ICT DIVISION
FUTURE IS HERE

BANGLADESH COMPUTER COUNCIL
COMPUTER FOR EVERYTHING

BGD e-GOV CIRT

TLP:CLEAR

AI-Branded Social Engineering, Phishing, and Malware Delivery Campaigns

TLP: CLEAR
Distribution: Public
Advisory on: AI-Branded Social Engineering, Phishing, and Malware Delivery Campaigns
Severity: High
Threat Type: AI-Branded Social Engineering, Phishing, Adversary-in-the-Middle (AiTM), and Malware Distribution Campaigns
Date: 09 July 2026

Executive Summary

Threat actors are increasingly abusing the branding of popular AI services as social-engineering bait to drive phishing, credential theft, payment fraud, and malware delivery. Microsoft reported campaigns using the names and visual branding of ChatGPT, Claude, DeepSeek, Microsoft Copilot, and other AI tools to make malicious messages, download pages, and repositories appear trustworthy and timely.

The activity spans several delivery models, including phishing emails with urgent account or billing themes, adversary-in-the-middle token theft flows, malvertising chains that push users to fake downloads, and SEO-driven discovery of malicious GitHub repositories. Microsoft assessed that the observed incidents reflected abuse of trusted brand names as lures rather than compromise of the legitimate AI vendors themselves.

This threat is directly relevant to Bangladesh because the observed activity already affected in South Asia, while the lures depend on user curiosity, urgent billing prompts, cloud-hosted redirects, and fake software downloads that are equally plausible against Bangladeshi government staff, students, freelancers, software users, and mobile-first internet users. The campaigns are therefore best understood as transferable social-engineering playbooks that can be quickly localized to Bangla or English and reused against Bangladeshi targets with minimal attacker effort.

Bangladesh Impact

Bangladeshi users are at realistic risk in at least four scenarios.

First, individuals who use AI services for study, coding, translation, graphics, or office productivity may respond to fake subscription, appeal, or access-restoration messages because the lures imitate normal service workflows.

Second, users searching online for “free” or “downloadable” AI tools, browser plugins, cracked software, or model installers may be exposed to malicious GitHub repositories, SEO-manipulated pages, and malvertising chains similar to those described by Microsoft.

Third, government officials, university staff, financial-sector employees, outsourcing professionals, and software developers in Bangladesh may face elevated exposure because

these roles regularly open email attachments, use cloud identities, and rely on remote collaboration platforms where token theft can lead to mailbox compromise, business email compromise, and unauthorized data access.

Fourth, widespread use of personal Windows devices in home, education, and small-business environments increases the likelihood that fake installers or archive files will be executed outside centrally managed enterprise controls.

The campaign's activity in the South Asian region indicates that Bangladeshi users could also be targeted with minimal localization, such as Bengali content, local brands, or region-specific lures delivered through email, messaging applications, or social media. Potential impacts include payment card theft, credential compromise, AiTM-based session hijacking, infostealer infections (e.g., Vidar), account takeover, unauthorized access to government and financial services, data leakage, malware propagation, and reputational damage through the misuse of compromised accounts.

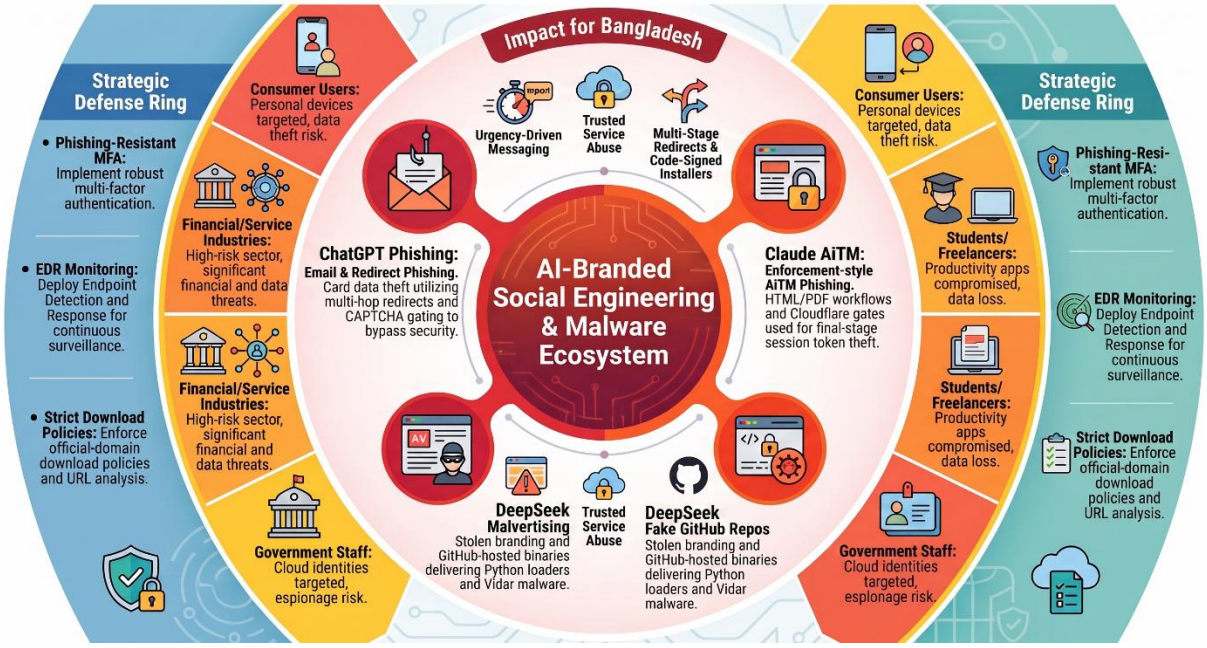

Figure: AI-Branded Social Engineering, Phishing, and Malware Delivery Campaigns Ecosystem

Threat Overview

Microsoft reported that threat actors are operationalizing public interest in AI as a durable social-engineering theme, combining curiosity and urgency with older but effective techniques such as spearphishing links, malicious attachments, redirector chains, and staged credential harvesting. The observed campaigns led to collection of payment card data, theft of credentials and access tokens, and delivery of malware including Vidar Stealer, Lumma Stealer, Hijack Loader, and Oyster. One cluster involved a ChatGPT-themed payment update lure, another abused Claude account-policy enforcement messaging for likely AiTM token theft, a separate campaign used an “Awesome AI Windows Plugin” malvertising lure to deliver

a signed malware chain, and another used a fake DeepSeek V4 GitHub repository to distribute infostealer payloads. Microsoft linked part of the malware-distribution activity to Storm-3075 and described the use of malware-signing services attributed to Fox Tempest to increase trust and reduce early detection rates.

Campaign Details

ChatGPT-themed phishing for card and personal data

On 5 May 2026, Microsoft detected a ChatGPT-themed phishing campaign that sent malicious URLs leading to pages that collected personal and credit card information. Microsoft stated that 4,500 emails in one activity cluster targeted South Africa and that related activity delivered as many as 100,000 emails in a single day across multiple sectors including higher education and professional services.

The phishing email used the sender display name "ChatGPT" and the subject line "**To ensure your ChatGPT Plus continues to work – please update your payment method**". The message warned that the account would be downgraded within seven days if payment details were not updated, a tactic designed to force hurried action before verification.

The redirection chain moved the victim through several legitimate or abused services before reaching the phishing kit: grupoconstat[.]bitrix24[.]com[.]br to awstrack[.]me to a Rebrandly URL and finally to legendarytrendsbay[.]shop/ChatGPT/. The final phishing site first displayed a custom CAPTCHA-like "Update payment" prompt and then collected the victim's first name, last name, address, and credit-card details including expiration date and verification code.

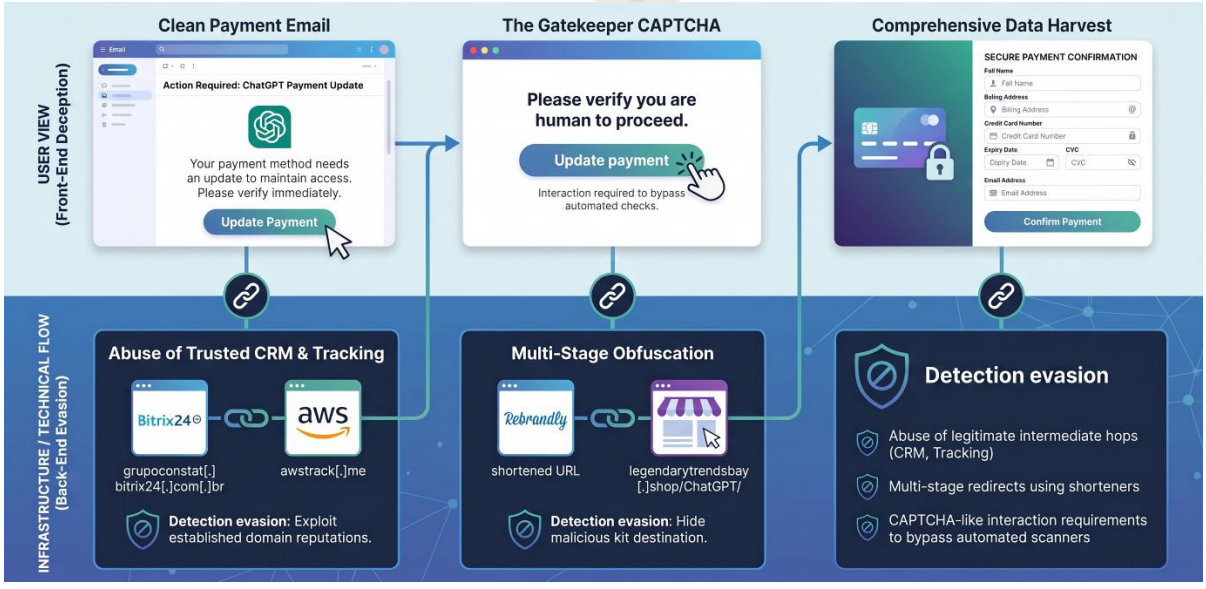

Figure: ChatGPT payment-update phishing chain

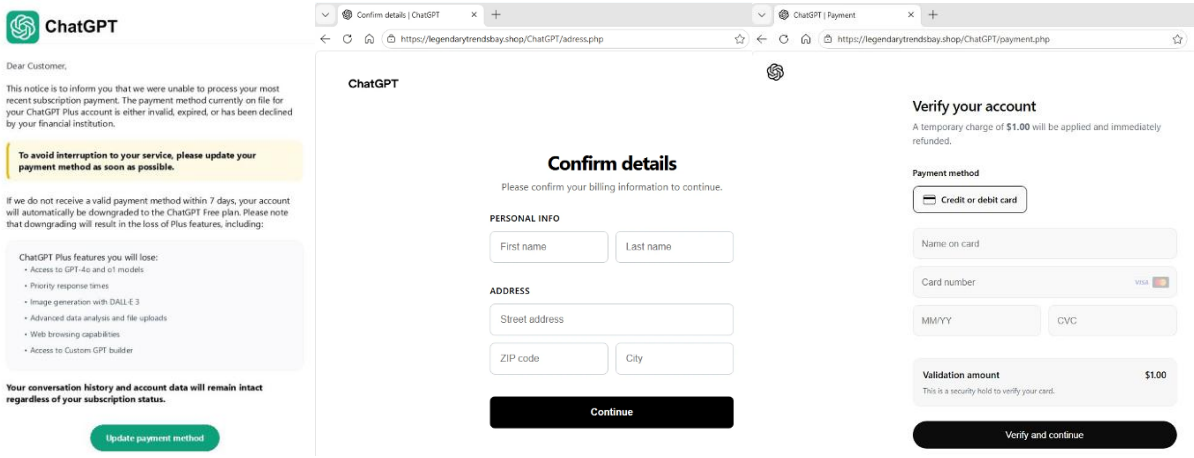

Figure: ChatGPT-themed phishing email and subsequent credential harvesting pages collecting personal identity information and credit card information [Source: Microsoft]

Claude-themed phishing leading to likely AiTM token theft

From 20 to 22 April 2026, Microsoft observed a campaign impersonating Anthropic-branded services and targeting more than 2,000 organizations, primarily in the United States (62%), the United Kingdom (18%), and India (9%). The campaign was especially concentrated in information technology (56%), other business entities (21%), and financial services (8%).
The messages used display names such as "**Anthropic Teams**" and "**Anthropic PBC**" and claimed the recipient's account violated acceptable-use policy and required immediate appeal action. The email body was delivered as HTML, carried Claude branding, and instructed recipients to review an attached PDF named Fill and Sign Claude Appeal Form.pdf.

That PDF told the user to copy an appeal ID and click a "Claude Appeal" link, which led to dash.awaydouble[.]org. The initial landing page displayed a Cloudflare verification prompt, then redirected to servicing.pureplantcravings[.]com, which showed an "**Account Appeal Notice**" and a one-time access code before steering the victim toward a final sign-in experience consistent with adversary-in-the-middle token interception; Microsoft noted the redirect logic differed for mobile and desktop users.

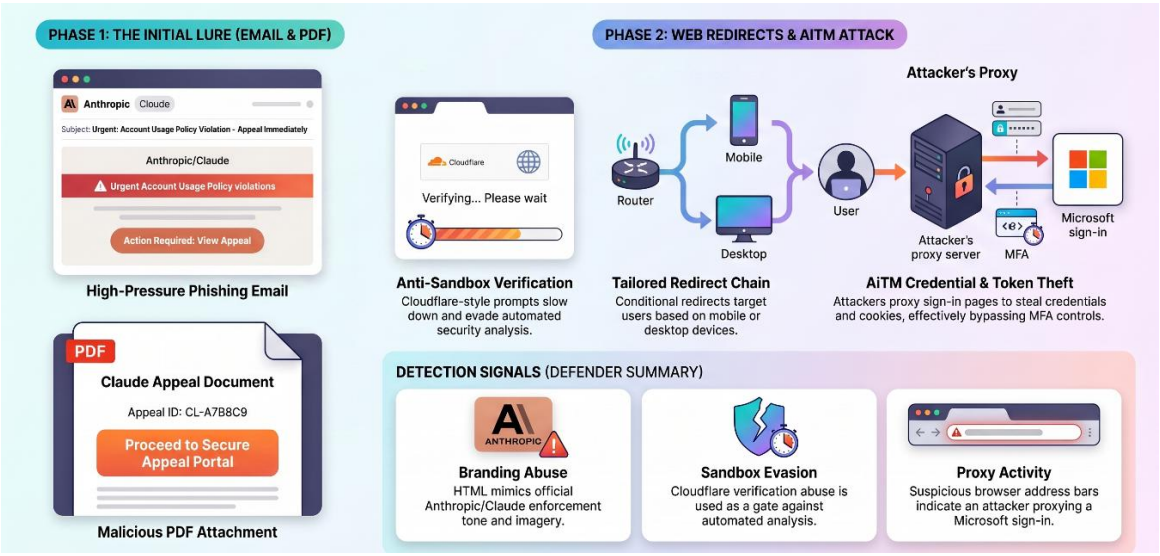

Figure: Claude appeal / AiTM credential and token theft

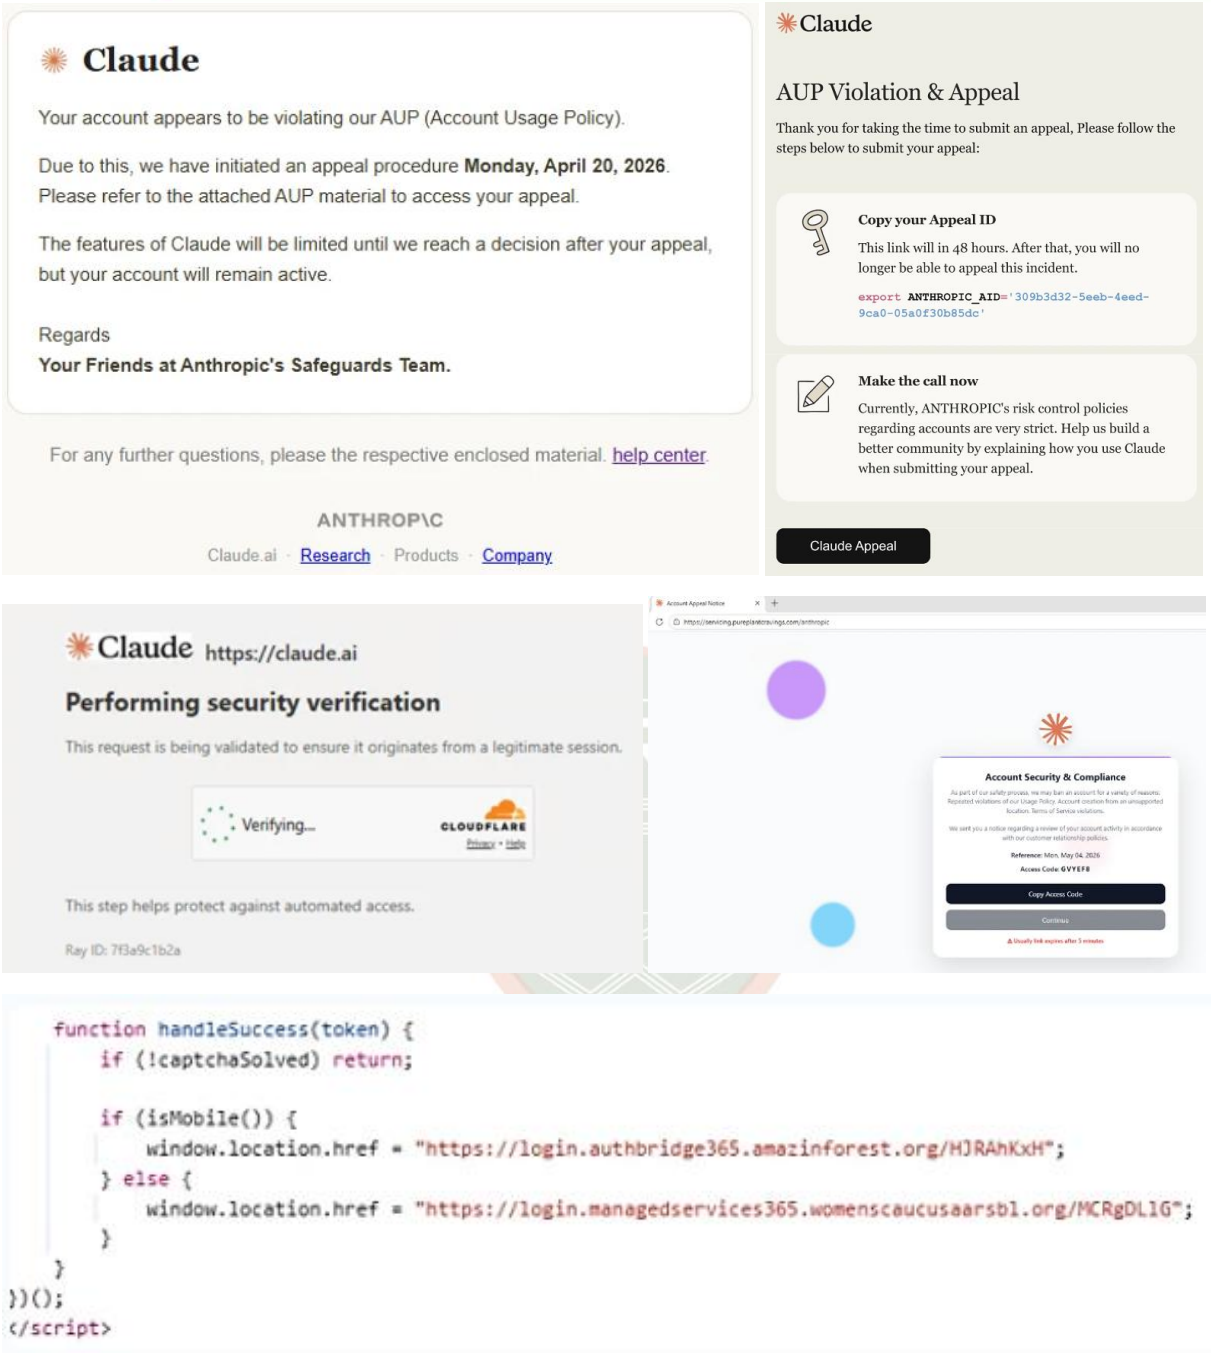

Figure: Claude-themed phishing campaign comprising a phishing email, malicious PDF attachment, CAPTCHA-gated landing page, branded intermediate page with an access code, and source code revealing device-specific redirect logic for mobile and desktop systems [Source: Microsoft]

"Awesome AI Windows Plugin" malvertising

Since at least early 2026, Microsoft observed malvertising campaigns using AI-themed terms such as "**Awesome AI Windows Plugin**" and "**Flux Pro AI**" in popup lures, malware filenames, and GitHub repository paths. Microsoft described these campaigns as moving from launch to mass impact within hours and noted that a single campaign run on 13 March 2026 targeted more than 66,000 devices.

Telemetry showed global distribution, with the top affected countries including Japan, South Africa, the United States, and France, and Microsoft assessed that most impacted devices were likely consumer endpoints because the attack traffic originated from free movie streaming activity. The infection flow began when users clicked popups or interacted with embedded players on streaming sites and were then redirected to a page promoting a fictitious “Awesome AI Windows plugin” download.

Victims were prompted to retrieve ProFluxeFlowAi-win-Setup.exe from a GitHub repository named shippingtechnologymovie in the AI-techVideos path. The executable was signed with a fraudulently obtained Microsoft-issued code-signing certificate with thumbprint 4f5c5b3ef45cfff7721754487a86aeff9a2e6e32; after launch, it displayed a “Continue” checkbox, then dropped pythonw.exe and LICENSE.txt into *\AppData\Local*, executed shellcode, contacted brokeapt[.]com, and delivered Vidar infostealer.

Microsoft further stated that signed malware tends to achieve lower detection rates early in the infection lifecycle and that the extra “Continue” click was likely intended to defeat automated sandbox analysis because the binary remained inert until the user interacted with it. These details are important for Bangladeshi defenders because users often interpret signed executables, GitHub hosting, and familiar brand names as indicators of legitimacy.

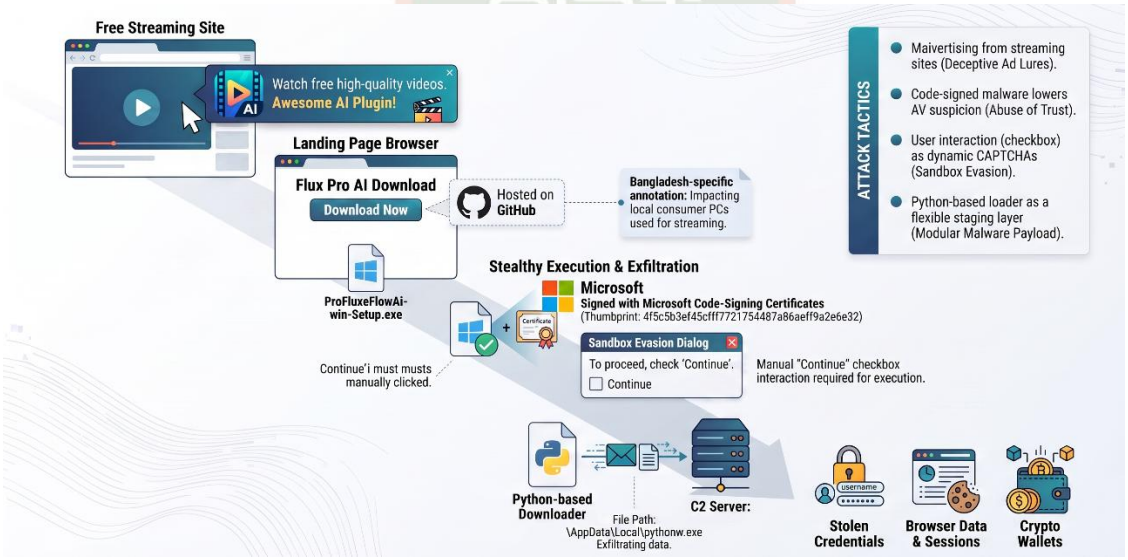

Figure: “Awesome AI Windows Plugin” malvertising

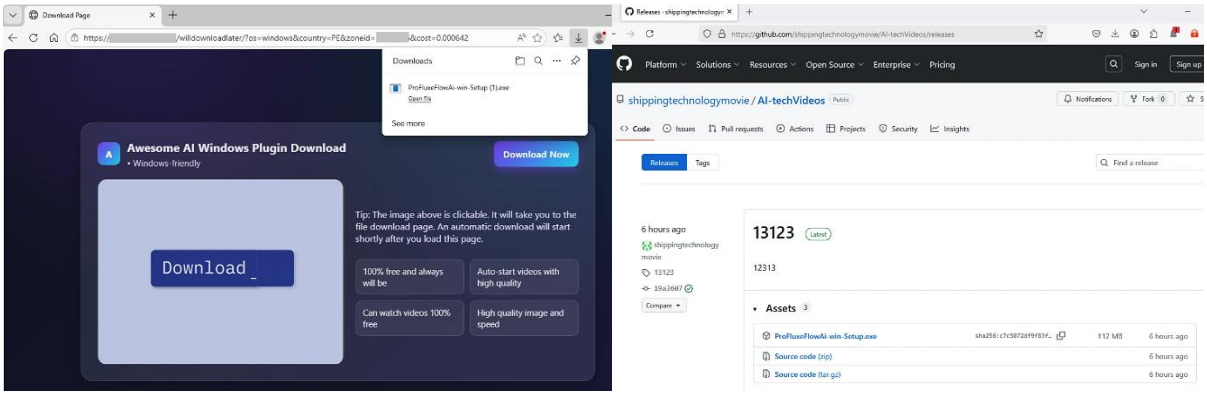

Figure: Malvertising page redirecting users to a purported "Awesome AI Windows Plugin" download and malware hosted in the "AI-techVideos" directory of the "shippingtechnologymovie" GitHub repository [Source: Microsoft]

Fake DeepSeek V4 GitHub repository distributing infostealers

In April 2026, Microsoft identified a fake DeepSeek V4 GitHub organization and repository that abused interest in a newly previewed model to distribute malware through GitHub release assets. Microsoft reported that within hours of the official preview, the threat actor created the GitHub organization DeepSeek-V4, the repository deepseek-V4, and the release tag deepseek-V4, then decorated the repository with stolen branding, real benchmark data, and SEO-optimized topics.

The release page hosted malicious archives such as **deepseek-v4-pro_x64.7z** and **deepseek-v4-flash_x64.7z**. Microsoft stated that users could discover the repository via GitHub search, web search, social sharing, and AI-assisted search results, aided by topic tags like deepseek-v4-install and an llms.txt file that repeated installer-oriented SEO language.

The archives contained a loader, including SHA-256 *5455341ed1bbe75a664fca2dd0794c508e1874f75360253a7ff5bc119bc92d80*, which Microsoft observed downloading and installing Vidar and potentially other malware. Microsoft also noted that the operator rotated archive content while preserving filenames and reused the same shared loader across lures that impersonated GPT-5.5, Claude Code, Kimi, Seedance, Gemma, GrokCLI, Manus AI, and FraudGPT, indicating a broader rotating brand-abuse ecosystem rather than a single DeepSeek-only campaign.

Microsoft observed that the repository accumulated 91 stars and 27 forks within four days, though the authenticity of that engagement was not independently confirmed. On closer review, the repository contained only superficial project materials such as README, LICENSE, llms.txt, and stub directories without real model code, and it showed inconsistent licensing claims and bursty commits from a single author, all of which are practical analytic clues for defenders and users reviewing suspicious repositories.

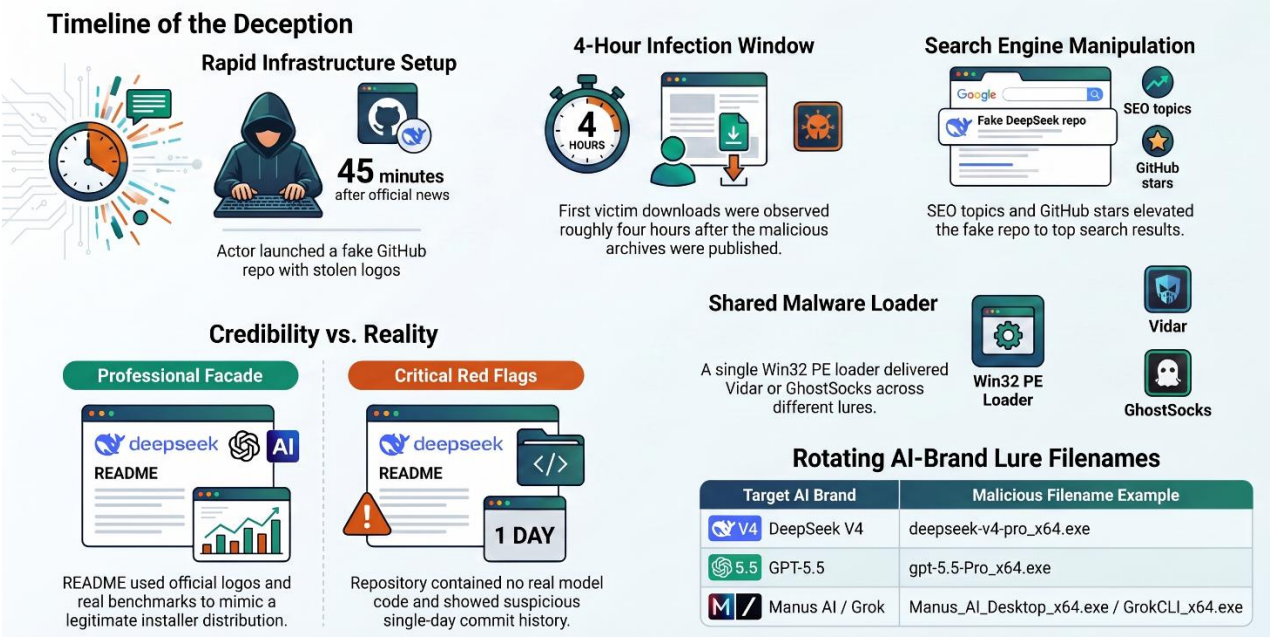

Figure: Fake DeepSeek V4 GitHub repo and rotating AI-brand ecosystem

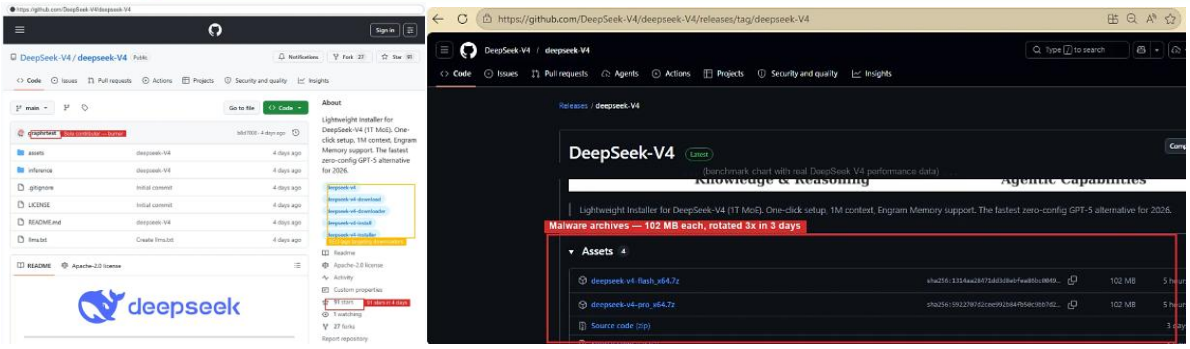

Figure: Fake DeepSeek-V4 GitHub repository using stolen branding, SEO optimization, a burner contributor account, and a fraudulent release page with authentic benchmark charts and malicious archive files. [Source: Microsoft]

Detection Opportunities

For Bangladesh, the highest-risk overlap is the combination of cloud identity use, unmanaged personal devices, interest in AI tools, and dependence on email and browser-based workflows. That mix makes both public-sector and consumer users vulnerable to the same tactics, even when the original campaign was first observed in another country.

- Monitor for AI-branded emails using billing, policy-violation, or appeal language, especially messages that create urgency with phrases such as payment update, account suspension, or policy review.

- Flag emails whose primary call-to-action is a button or link that chains through multiple domains, including CRM services, email-tracking domains, and URL shorteners before resolving to a final destination.

- Detect PDF attachments referencing AI services that instruct users to copy appeal IDs, verification codes, or account references and then click embedded links to unknown domains.

- Apply time-of-click URL inspection and link detonation to uncover late-stage redirects such as *Bitrix24* to *awstrack[.]me* to *Rebrandly* to a compromised phishing host.

- Hunt in DNS, proxy, firewall, and EDR telemetry for communications with known malicious infrastructure, including *dash.awaydouble[.]org*, *servicing.pureplantcravings[.]com*, *brokeapt[.]com*, *pan.ssffaa19[.]xyz*, and **pan.rongtv[.]xyz**.

- Alert on execution of newly downloaded AI-themed installers or archives from Downloads, Desktop, Temp, or other user-writable paths, especially when execution follows web browsing or GitHub release downloads.

- Detect suspicious creation and execution of pythonw.exe, script files, or text files such as *LICENSE.txt* inside *\AppData\Local*, followed by shellcode-like behavior or unexpected outbound network activity.

- Monitor for sign-in anomalies associated with AiTM activity, including impossible travel, unfamiliar IP addresses or devices, sudden session-token reuse, and rapid mailbox or cloud-app access after a phishing event.

- Review GitHub downloads and developer workflows for installer-style archives or binaries from repositories that use strong branding and marketing text but provide little or no legitimate source code.

- Prioritize hunting for shared malware pivots, including the loader hash *5455341ed1bbe75a664fca2dd0794c508e1874f75360253a7ff5bc119bc92d80* and the code-signing certificate thumbprint *4f5c5b3ef45cfff7721754487a86aeff9a2e6e32*, because Microsoft linked them to broader AI-themed malware-delivery activity.

MITRE ATT&CK Mapping

Tactic	Technique ID	Technique Name
Initial Access	T1566.002	Phishing: Spearphishing Link
Execution	T1204.001	User Execution: Malicious Link
Execution	T1204.002	User Execution: Malicious File
Credential Access	T1056.003	Input Capture: Web Portal Capture
Credential Access	T1556	Modify Authentication Process
Credential Access	T1528	Steal Application Access Token
Command and Control	T1105	Ingress Tool Transfer
Defense Evasion	T1036.005	Masquerading: Match Legitimate Name or Location
Defense Evasion	T1027	Obfuscated Files or Information

Indicators of Compromise (IoC)

Indicator	Type	Notes
hxxp://dash.awaydouble[.]org/0v2auth	URL	URL embedded in Claude-themed PDF lure
hxxps://github[.]com/shippingtechnologymovie/AI-techVideos/releases/download/13123/ProFluxeFlowAi-win-Setup.exe	URL	Malware-hosting GitHub release asset
hxxps://github[.]com/DeepSeek-V4/deepseek-V4/releases/download/deepseek-V4/deepseek-v4-pro_x64.7z	URL	Malicious DeepSeek-themed archive
dash.awaydouble[.]org	Domain	Claude-themed phishing infrastructure
servicing.pureplantcravings[.]com	Domain	Intermediate Claude-themed landing page
brokeapt[.]com	Domain	Python-loader C2 domain
pan.ssffaa19[.]xyz	Domain	Vidar C2
pan.rongtv[.]xyz	Domain	Vidar C2
791efb555eefb7215e96659a1353a97416743b66bdd72705493129c64057d40e	SHA-256	Hash for *Fill and Sign Claude Appeal Form.pdf*
c7c5072df9f83f4c440a5c3bb4be1d5f6c67bbf78f196406ca20d27b43b975b8	SHA-256	Hash for *ProFluxeFlowAi-win-Setup.exe*
4f5c5b3ef45cfff7721754487a86aeff9a2e6e32	Signer SHA-1	Fraudulently obtained signing certificate thumbprint
0a26238f6c516de5885457c93042531aa59bc206a9537cebf5267cedc6c68531	SHA-256	*deepseek-v4-pro_x64.7z* (v1)
8610d4fb0ec5b525071c2aaec4df0f8fcbb3673aba58a7e1959fc44e83c0e2ca	SHA-256	*deepseek-v4-flash_x64.7z* (v1)

99231deb373997364381d1eb513d2d42231d418c3a2db9007c5af9bd56ab9371	SHA-256	*deepseek-v4-flash_x64.7z* (v2)
25270cc429ada8028b5b33220ed412c47907ecceea7377d608fac5af01bed56a	SHA-256	*deepseek-v4-pro_x64.7z* (v2)
56d722b0331bf0aaa86bb37483486c6dff6ad9427fc473ed7c3226c21a9bdd23	SHA-256	Extracted DeepSeek-themed PE
5455341ed1bbe75a664fca2dd0794c508e1874f75360253a7ff5bc119bc92d80	SHA-256	Shared loader across multiple AI-themed lures

Recommendations

- Block and hunt for the listed IOCs across email, web proxy, DNS, EDR, firewall, and SIEM layers, and retain historical logs to identify prior exposure.
- Enforce MFA universally, with strong preference for phishing-resistant methods on administrative, financial, and external-facing accounts
- Enable time-of-click URL analysis, safe-link rewriting, attachment detonation, retrospective email purge, and detection of HTML-based phishing content in mail security controls.
- Restrict user execution from Downloads, Temp, archive-extraction directories, and other user-writable paths where practical, and apply application-control policies for unknown or low-reputation binaries.
- Inspect AI-related archive downloads and GitHub release assets before permitting execution on government or critical-sector devices.
- Monitor for abnormal sign-ins, token replay, impossible-travel events, new session cookies, and suspicious OAuth or cloud-authentication behavior consistent with AiTM phishing.
- Conduct targeted outreach to universities, software firms, financial institutions, and government offices because these sectors closely match the user behaviors and industry exposure patterns observed in Microsoft's reporting.

References

- https://www.microsoft.com/en-us/security/blog/2026/06/08/ai-brands-as-bait-how-threat-actors-are-using-the-ai-hype-in-social-engineering/
- https://exchange.xforce.ibmcloud.com/osint/guid:53d8878dddbf4ba3ae41cec409365b54