CSAI Foundation | Cloud Security Alliance

Forg365: AI Lure Generation in an M365 Phishing Kit

2026-07-10

AI-assisted Rapid Research

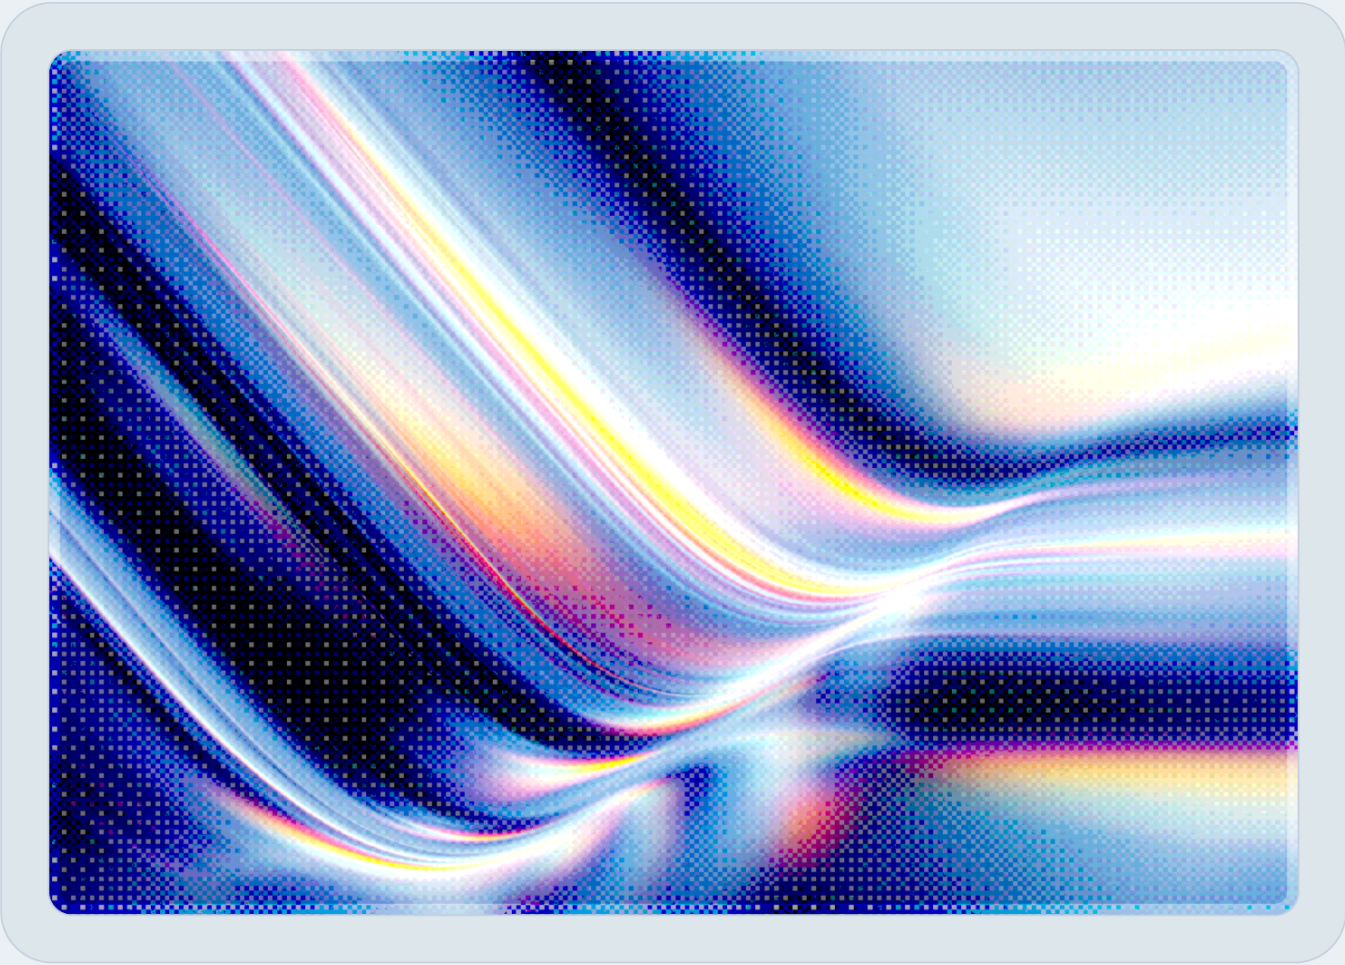

CSAI

CSA cloud security alliance®

Key Takeaways

Forg365 is a newly identified phishing-as-a-service (PhaaS) platform that targets Microsoft 365 accounts and folds AI-assisted lure generation directly into the operator's administrative dashboard, letting even unskilled buyers draft, refine, and localize phishing emails in the target's language and register without leaving the panel they use to run the rest of the campaign [1]. The platform combines two credential-theft paths already familiar to Microsoft 365 defenders: adversary-in-the-middle (AiTM) session hijacking and OAuth device-code phishing, the latter of which CSA has previously documented compromising more than 340 organizations in a single campaign [2]. Persistence is maintained through a companion browser extension, ForgCookie, that silently refreshes stolen Microsoft SSO cookies so operators retain access without needing the victim to re-authenticate [1]. Because Forg365 packages reconnaissance, lure drafting, delivery infrastructure, and post-compromise account monitoring into one commercial product, it likely lowers the skill and time investment needed to run a sustained business email compromise operation against Microsoft 365 tenants, though no data yet quantifies the effect on attacker success rates.

Background

Security researchers at ZeroBEC, an email security vendor, identified Forg365 after obtaining access to its administrative panel, which the researchers describe as a full campaign-management console rather than a simple phishing-page generator [1]. The panel lets an operator create new campaigns, manage phishing links, configure OAuth applications and SMTP sending profiles, and track harvested tokens and cookies from a single interface, and appears to mirror the feature set of commercial software-as-a-service products more than the ad hoc phishing kits security teams have dealt with for the past decade [1]. ZeroBEC was unable to establish a definitive link between Forg365 and other known PhaaS brands such as Kali365 or Sneaky2FA, despite functional overlap, suggesting either a new entrant building on established techniques or a rebrand that has not yet been attributed with confidence [1].

Forg365 arrives amid a broader surge in OAuth-based attacks against Microsoft 365. CSA's March 2026 research note on device-code phishing documented a campaign, tracing back to techniques Microsoft first attributed to the Russian-linked actor Storm-2372 in mid-2024, that compromised more than 340 organizations across five countries between February and March 2026 after the technique was commoditized into a platform called EvilTokens [2][3]. That campaign reached commercial availability

roughly eighteen months after Storm-2372's original activity; Forg365 followed a further five months after EvilTokens, meaning the full path from nation-state technique to a second, AI-enhanced commercial kit spanned closer to two years rather than eighteen months. Forg365 packages the same device-code flow alongside AiTM capture and adds AI-assisted content generation as a differentiating feature [1]. CSA's broader analysis of offensive AI use has found that large language models are increasingly being applied to planning realistic attack scenarios, primarily within penetration-testing and red-teaming workflows [4]; the appearance of AI drafting tools inside a criminal service platform such as Forg365 extends that same pattern into social-engineering content generation, though no independent testing has yet evaluated the persuasiveness of Forg365's output specifically.

Microsoft 365 remains an outsized target for this class of attack because a single compromised identity often grants access to email, files, calendars, and connected line-of-business applications through federated authentication, often making account takeover more valuable to an attacker than compromising an isolated endpoint [3]. The device-code flow itself, standardized in RFC 8628 for devices that lack a full browser, such as smart TVs or CLI tools, is legitimate Microsoft functionality that phishing operators have learned to abuse by directing victims to authorize an attacker-controlled device rather than their own [2]. Because the victim completes the real Microsoft authentication and MFA challenge on the attacker's behalf, the resulting refresh token is functionally indistinguishable from one issued during a legitimate login, and it persists across password resets, which is what makes the technique attractive to PhaaS operators in the first place [2].

Security Analysis

In CSA's assessment, Forg365's most consequential feature is not a new credential-theft technique but the integration of an AI drafting assistant into the same console an operator uses to manage delivery infrastructure and post-compromise activity. Reporting describes the feature as allowing operators to "create the malicious emails, prepare the text, and refine the messages" without switching tools, which collapses what used to be a separate content-creation step, often outsourced or copy-pasted from templates, into the campaign workflow itself [1]. In our assessment, this matters less because AI-written phishing text is necessarily more sophisticated than human-written text, and more because it likely removes a bottleneck: an operator can iterate on wording, tone, and targeting in real time and can plausibly generate lures in additional languages or industry-specific jargon without recruiting a native-language collaborator. That would lower the effective skill floor for running a credible business email compromise campaign and increase the volume of distinct, non-reused lure variants a single operator can push through email filters that rely partly on detecting repeated or templated content.

The platform's two credential-capture paths reflect a dual-track approach to session theft. In the AiTM path, Forg365 proxies the victim's authentication traffic to and from genuine Microsoft infrastructure, capturing the resulting session cookie as it passes through, a technique that defeats MFA generally because the attacker rides the legitimate authenticated session rather than needing to intercept a password or a one-time code [1]. In the device-code path, the victim is shown a Microsoft-branded verification page and unknowingly authorizes an attacker's device through the legitimate OAuth device authorization grant, again producing a valid, MFA-satisfied refresh token that Microsoft's systems cannot easily distinguish from a legitimate device enrollment [2]. Both paths converge on the same outcome that CSA highlighted in its device-code research: multifactor authentication does not stop the attack because the victim, not the attacker, is the one completing the MFA challenge [2].

Persistence is handled by ForgCookie, a browser extension built for Chrome, Edge, and Brave that, according to reporting, silently triggers OAuth flows to refresh Microsoft SSO cookies before they expire, keeping the operator's access alive without requiring the victim to authenticate again [1]. Paired with an account-intelligence dashboard that includes keyword monitoring across compromised mailboxes, the platform gives operators an alerting capability that flags mailboxes containing terms of interest, invoice, wire, password, and the like, letting a single operator triage which of many compromised accounts is worth acting on manually [1]. This shifts Forg365 from a one-time credential harvester toward a managed-access product, closer in structure to commercial remote-access tooling than to a disposable phishing template.

Operationally, Forg365 leans on legitimate cloud services to blend in with normal traffic: Amazon SES for outbound email delivery, Cloudflare Pages for hosting phishing landing pages, and the open-source Gophish framework for campaign orchestration [1]. Layered on top is an anti-analysis capability the researchers describe as including AES-encrypted redirectors, bot detection, debugger traps, sandbox checks, and polymorphic code, features aimed at frustrating researcher and automated-scanner access generally, even though ZeroBEC's own access to the panel evidently circumvented these defenses [1]. Using reputable cloud providers for delivery and hosting also complicates blunt domain- or IP-based blocking, since defenders risk false positives if they block Amazon SES or Cloudflare Pages outright.

Separately, reporting on a related wave of device-code abuse describes the extortion group "Pink" combining vishing calls, impersonating IT staff, with a fake Microsoft Entra passkey enrollment page, timing the fraudulent passkey setup to coincide with a real account-takeover attempt in the background [5]. While this is a single reported case rather than an established trend, and not confirmed as part of Forg365 specifically, it suggests device-code and passwordless-authentication abuse can be paired with voice-based social engineering. Defenders should watch for AI-assisted lure generation being combined with live voice pretexting, though no reporting yet documents that specific combination.

The commercial packaging of Forg365 also has implications for the economics of the PhaaS market itself. Platforms that bundle lure drafting, delivery, session capture, and post-compromise monitoring into one subscription reduce the number of specialized actors a criminal operation needs to coordinate, effectively vertically integrating a supply chain that previously required separately sourcing a phishing kit, a spam delivery service, and a manual triage step for harvested credentials [1]. That integration is likely to compress the time between a technique's first appearance in state-sponsored tooling and its availability as a point-and-click product for lower-skilled operators, a trend already visible in the roughly two-year span from Storm-2372's original device-code campaigns to EvilTokens and now Forg365, a path that included a further five months between the two commercial kits themselves [2][3]. Defenders evaluating their exposure should treat the pace of that commoditization, rather than any single platform's feature list, as the more durable signal, since the underlying OAuth abuse techniques will likely outlive Forg365 as a named brand.

Recommendations

Immediate Actions

Security teams should audit whether device-code authentication is enabled in their Microsoft Entra ID tenant and, if it is not required for legitimate use cases such as CLI tools or shared devices, block it through a Conditional Access policy that targets the Authentication Flows condition with the Device Code Flow option selected and the Grant action set to Block, which CSA's prior research identified as the highest-impact sustained control against this attack class [2]. Teams should also query Entra ID sign-in logs for device-code authentication events over the past 90 days to identify any unauthorized use that may already have occurred, and cross-reference results against known attacker infrastructure where indicators are available [2]. Any account showing signs of compromise should have all active sessions and refresh tokens revoked immediately using the `revokeSignInSessions` API, since password resets alone do not invalidate tokens obtained through this technique [2].

Short-Term Mitigations

Beyond device-code controls, organizations should review OAuth application consent grants and mailbox rules for anomalies, since AiTM-derived sessions are frequently used to register malicious OAuth applications or create forwarding rules that persist even after the original session is revoked [3]. Email security tooling should be evaluated for its ability to detect AI-generated lure content specifically, given that Forg365 and similar platforms are optimizing for volume and variation rather than relying on reused templates that legacy filters were tuned to catch [1]. Security awareness training should be updated to

cover device-code verification screens and fake passkey enrollment prompts as distinct phishing vectors, since much existing training has historically emphasized password-harvesting pages over emerging OAuth-based vectors.

Strategic Considerations

Organizations should treat the device-code authentication flow as a legacy protocol to be eliminated rather than merely monitored, consistent with CSA's prior guidance, given that its abuse has now been commoditized into at least two distinct commercial platforms within roughly six months of each other [2]. More broadly, the appearance of AI drafting assistants inside criminal service platforms should inform how security teams evaluate their own AI-assisted defenses: content-based phishing detection that depends on stylistic tells of human-written scam email, awkward phrasing, poor localization, will need to be reassessed as those tells disappear from AI-assisted lures [4]. CISOs should also factor the rising sophistication of managed-access PhaaS platforms, which combine content generation, delivery, and persistence into one product, into vendor risk conversations with any third party granted federated access into the organization's Microsoft 365 tenant, since a single compromised partner account can become a durable foothold rather than a one-time incident.

CSA Resource Alignment

CSA's March 2026 research note, "OAuth Device Code Phishing Hits 340+ Microsoft 365 Organizations," is the most directly relevant prior publication, since Forg365 packages the identical device-code abuse technique that note documented as compromising organizations across five countries, and its Conditional Access blocking guidance applies without modification to the Forg365 threat [2]. CSA's "Using AI for Offensive Security" report provides useful framing for understanding why AI-assisted lure generation matters operationally: the report anticipated that large language models would be applied to planning realistic attack scenarios within penetration-testing and red-teaming workflows, and Forg365's embedded AI drafting assistant shows that same offensive-AI pattern extending into commercial phishing content generation [4]. Finally, CSA's AI Controls Matrix (AICM v1.1) offers the governance and monitoring controls organizations should apply when evaluating their own exposure, particularly in the domains covering identity and access management and threat and vulnerability management, since Forg365's abuse of OAuth device-code and session-token flows falls squarely within the identity-centric risks those domains are designed to address [6].

References

[1] Toulas, Bill. "New Forg365 phishing platform uses AI to target Microsoft 365 accounts." BleepingComputer, July 2026.

[2] Cloud Security Alliance. "OAuth Device Code Phishing Hits 340+ Microsoft 365 Organizations." CSA AI Safety Initiative, March 2026.

[3] Microsoft Security. "Storm-2372 conducts device code phishing campaign." Microsoft Security Blog, February 2025.

[4] Cloud Security Alliance. "Using AI for Offensive Security." CSA AI Technology and Risk Working Group, August 2024.

[5] HostDir. "Forg365 PhaaS and Fake Passkey Attacks Target Microsoft 365 Accounts." HostDir Blog, July 2026.

[6] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." CSA, 2026.