compromised environments.
Discovery	T1046 Network Service Discovery	The actor performed network service discovery to identify exposed remote-access, database, and internal services for follow-on activity.
Collection	T1039 Data from Network Shared Drive	The actor recursively collected data from network shares, focusing on user and application data while avoiding standard system directories.
Command and Control	T1071.001 Application Layer Protocol: Web Protocols	The actor used web-based communication channels for sniffer management, dashboard access, and result aggregation.

DragonForce Ransomware Group Abuses Microsoft Teams for Covert Command and Control

Tactics	Techniques	Observed Activity
Initial Access	Exploit Public-Facing Application	Likely exploitation of SQL/MSSQL servers for entry
Execution	DLL Sideloading	Malicious DLL executed via legitimate applications
Persistence	Account Manipulation	Creation of additional user accounts and system configuration changes
Defense Evasion	BYOVD	Use of vulnerable drivers to disable security tools
Defense Evasion	Process Injection	Injection into legitimate processes for stealth
Credential Access	Credential Dumping	Browser credential theft and authentication data collection
Discovery	Network Scanning	Enumeration of network and Active Directory
Lateral Movement	Valid Accounts	Movement using compromised credentials
Command and Control	Application Layer Protocol	Use of Teams TURN relay and QUIC for covert C2
Command and Control	Proxy/Relay Abuse	Leveraging legitimate Microsoft relay infrastructure

Gentlemen Ransomware Group Uses Advanced EDR Killer Framework

Tactics	Techniques	Observed Activity
Execution	T1059.003: Command and Scripting Interpreter: Windows Command Shell	GentleKiller and related tools are console-based executables that run visibly and emit debug strings during execution.
Execution	T1106: Native API	User-mode components interact directly with kernel drivers via DeviceIoControl and other native Windows APIs to perform privileged actions.
Persistence	T1543.003: Create or Modify System Process: Windows Service	The EDR killers install and start vulnerable or malicious drivers as services prior to exploitation.
Stealth	T1036: Masquerading	Gentlemen's EDR killers are protected by impersonating legitimate vendors through filenames, version information, icons, and copied digital certificates.
Stealth	T1036.001: Masquerading: Invalid Code Signature	The protection applied to Gentlemen's EDR killers adds an invalid code signature as part of the impersonation strategy.
Stealth	T1027: Obfuscated Files or Information	Some executables are protected with packers (e.g., Enigma, Themida) and custom control-flow obfuscation.
Defense Impairment	T1685: Disable or Modify Tools	GentleKiller and other EDR killers that Gentlemen is in possession of aim to bypass security products such as EDRs.

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 eyes and ears of individual recipients only, no further.	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 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
Account Takeover	When an attacker gains control of a legitimate user account and can act as that user inside systems or applications.
Active Directory (AD)	A Microsoft service used by organizations to manage users, systems, and access rights across the enterprise.
ActiveMQ	A software component referenced in the Atlassian updates that can introduce security risk if vulnerable.
Adler-32	A checksum-style algorithm referenced in the OXLOADER analysis for API resolving.
Admin Account	A user account with elevated privileges that can change settings, add users, or manage systems.
AI Assistant	A built-in assistant feature that uses artificial intelligence to interpret prompts, retrieve information, or perform actions.
AI Assistant Prompt Injection	A manipulation technique where attacker-controlled input causes an AI assistant to follow harmful or unintended instructions.
Anti-Forensic	Any technique designed to hide attacker activity or reduce evidence left on disk or in memory.
API	Application Programming Interface; a software interface that lets systems or applications exchange data or perform actions.
API Endpoint	A specific application path or function that accepts requests and returns data or results.
API Monitor	A legitimate tool name that OXLOADER used as part of its disguise.
App Protect	A protection-related NGINX product line mentioned in the F5 advisory.
App Protect DoS	An NGINX-related component intended to help defend against denial-of-service activity.
App Protect WAF	An NGINX-related web application firewall component referenced in the F5 advisory.
AppServer Logs	Application logs that may help identify suspicious activity such as file uploads or unauthorized access.
Arbitrary Code Execution	The ability for an attacker to make a system run commands or code of the attacker's choice.
Ary-Attack	A project name indicated in source code associated with the AryStinger malware family.
AryStinger	A malware family used to compromise old routers and NAS devices and turn them into reconnaissance and relay infrastructure.
Asset Mapping	The process of identifying exposed devices or systems across the internet or within an environment.
Authentication Bypass	A weakness that lets an attacker access a system or function without completing normal login or identity checks.
Axios	A software dependency referenced in multiple Jira, Bamboo, and npm-related issues in the newsletter.
Backdoor	Hidden or unauthorized access left on a system so an attacker can return later.
Backdoor.Turn	A custom Go-based backdoor used in the DragonForce campaign to conceal command-and-control traffic through Microsoft Teams relays.
BAMBOO	An Atlassian product used for build and deployment workflows that was affected by June 2026 security fixes.
Batch File / BAT	A Windows script file used to automate commands; often abused in malware delivery chains.
Behavior-Based Detection	A security approach that looks for suspicious actions or patterns rather than only known file signatures.
Bing SSRF	A technique in the SearchLeak chain that abused Bing infrastructure to fetch attacker-controlled URLs and assist data exfiltration.
Bitbucket	An Atlassian product included in the June 2026 security updates.

Blockchain	A distributed ledger system; in one campaign it was used by malware to retrieve command-and-control information.
Bot	A compromised system or program that receives instructions from an attacker and performs assigned tasks.
BotID	A unique identifier assigned to an infected device so the attacker can track or manage it.
Browser Credential Theft	Theft of usernames, passwords, cookies, or session data stored in web browsers.
Build ID	A value used by malware or software to identify a specific version or campaign configuration.
BYOVD	Bring Your Own Vulnerable Driver; a method where attackers use legitimate but vulnerable drivers to disable protections or gain stronger access.
C2 / Command and Control	The infrastructure attackers use to communicate with and control compromised systems.
Calendar Data	Meeting-related content such as invites, titles, notes, or schedules that may be accessible in collaboration platforms.
CASTLESTEALER	An infostealer delivered by OXLOADER and designed to collect sensitive data from infected systems.
Catalyst SD-WAN Manager	Cisco's SD-WAN management platform, formerly known as vManage.
ChaCha20-Poly1305	An encryption method used by Prinz Eugen ransomware to encrypt files while keeping integrity checks.
Check Poll Loop	A repeated communication cycle used by malware to check with its command-and-control server for new tasks.
CI/CD	Continuous Integration / Continuous Delivery; automated software build and deployment pipelines that can be exposed in supply chain attacks.
Cisco TAC	Cisco Technical Assistance Center; Cisco's support function referenced for incident response when compromise is suspected.
Cloud-Pro	A cloud-based SD-WAN deployment model referenced in the Cisco vulnerability entry.
CMD	Windows command shell used to run commands and scripts.
Coherence	An Oracle product family affected by critical vulnerabilities in the June 2026 Oracle update.
COM Hijack	A Windows persistence or execution technique that abuses how system components load or call objects.
Compromise Assessment	A review performed to determine whether a system has already been breached or altered by an attacker.
Confluence	An Atlassian collaboration platform included in the June 2026 security updates.
Control-Flow Flattening	An obfuscation technique that makes malicious code harder to analyze by disguising the order of execution.
Copilot Enterprise Search	A Microsoft 365 feature that searches across user-accessible enterprise content such as email and files.
Cross-Site Scripting (XSS)	A web vulnerability that allows malicious script content to run in another user's browser.
Crowd	An Atlassian identity-related product included in the June 2026 updates.
Crowd Data Center	The enterprise deployment model of Crowd referenced in the Atlassian advisory.
CRT Initialization	An early application startup phase that malware can abuse to begin execution before normal program logic.
CSC	UAE Cyber Security Council
CSP / Content Security Policy	A browser security control that restricts which external sources a web page can load content from.
CVE	Common Vulnerabilities and Exposures; the standard identifier used to track publicly disclosed security flaws.

CVSS	Common Vulnerability Scoring System; a severity scoring method used to rate vulnerabilities.
Data Center	An enterprise product deployment model referenced throughout the vulnerability entries, especially for Atlassian and other platforms.
Data Exfiltration	Unauthorized removal or transfer of data from a system to attacker-controlled infrastructure.
Defender Exclusion	A configuration that tells Microsoft Defender to ignore certain files or paths, which attackers may abuse to reduce detection.
Defense Evasion	Actions taken by attackers or malware to avoid being detected or stopped by security controls.
Dependency	A software component or library that another product relies on; many vulnerabilities in the newsletter originated in such dependencies.
Detached Hidden Process	A background process launched in a way that reduces visibility to the user.
D-Link	A router vendor heavily represented among AryStinger-infected devices in the reported campaign.
DLL	Dynamic-Link Library; a Windows component that applications load, which attackers can abuse to run malicious code.
DLL Sideloading	A technique that makes a trusted application load a malicious DLL placed in an expected path.
DNS	Domain Name System; the service that translates domain names into IP addresses and can be abused in scanning or redirection activity.
DonutLoader	A shellcode packaging tool used to deliver payloads for in-memory execution, helping malware avoid writing a standard executable to disk.
DoS	Denial of Service; an attack or weakness that can disrupt or degrade system availability.
DragonForce	A ransomware group that used Microsoft Teams relay infrastructure and driver-based evasion in one of the campaigns.
Dropbear	A lightweight SSH server that AryStinger deploys on compromised routers to maintain remote access.
Edge Device	A network-facing device such as a router, firewall, or gateway that connects an organization to the internet.
EDR	Endpoint Detection and Response; endpoint security technology used to detect and respond to suspicious behavior.
EDR Killer	A tool designed to disable or interfere with endpoint security products before or during an intrusion.
Elastic Defend	A defensive product mentioned in the OXLOADER reporting as detecting behavior in the attack chain.
Enterprise Manager	An Oracle product family affected by the June 2026 Oracle patch update.
Enterprise Search	A search capability that retrieves information across business applications and data sources.
Ethereum	A blockchain platform used by EtherRAT to retrieve command-and-control server details.
EtherRAT	A Node.js-based remote access trojan that executes attacker-supplied code and uses Ethereum infrastructure to discover its C2 server.
Executor	In the AryStinger reporting, the name given to each infected node that performs tasks assigned by the attacker.
FedRAMP	A government cloud environment model referenced in the Cisco SD-WAN Manager vulnerability coverage.
File Upload Vulnerability	A weakness in how an application processes file uploads, which can let attackers create or overwrite files on a system.
Fisheye/Crucible	Atlassian products included in the June 2026 security update coverage.
FortiBleed	A campaign targeting FortiGate devices through credential reuse, brute force, and further credential harvesting.

FortiGate	A firewall and VPN product family referenced in the FortiBleed campaign reporting.
Fusion Middleware	An Oracle product family included in the June 2026 Oracle security update.
Gateway Fabric	An NGINX-related product referenced in the F5 advisory.
GentleKiller	The in-house EDR-killer framework maintained by the Gentlemen ransomware group.
Gentlemen	A ransomware-as-a-service group that provides affiliates with operator-managed tools to disable endpoint protections.
GitHub Actions OIDC	A publishing method referenced in the npm compromise analysis to distinguish normal package publication from suspicious manual publication.
Go / Golang	A programming language used by several malware families and ransomware samples mentioned in the newsletter.
Google Ads Malvertising	Use of sponsored search advertisements to redirect victims to malicious download pages.
Graph Permissions	The access rights that determine what Microsoft 365 content a user or service can retrieve.
Group Policy / Policy Setting	A control or configuration setting used to determine how systems or security tools behave.
gRPC	A modern communications protocol used by applications and services, referenced in the F5 advisory as part of the affected traffic handling.
HavocKiller	A third-party EDR killer incorporated into the Gentlemen toolkit.
Heartbeat	A regular status message sent by malware to its command-and-control server to confirm it is still active.
Hidden Process	A process launched without a visible user interface to reduce the chance of being noticed.
HTTP Request Smuggling	A weakness where differences in request handling between systems can let attackers bypass controls or confuse traffic processing.
HTTP/2	A web communication protocol referenced in the F5 advisory as part of the affected handling paths.
HTTP/3	A newer version of the web protocol that uses QUIC and was referenced in the F5 advisory.
HTTPS	Encrypted web traffic used for legitimate communications and also by malware command-and-control channels.
Hyperlink-Based Delivery	A delivery method where the victim clicks a link that starts the attack chain.
Identity and Role Information	Details about who a user is and what level of access they have, which can help attackers select higher-value targets.
Identity Provider	A system or service that handles user authentication and access to connected applications.
Improper Authorization	A flaw where a system fails to correctly enforce what an authenticated user is allowed to do.
Ingress Controller	A component used to manage access into services or applications, referenced in the F5 advisory.
Injection	A vulnerability type where attacker-controlled input is interpreted as commands or code rather than harmless data.
In-Memory Execution	Running malicious code directly in system memory rather than as a normal file on disk.
Insecure Deserialization	A weakness where serialized input is unsafe to process and can lead to remote code execution or other compromise.
Instance Manager	An NGINX management product referenced in the F5 advisory.
Intranet Scanning	Scanning internal systems and services from within a compromised environment to map targets for later action.
IOC / Indicator of Compromise	A sign that a system may have been attacked, such as a suspicious file name, process, account, or connection.

JAR / Java Artifact	A Java application package or code component that can be loaded or executed in enterprise environments.
JavaScript Loader	A script stage used to decrypt and launch later malware components.
JD Edwards	An Oracle product family affected by the June 2026 Oracle update.
Jira Service Management	An Atlassian service management product included in the June 2026 updates.
Jira Software	An Atlassian product included in the June 2026 updates.
Lateral Movement	The process of moving from one compromised system to other systems inside the same environment.
Loader	Malware whose main role is to unpack, stage, or deliver a second payload.
Low Detection Rate	A situation where malware or infrastructure is not widely recognized by security tools, making early discovery harder.
M365	Microsoft 365; a productivity and collaboration environment referenced in multiple newsletter entries.
M365 Manager Plus	A ManageEngine product affected by the AD360-related SSO vulnerability.
MaaS	Malware-as-a-Service; malware offered or maintained as a reusable service for different users or operators.
Malvertising	Malicious advertising used to redirect users to malware or phishing content through online ads.
Managed Deployment	A service model where a vendor hosts or operates the environment on behalf of the customer.
Mastra	The npm package scope that was compromised in the supply chain attack attributed to Sapphire Sleet.
MFA	Multi-Factor Authentication; a security control that requires more than one step to verify identity during login.
MITRE ATT&CK	A framework used to describe attacker goals, behaviors, and techniques in a consistent structure.
Mixed Boolean-Arithmetic (MBA)	A code-obfuscation technique that makes logic harder to read and analyze.
MSI	Microsoft Installer package; a Windows installation format commonly used for software delivery and sometimes abused for malware.
MSSQL	Microsoft SQL Server; one of the possible entry points mentioned in the DragonForce campaign.
NAS	Network-Attached Storage; a storage device on a network that was targeted by the Standard version of AryStinger.
Netty	A networking component referenced in multiple Atlassian dependency-related vulnerabilities.
NGINX Open Source	The core open source NGINX product referenced in the F5 advisory.
NGINX Plus	The commercial version of NGINX affected by the F5 advisory.
Node[.]js	A JavaScript runtime used for development and server-side execution, and abused in multiple malware chains in the newsletter.
npm	Node Package Manager; the software package ecosystem affected in the Mastra supply chain compromise.
Obfuscation	Techniques used to make malicious code or logic harder to analyze, detect, or understand.
Opaque Predicate	An obfuscation technique that introduces misleading conditions or branches to confuse analysis tools.
Oracle Coherence	An Oracle product family affected by multiple unauthenticated critical vulnerabilities in the June 2026 patch update.
Oracle CSPU	Oracle's Critical Security Patch Update, the vendor's periodic release for security fixes.
OXLOADER	A previously undocumented Windows loader that used layered obfuscation and stealthy staging to deliver CASTLESTEALER.

Parameter-to-Prompt Injection (P2P)	A vulnerability where a normal input field or URL parameter is treated as instructions by an AI system.
PBKDF2	A password hashing approach referenced in the Fortinet guidance for strengthening administrator credential storage.
PeopleSoft	An Oracle product family included in the June 2026 patch update.
Persistence	A method used by malware or attackers to remain present on a system after reboot or logon.
Phantom Stealer	A commercially offered infostealer delivered through phishing and designed to harvest credentials and other sensitive data.
Phishing	A deceptive technique used to trick users into clicking links, opening attachments, installing software, or revealing information.
Phishing Overlay	A fake on-screen login or prompt placed over a trusted application or screen to capture user input.
Port Scanning	A reconnaissance method used to identify open services or reachable systems on a network.
PostgreSQL	A database platform administered through pgAdmin, which was affected by critical vulnerabilities.
Postinstall Hook	A script that runs automatically after a software package is installed and can be abused in supply chain attacks.
PowerShell	A Windows scripting and automation tool often used legitimately by administrators and also abused by attackers.
Predictable Authentication Token	A weakness where tokens or tickets can be guessed or predicted, allowing account misuse.
Privilege Escalation	The process of gaining higher access rights than originally held.
Process Injection	A technique in which malicious code is inserted into another running process for stealth or privilege reasons.
Protobuf	Protocol Buffers; a compact data format used by some malware to communicate with command-and-control infrastructure.
Prototype Pollution	A software flaw that lets attacker-controlled data alter application behavior, particularly in JavaScript-based environments.
Proxy / Forwarding Channel	A relay path that attackers use to pass traffic through compromised systems and hide the real source of activity.
QUIC	A network transport protocol used with HTTP/3 and referenced in the F5 advisory.
RaaS	Ransomware-as-a-Service; a model where operators provide ransomware and supporting services to affiliates.
RAT	Remote Access Trojan; malware that lets an attacker remotely control an infected device.
React Router	A software dependency referenced in the Atlassian vulnerability coverage.
Read-Only Transaction Bypass	A failure to properly enforce read-only restrictions, allowing actions beyond what should be permitted.
Reflective Code Loading	Loading code directly into memory without a normal executable being written and run from disk.
Registry Run Key	A Windows registry setting commonly abused to automatically launch malware when the user logs in.
Remote Code Execution (RCE)	A vulnerability impact where an attacker can run code of their choice on a target system.
Remote Management Tool	Software used to administer systems remotely; attackers often abuse such tools because they resemble normal administration activity.
Reprompt	A previously disclosed AI assistant vulnerability referenced as an earlier example of prompt-based abuse.
Rokarolla	An Android banking trojan that uses overlays, device permissions, and broad control features to steal banking and cryptocurrency-related information.

Root Privileges / Root-Level Access	The highest level of control on a system, allowing broad or total administrative access.
Router Botnet	A group of compromised router devices controlled together and used for scanning, relaying, or attack activity.
Run Key	A Windows startup persistence location that launches a program when a user signs in.
SearchLeak	The name given to the Microsoft 365 Copilot Enterprise Search vulnerability chain that enabled silent data exfiltration.
Security Code / OTP	A one-time security code often used in login, reset, or verification processes and explicitly mentioned as exposed in the SearchLeak scenario.
Semantic Versioning (SemVer)	The versioning system used in software packages that can affect which dependency version gets installed automatically.
Session Token	A value that represents a logged-in user session and can allow continued access if stolen.
SharePoint	A Microsoft platform for storing and sharing documents that was included in the SearchLeak exposure discussion.
Shellcode	Small code designed to run directly in memory, often used in staged malware delivery.
Silent Data Exfiltration	Theft of data in a way that produces little or no visible warning to the user.
Single Sign-On (SSO)	A login method that allows one authentication event to access multiple connected applications.
Smart Contract	Code stored on a blockchain that can be called or queried by applications, including malware in the EtherRAT case.
SSRF	Server-Side Request Forgery; a vulnerability that makes a server request data from an attacker-chosen location.
Stored XSS	Cross-site scripting where malicious content is saved by the application and later executed when viewed by another user.
Storj	A legitimate file-sharing/storage service abused in the OXLOADER delivery chain to host malware-related files.
Supply Chain Attack	A compromise that targets software packages, updates, or dependencies so downstream users are affected indirectly.
Suspicious.war	A suspicious file name specifically listed as an indicator of compromise in the Cisco SD-WAN Manager vulnerability advisory.
Telemetry	Data collected from systems, users, or applications to help monitor behavior or identify malicious activity.
Threat Actor	An individual or group carrying out malicious cyber activity.
ThreatDown	A security vendor referenced in the Prinz Eugen reporting.
ThrottleBlood	A third-party EDR killer integrated into the Gentlemen toolkit.
TLS	Transport Layer Security; the protocol used to secure network communications.
TLS Certificate Verification	The process of confirming that a secure connection is using a trusted certificate; disabling it weakens trust validation.
Tomcat	A web application component referenced in multiple Atlassian vulnerabilities.
TURN Relay	A relay service used in real-time communications that was abused to mask Backdoor.Turn traffic in the DragonForce campaign.
Typosquat	A malicious package or domain that imitates a legitimate name using a closely similar spelling.
UAC	User Account Control; a Windows security feature that prompts when elevated privileges are requested.
Unified Directory	An Oracle directory product affected by high-severity vulnerabilities in the June 2026 patch update.
Unsupported Infrastructure	Older systems or devices that no longer receive regular vendor updates or maintenance.

Varonis Threat Labs	The research team that reported the SearchLeak chain affecting Microsoft 365 Copilot Enterprise Search.
Virtual Machine / VM	A software-based computing environment often used for testing or analysis, and commonly checked by malware for evasion purposes.
Virtualization / Sandbox Evasion	Techniques used by malware to detect whether it is running in a test or analysis environment and stop execution if so.
Vulnerable Driver	A legitimate driver with known weaknesses that attackers can abuse to disable defenses or gain deeper access.
Wealth Management Data	Sensitive business or client information that may be exposed if enterprise collaboration and storage systems are abused.
Web Application Firewall (WAF)	A security control designed to inspect and protect web application traffic.
WebCenter	An Oracle product family affected by critical vulnerabilities in the June 2026 patch update.
WebLogic Server	An Oracle middleware product affected by multiple critical remotely exploitable vulnerabilities.
WMI	Windows Management Instrumentation; a Windows interface often used for system information gathering.
XOR	A basic mathematical operation often used in malware to encode, decode, or obfuscate content.