CSAI Foundation | Cloud Security Alliance

ConsentFix v3: Automated Azure OAuth MFA Bypass

Pre-Consented First-Party App Abuse Automates Account Takeover in Microsoft Entra

2026-05-04

AI-assisted Rapid Research

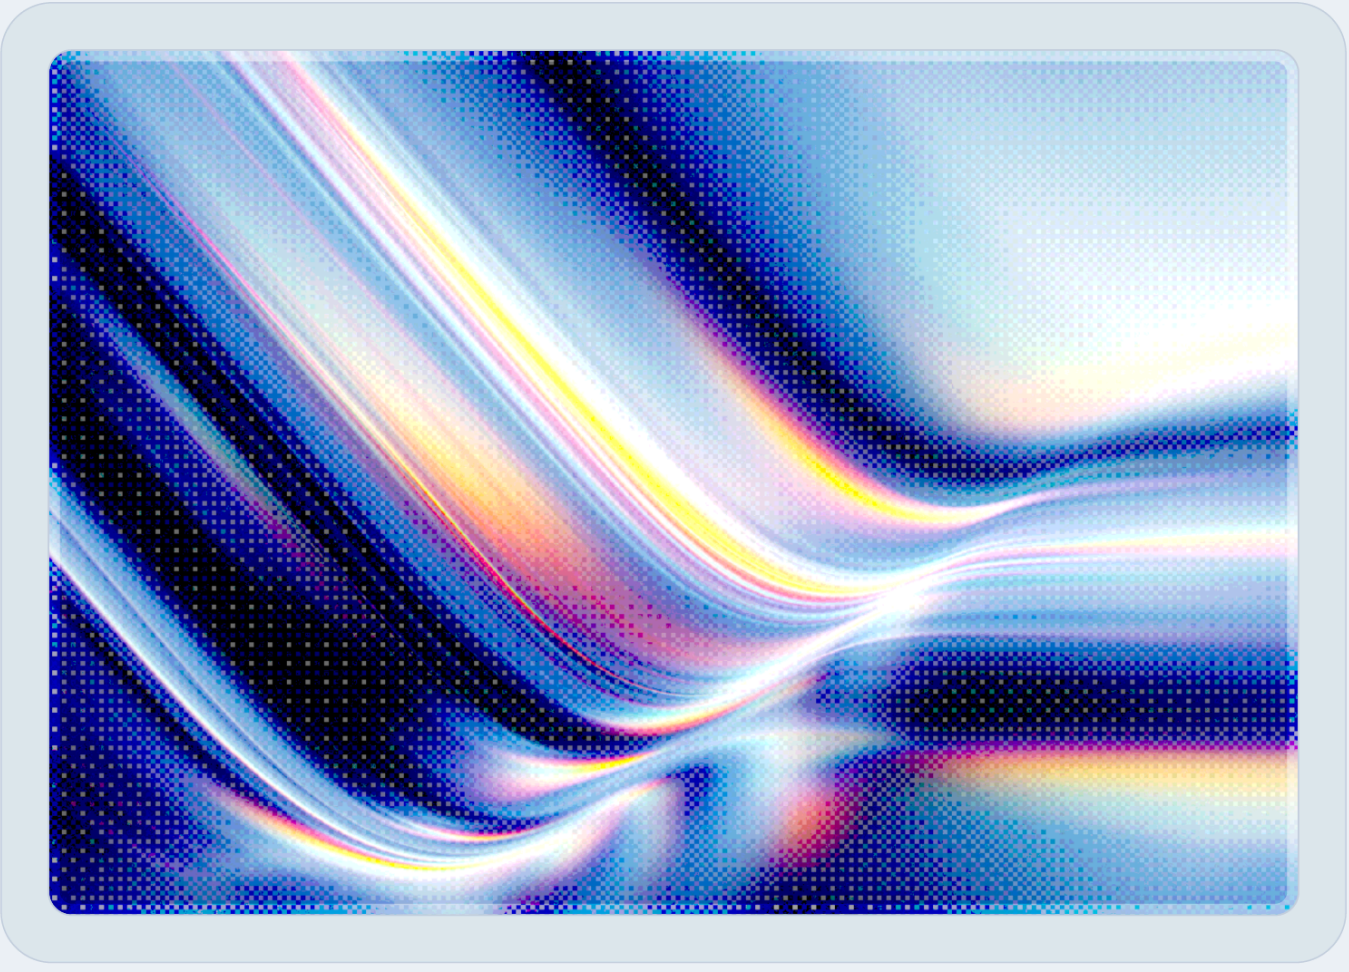

CSAI

CSA cloud security alliance®

Key Takeaways

- ConsentFix v3 is a criminal toolkit, distributed on the XSS forum in late April 2026, that automates the full OAuth authorization-code phishing kill chain against Microsoft Entra ID — from victim identification and phishing infrastructure setup through token capture, exchange, and post-exploitation — without requiring manual steps on the attacker's side after initial deployment [1][2].
- The attack exploits a structural trust property of Microsoft's identity platform: first-party applications such as Azure CLI are pre-consented across all Entra tenants, cannot be blocked or removed by tenant administrators, and are exempt from the standard OAuth consent prompts that would alert a victim to anomalous delegated-permission requests [3][4].
- Because ConsentFix abuses a legitimate OAuth flow rather than stealing passwords, a victim's completed MFA challenge actually advances the attack — the authorization code produced at the end of a successful authentication is the artifact the attacker captures and converts into durable account access [2][3].
- The Family of Client IDs (FOCI) mechanism — a behavior in Azure AD not officially documented by Microsoft but first described by Secureworks researchers in 2022 — allows a refresh token obtained from one family member application to be redeemed as any other application in the family, expanding the access surface beyond what the original consent grant appeared to authorize [5][6].
- Stolen tokens are imported into Specter Portal, a post-exploitation framework for Microsoft environments, enabling access to Exchange Online, SharePoint, Teams, and any resource covered by the victim's delegated permissions without further interaction with the victim's device [1][2].
- The most effective current technical control is Token Protection via Conditional Access, which binds tokens cryptographically to the originating device and renders an externally redeemed authorization code unusable; this control requires Entra ID Premium P1 or P2 licensing and does not yet cover browser-based applications [7][8][12].
- Organizations that have not explicitly audited the scope of first-party app consent grants, deployed Token Protection, or enabled monitoring of non-interactive sign-in logs are structurally exposed to this attack class regardless of standard MFA enforcement posture — including TOTP, push notification, and SMS-based methods — because the attack captures the authorization code produced after authentication completes [2][3].

Background

The ConsentFix attack family emerged publicly in December 2025 when Push Security, a browser-security firm, published research describing a novel variant of the "ClickFix" social-engineering class. Where ClickFix tricks users into pasting attacker-crafted content into terminals or run dialogs, ConsentFix redirects that social engineering pattern toward the OAuth authorization code flow [3][9]. The specific mechanism Push Security documented instructed victims to sign into Azure CLI through a legitimate Microsoft login page — completing any MFA challenge along the way — and then paste the resulting localhost redirect URL into an attacker-controlled site. That URL contains a live OAuth authorization code, which the attacker's infrastructure immediately exchanges for access tokens and refresh tokens via Microsoft's token endpoint.

Push Security's analysis established that Azure CLI is an especially attractive target for this technique because it occupies a privileged position within Microsoft's identity infrastructure. As a first-party application, it is implicitly trusted across every Entra tenant without any administrator approval, cannot be disabled or removed from the tenant, and is permitted to request delegated permissions without triggering the consent prompt that appears for third-party applications [3][4]. A victim authenticating through the Azure CLI flow sees what appears to be a normal Microsoft login screen, with no visual signal that the resulting authorization grant will be handed to an attacker.

In early 2026, security researcher John Hammond published a refined variant, commonly referred to as ConsentFix v2, that replaced the copy-paste instruction with a drag-and-drop interaction to reduce friction in the phishing flow and make the social engineering more convincing [1][2]. The progression from v1 to v2 follows the pattern observed in prior fix-class attacks, where delivery mechanisms evolve as defenders develop awareness of the preceding variant and users become familiar with earlier social engineering cues.

ConsentFix v3 represents a more significant evolution. In late April 2026, a member of the XSS criminal forum — a Russian-language platform whose membership has been linked to ransomware groups including REvil, LockBit, and Conti [14] — published a detailed toolkit post that reads, in Push Security's characterization, more like a vendor security blog than a forum advertisement [1]. The post introduces automation through a Pipedream-based backend, adds FOCI token escalation and Primary Refresh Token (PRT) device registration as post-exploitation paths, and describes a multi-platform infrastructure using Cloudflare Workers, ZoomInfo, DocSend, Hunter.io, and Tutanota to execute the campaign from

reconnaissance through exfiltration [1][2]. The shift from a manual proof-of-concept to an automated pipeline lowers the skill threshold for conducting the attack and increases the number of simultaneous targets a single operator can pursue.

Security Analysis

The Structural Vulnerability: First-Party App Trust

The ConsentFix attack class exploits a structural property of the trust architecture Microsoft has built into Entra ID for its own applications. Standard OAuth security guidance requires that applications explicitly request and receive user or administrator consent before obtaining delegated access to protected resources. Microsoft enforces this requirement for third-party applications, presenting a consent screen that lists the permissions being requested and requiring the user to affirmatively authorize them. First-party Microsoft applications are exempt from this requirement: their permissions are pre-authorized at the platform level, invisible to the user at authentication time, and unavailable to tenant administrators to revoke or restrict [3][4].

Azure CLI's implicit authorization extends to requesting delegated access to Microsoft Graph, Exchange Online, SharePoint, and Azure Resource Manager, among other services [3][5]. This means an authorization code obtained through Azure CLI carries permissions that the victim may not realize they have authorized, and that a tenant administrator reviewing their OAuth grant inventory would not see listed under explicit user-consented applications. The attack exploits the gap between what the identity platform displays to users and what the underlying delegated grant actually permits.

FOCI: Expanding the Post-Compromise Access Surface

Compounding this structural issue is a largely undocumented mechanism called the Family of Client IDs (FOCI). Documented initially by Secureworks researchers in 2022, FOCI groups a set of Microsoft first-party public-client applications — including Azure CLI, Azure PowerShell, Microsoft Teams, and others — into a shared token family [5][6]. A refresh token obtained from any member of this family can be redeemed for access tokens impersonating any other family member, for any scope that the chosen family client has been granted consent to request.

The practical implication for ConsentFix v3 is that an attacker who captures an authorization code through an Azure CLI phishing flow can use the resulting refresh token not only as Azure CLI, but as any other FOCI-family application. Access that superficially appears limited to CLI operations expands to the

full scope of the family, and the OAuth standard's expectation that refresh tokens be bound to a specific client and scope does not apply to Azure AD's implementation of FOCI [5]. Token redemption events for family refresh tokens appear in Entra logs under non-interactive sign-ins — a category that many organizations monitor less closely than interactive authentication events — and carry no indicator flagging the FOCI mechanism as active [5][10].

ConsentFix v3: The Automated Attack Pipeline

The toolkit described in the XSS forum post builds the complete attack pipeline from a single orchestration layer. The operator begins by using ZoomInfo or comparable data sources to enumerate employees at the target organization, harvesting names, roles, and email addresses to support convincing impersonation in the phishing emails [1]. The phishing messages are AI-assisted, highly personalized based on harvested profile data, and embed malicious links inside PDFs hosted on DocSend — a tactic that exploits DocSend's legitimate reputation to pass organizational spam filters and email security gateways [2].

When a victim clicks the embedded link and is directed through the phishing flow, the critical automation occurs at Pipedream, a free serverless integration platform. Pipedream serves three roles in the v3 architecture: it is the webhook endpoint that receives the victim's OAuth authorization code, the automation engine that immediately exchanges that code for a refresh token via Microsoft's token endpoint, and the central collector that makes captured tokens available to the operator in near real time [1][2]. Microsoft has implemented a short authorization code lifetime — currently one minute — to limit the window for code capture, but the v3 toolkit's near-real-time Pipedream automation operates well within that window [8].

Following token capture, the toolkit incorporates two post-exploitation escalation paths. The primary path involves importing the captured tokens into Specter Portal, a post-exploitation framework for Microsoft 365 environments, which provides an operator interface for reading email, accessing SharePoint and OneDrive files, and interacting with other resources covered by the token's scope [2]. The advanced path — which Push Security characterizes as closer to red-team tradecraft than operational criminal tooling — involves targeting the Microsoft Authentication Broker application to request a Primary Refresh Token and then leveraging that PRT to register a new device under the victim's identity within Entra ID, achieving persistent durable access that survives password resets and token revocations [1].

MFA Bypass Mechanics

The MFA bypass in ConsentFix is architectural rather than technical. No MFA implementation is bypassed in the conventional sense of being circumvented, forged, or exhausted. Instead, the victim's successful completion of MFA is a prerequisite for the attack: the authorization code is issued only after the identity provider has verified the user's identity, including any MFA factors. The attacker captures the authorization code that results from a completed, legitimate authentication session and redeems it from their own infrastructure. The MFA challenge was satisfied by the victim; the attacker inherits the resulting token without having passed any authentication test themselves [3][4].

Completing MFA — whether via TOTP, push notification, hardware key, or passkey — does not prevent ConsentFix from succeeding if the victim follows the phishing flow to completion. The authentication factor protects the sign-in event; it does not protect the resulting authorization code from capture. Phishing-resistant methods such as FIDO2 passkeys retain value against other attack classes and complement Token Protection in this scenario, but they do not independently close this vulnerability. Standard Conditional Access policies that evaluate device compliance and sign-in context at authentication time do not prevent post-capture token reuse, because the authentication event itself was legitimate; policies enforcing Continuous Access Evaluation (CAE) or Token Protection impose conditions on ongoing token use and provide substantially better coverage, as discussed in the Recommendations section [7][8]. Standard geographic anomaly detection may flag token usage from unexpected locations, but automated Conditional Access rules that trigger on sign-in anomalies rather than token usage anomalies will not fire during the capture phase [2][7].

Recommendations

Immediate Actions

Organizations operating Microsoft Entra ID environments should treat ConsentFix v3 as an active threat requiring immediate posture assessment. The first priority is auditing non-interactive sign-in logs for anomalous OAuth authorization code exchanges — particularly any exchange involving Azure CLI or other FOCI-family applications originating from IP addresses, geographies, or ASNs inconsistent with the organization's known user population. Pipedream's infrastructure (pip.me domains, pipedream.net) and similar serverless platform endpoints appearing as OAuth redirect targets in sign-in logs should be treated as indicators of active compromise.

Security teams should brief employees on the ConsentFix social engineering pattern: no legitimate IT system or Microsoft service will instruct a user to paste a URL from their browser into an external site, nor to drag-and-drop browser content to a third party. Awareness communications should specifically address the PDF-in-DocSend delivery mechanism and the pattern of highly personalized phishing messages asking for "account verification" or similar pretexts that trigger an Azure CLI login flow.

Short-Term Mitigations

The highest-impact technical control available today is Token Protection via Conditional Access. When Token Protection is enforced, access tokens are cryptographically bound to the device on which the authentication session originated through a proof-of-possession mechanism. An authorization code captured from a victim's browser cannot be successfully redeemed by an attacker operating on a different device, because the broker-signed device proof required to complete the exchange is absent. Organizations should create a Conditional Access policy in report-only mode targeting Exchange Online, SharePoint Online, and Microsoft Teams, monitor for incompatibility exceptions over a pilot period, and then enforce the policy organization-wide [7][8]. Token Protection requires Entra ID Premium P1 or P2 licensing and is currently limited to native applications on Windows 10 or later; browser-based applications are not yet covered [7][12].

Beyond Token Protection, organizations should enable Conditional Access policies that restrict authorization code redemption to compliant or Hybrid Azure AD Joined devices. While this does not prevent code capture, it limits the infrastructure from which a captured code can be exchanged. Configuring Continuous Access Evaluation (CAE) ensures that revocation signals propagate quickly to supported applications, reducing the window in which a stolen token remains usable following detection.

Microsoft's Entra audit logs capture OAuth code grants under the "Add delegation" event; enabling alerts on this event for first-party applications — particularly Azure CLI — provides an additional detection layer without requiring premium licensing. Security teams using SIEM platforms should implement the detection rules published by NVISO in January 2026, which identify ConsentFix patterns specifically in Entra sign-in logs [10].

Strategic Considerations

The ConsentFix attack class reflects a structural design decision in Microsoft's identity architecture — the implicit trust and permanent tenant-wide consent granted to first-party applications — that no individual customer can remediate unilaterally. Organizations should formally evaluate whether their enterprise risk tolerance is compatible with this constraint and, where it is not, engage Microsoft through their enterprise

account teams to advocate for administrative controls that would allow tenant administrators to restrict or monitor first-party application consent grants, publish the complete FOCI family membership list, and provide Token Protection coverage for browser-based applications.

Based on the trajectory of comparable attack classes, wider operational adoption of ConsentFix v3 is a plausible near-term scenario that security teams should plan for, rather than a background risk. Organizations that have deferred implementing Token Protection or non-interactive sign-in monitoring should treat that trajectory as an operational deadline.

At the identity architecture level, the FOCI mechanism warrants explicit treatment in access governance programs. Because FOCI refresh tokens grant broader application access than the original consent grant implies, organizations with least-privilege access policies should treat any FOCI-family application grant as equivalent to granting the combined permission set of all family members. This framing clarifies the actual risk of what appears to be limited CLI access and may influence decisions about conditional access scoping, privileged identity management enrollment, and role-based access control for sensitive resources.

CSA Resource Alignment

ConsentFix v3 maps directly to several established CSA frameworks and guidance documents. Within the Cloud Controls Matrix (CCM), the attack engages Identity and Access Management (IAM) domain controls — specifically IAM-02 (Strong Authentication), IAM-08 (User Access Restriction), and IAM-09 (User Access Reviews) — as well as Infrastructure and Virtualization Security controls concerning API protection and access governance.

CSA's Zero Trust guidance is directly applicable. The ConsentFix attack class succeeds precisely because identity verification at authentication time is treated as sufficient to authorize ongoing resource access: once the token is issued, its origin and binding are not continuously evaluated. A Zero Trust posture that enforces continuous validation of device compliance, network context, and behavioral signals — rather than treating a successfully authenticated session as implicitly trustworthy — materially limits post-compromise lateral movement even when an attacker holds a valid token.

The CSA white paper Using Zero Trust to Counter Identity Spoofing and Abuse [13] addresses the taxonomy of identity abuse patterns relevant here, including delegation abuse and token theft. CSA's ongoing research on OAuth trust abuse in AI SaaS integrations addresses the structural trust-chain risk

that ConsentFix v3 operationalizes from a different phishing vector; the applicable mitigations — Token Protection, OAuth grant auditing, and behavioral monitoring of non-interactive sign-ins — are consistent across both attack surfaces.

For organizations implementing STAR self-assessments, ConsentFix v3 raises specific concerns under the Identity and Access Management category: assessment teams should verify that token binding controls are in scope, that first-party application grant inventories exist and are reviewed periodically, and that non-interactive sign-in monitoring is included in the security operations program rather than treated as a lower-priority logging category.

References

[1] Push Security. "Investigating a new criminal toolkit for ConsentFix." Push Security Blog, May 2026.

[2] BleepingComputer. "ConsentFix v3 attacks target Azure with automated OAuth abuse." BleepingComputer, May 2026.

[3] Push Security. "ConsentFix: Browser-native ClickFix hijacks OAuth grants." Push Security Blog, December 2025.

[4] eSecurity Planet. "Azure CLI Trust Abused in ConsentFix Account Takeovers." eSecurity Planet, 2026.

[5] Secureworks. "Family of Client IDs Research." GitHub, 2022.

[6] Payatu. "Microsoft's Family of Client IDs (FOCI): Convenience vs. Compromise." Payatu Blog, 2023.

[7] Microsoft. "How Token Protection Enhances Conditional Access Policies." Microsoft Learn, 2025.

[8] AdminDroid. "What is ConsentFix Attack and How to Mitigate it in Microsoft 365." AdminDroid Blog, 2026.

[9] glueckkanja. "ConsentFix: How a New OAuth Attack Bypasses Microsoft Entra Conditional Access." glueckkanja Blog, December 2025.

[10] NVISO. "ConsentFix (a.k.a. AuthCodeFix): Detecting OAuth2 Authorization Code Phishing." NVISO Blog, January 2026.

[11] Mitiga. "ConsentFix OAuth Phishing Explained: How Token-Based Attacks Bypass MFA in Microsoft Entra ID." Mitiga Blog, 2026.

[12] Microsoft. "Token Protection Deployment Guide - Windows." Microsoft Learn, 2025.

[13] Cloud Security Alliance. "Using Zero Trust to Counter Identity Spoofing and Abuse." CSA, 2023.

[14] Flare. "Top Russian-Language Cybercrime Forums to Monitor." Flare Blog, 2025.

