INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-03-29 18:33 UTC

AI-Augmented OAuth Phishing Campaign Compromises 344 Organizations via Microsoft Cloud Account Abuse

THREAT CAMPAIGN | **HIGH** | CVSS 8.1

SCC Item ID	SCC-CAM-2026-0123
Type	Threat Campaign
Severity	HIGH
CVSS Base Score	8.1
Affected Products	Microsoft 365, Microsoft Entra ID, Microsoft OAuth 2.0 authorization endpoints (cloud accounts leveraging OAuth token flows)
Published	2026-03-29
Discovery Source	Gemini

Executive Summary

An AI-assisted phishing campaign has compromised Microsoft cloud accounts across at least 344 organizations in construction, law, healthcare, and government sectors. Attackers abuse legitimate OAuth 2.0 authorization flows to harvest access tokens and refresh tokens, gaining persistent access to Microsoft 365 services without ever stealing passwords. The business risk is significant: compromised accounts can expose sensitive data, enable lateral movement across Microsoft 365 tenants, and persist even after password resets unless OAuth token grants are explicitly revoked.

Technical Analysis

This campaign exploits OAuth 2.0 redirect URI abuse and token delegation mechanisms in Microsoft Entra ID and Microsoft 365 environments. Attackers craft malicious OAuth authorization requests that direct victims to attacker-controlled applications, capturing access tokens and refresh tokens rather than credentials. Because refresh tokens enable long-lived session persistence, account access survives password changes unless token grants are explicitly invalidated. No CVE has been assigned; the technique abuses legitimate OAuth protocol features. Relevant CWEs include CWE-601 (Open Redirect), CWE-352 (CSRF), CWE-346 (Origin Validation Error), and CWE-287 (Improper Authentication). MITRE ATT&CK coverage spans T1566 and T1566.002 (Phishing/Spearphishing Link), T1528 (Steal Application Access Token), T1539 (Steal Web Session Cookie), T1556.006 (Hybrid Identity), T1550.001 (Use Alternate Authentication Material: Application Access Token), and T1078/T1078.004 (Valid Accounts: Cloud Accounts). AI-generated lure content is attributed to increased

phishing plausibility and victim interaction rates. Sources: Microsoft Security Blog (2026-03-02), The Register (2026-03-03).

Action Checklist

1. Step 1: Containment, Immediately audit third-party OAuth application consents granted in Microsoft Entra ID. Open Entra ID > Enterprise Applications > All Applications and filter for user-consented apps added in the last 90 days. Revoke consent for unrecognized or suspicious OAuth applications. Suspend accounts with confirmed token compromise and revoke all refresh tokens for those accounts using the Entra ID 'Revoke Sign-In Sessions' function or the Microsoft Graph API revokeSignInSessions endpoint. Reference: Microsoft Security Blog (2026-03-02).
2. Step 2: Detection, Query Microsoft Entra ID sign-in logs (Entra ID > Monitoring > Sign-in Logs) for OAuth authorization flows originating from unfamiliar application IDs, especially those with delegated permissions for Mail.Read, Files.ReadWrite, or Contacts.Read. In Microsoft Sentinel or Defender for Cloud Apps, hunt for AppConsentGranted events and OAuthApp policy alerts. Look for refresh token usage from anomalous IP geolocations or user agents inconsistent with the legitimate application. Behavioral indicator: legitimate user accessing Microsoft 365 resources but with an application client ID not matching your approved application inventory.
3. Step 3: Eradication, Remove unauthorized OAuth application registrations from Entra ID. Revoke delegated permission grants using PowerShell (Remove-MgOauth2PermissionGrant) or the Entra admin portal. Enforce admin consent requirements for all OAuth application authorizations by enabling the 'Require admin consent for new application registrations' policy in Entra ID > User Settings > Enterprise Applications. This eliminates the user-consent path attackers rely on. Reference: Microsoft Entra ID application consent configuration documentation.
4. Step 4: Recovery, After revoking compromised token grants and removing malicious app registrations, validate by re-querying the OAuth consent audit logs to confirm no residual grants exist for the identified malicious application IDs. Re-enable affected accounts only after token revocation is confirmed. Monitor the re-enabled accounts via Entra ID risky sign-ins and Defender for Cloud Apps anomaly detection for at least 30 days post-remediation. Validate that Conditional Access policies requiring compliant devices and MFA are enforced for all cloud account access.
5. Step 5: Post-Incident, This campaign exposes a control gap in OAuth application consent governance. Implement a formal OAuth application allowlist and enforce admin-only consent as a standing policy. Evaluate deployment of Microsoft Entra ID's App Governance add-on for continuous OAuth permission monitoring. Review phishing awareness training to include OAuth consent phishing scenarios, which bypass credential-focused training. Map control improvements to NIST SP 800-53 AC-2 (Account Management), AC-17 (Remote Access), and SI-4 (Information System Monitoring).

IR / Forensic Enrichment

Triage Priority	IMMEDIATE

Escalation Criteria	Escalate immediately to legal and compliance leadership if any compromised account had delegated Mail.Read or Files.ReadWrite access to mailboxes or SharePoint sites containing PHI (HIPAA breach notification threshold), PII subject to state notification laws, or attorney-client privileged matter files, given the campaign's confirmed targeting of healthcare, law, and government sectors and the absence of a CVSS-defining software vulnerability — token-based persistence means compromise duration may extend 90+ days prior to detection.
Recovery Notes	Re-enable affected accounts only after confirming via 'Get-MgUserOauth2PermissionGrant' that zero delegated grants remain for the malicious application IDs and that 'Revoke-MgUserSignInSession' has been executed and reflected in sign-in logs as a forced re-authentication event. Monitor all formerly compromised accounts for a minimum 30-day window using Entra ID Identity Protection 'Anomalous Token' and 'Unfamiliar Sign-in Properties' risk detections, as OAuth refresh tokens issued before full revocation may persist in attacker-controlled caches and produce delayed re-use attempts. Validate that Conditional Access policies enforcing MFA and compliant-device requirements apply to all OAuth 2.0 authorization code flows, not just interactive browser sign-ins, as this campaign specifically exploits the delegated authorization path that misconfigured policies leave unguarded.
Forensic Artifacts	Entra ID Unified Audit Log — 'Consent to application' and 'Add delegated permission grant' operations: contain victim UPN, malicious AppId, OAuth scope strings (Mail.Read, Files.ReadWrite, Contacts.Read), and grant timestamp — the primary evidence chain linking phishing delivery to authorized data access | Entra ID Sign-In Logs — non-interactive token refresh entries: 'tokenIssuanceType', 'appId', 'ipAddress', 'location', 'deviceDetail/userAgent' fields record every instance the attacker reused a harvested refresh token to silently obtain new access tokens for Microsoft 365 services without user interaction | Exchange Online Unified Audit Log — 'MailItemsAccessed' operations filtered by the malicious AppId: documents the exact mailbox folders, message IDs, and timestamps of attacker mail access performed under the delegated Mail.Read grant, establishing data exfiltration scope for breach notification purposes | Entra ID App Registration and Enterprise Application objects for malicious apps: the 'replyUrls' (redirect URIs) field contains the attacker's token-harvesting callback infrastructure; 'requiredResourceAccess' and 'oauth2Permissions' fields document the full permission scope requested — preserve full JSON export via 'Get-MgApplication -ApplicationId | ConvertTo-Json -Depth 10' | Microsoft Graph API activity logs (if Microsoft Purview Audit Premium is licensed) or Exchange Online/SharePoint Online audit logs filtered on the malicious ServicePrincipalId: reconstruct the complete post-compromise activity timeline showing which files were read or downloaded from OneDrive/SharePoint under the Files.ReadWrite delegated grant, supporting both forensic scoping and regulatory notification obligations

Per-Action IR Details

Step 1: Containment — Immediately audit third-party OAuth application consents granted in Microsoft Entra ID. Navigate to Entra ID > Enterprise Applications > All Applications and filter for user-consented apps added in the last 90 days. Revoke consent for unrecognized or suspicious OAuth applications. Suspend accounts with confirmed token compromise and revoke all refresh tokens for those accounts using the Entra ID 'Revoke Sign-In Sessions' function or the Microsoft Graph API revokeSignInSessions endpoint. Reference: Microsoft Security Blog (2026-03-02).

NIST Phase: Containment

Reference: NIST 800-61r3 §3.3 — Containment Strategy

Controls: NIST IR-4 (Incident Handling), NIST AC-2 (Account Management), NIST AC-17 (Remote Access), CIS 5.3 (Disable Dormant Accounts), CIS 6.2 (Establish an Access Revoking Process)

Compensating: Without Entra ID Premium or Sentinel, use the Microsoft Graph PowerShell SDK (free): run 'Get-MgUserOauth2PermissionGrant -UserId ' for each suspected user to enumerate delegated grants, then 'Revoke-MgUserSignInSession -UserId ' to invalidate all refresh tokens. Export the full OAuth grant list to CSV via 'Get-MgOauth2PermissionGrant | Export-Csv oauth_grants.csv' for offline review by two analysts cross-referencing against your approved app inventory. For tenants without a formal app inventory, cross-check client IDs against Microsoft's own first-party app ID list published in the Entra documentation.

Evidence: Before revoking any tokens, export the complete Entra ID sign-in log for the 90-day window (Entra ID > Monitoring > Sign-in Logs > Download) preserving the 'AppId', 'AppDisplayName', 'ConditionalAccessStatus', 'IPAddress', and 'UserAgent' columns — these fields record the exact OAuth client IDs and IP geolocations used by the attacker's malicious application during token harvest. Capture the 'AuditLogs' blade export filtered on 'Activity: Consent to application' to preserve the user-consent grant event including the timestamp, consenting user UPN, and the granted OAuth scopes (Mail.Read, Files.ReadWrite, Contacts.Read) before revocation destroys the forensic link between the malicious app registration and victim account.

Step 2: Detection — Query Microsoft Entra ID sign-in logs (Entra ID > Monitoring > Sign-in Logs) for OAuth authorization flows originating from unfamiliar application IDs, especially those with delegated permissions for Mail.Read, Files.ReadWrite, or Contacts.Read. In Microsoft Sentinel or Defender for Cloud Apps, hunt for AppConsentGranted events and OAuthApp policy alerts. Look for refresh token usage from anomalous IP geolocations or user agents inconsistent with the legitimate application. Behavioral indicator: legitimate user accessing Microsoft 365 resources but with an application client ID not matching your approved application inventory.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis

Controls: NIST SI-4 (System Monitoring), NIST AU-6 (Audit Record Review, Analysis, and Reporting), NIST AU-2 (Event Logging), NIST IR-5 (Incident Monitoring), CIS 8.2 (Collect Audit Logs)

Compensating: Without Sentinel or Defender for Cloud Apps, query the Unified Audit Log (UAL) directly via PowerShell: 'Search-UnifiedAuditLog -StartDate (Get-Date).AddDays(-90) -EndDate (Get-Date) -RecordType AzureActiveDirectory -Operations "Consent to application" | Select-Object -ExpandProperty AuditData | ConvertFrom-Json | Select UserKey, ObjectId, ModifiedProperties | Export-Csv consent_events.csv'. Then run a second query for 'Add delegated permission grant' operations. Cross-reference the returned AppId values against Microsoft's published list of known-malicious OAuth app IDs from this campaign (if released by CISA or the Microsoft Security Blog advisory). For refresh token geolocation anomalies without SIEM, export sign-in logs via Graph API ('GET /auditLogs/signIns?$filter=appId eq '''') and manually review the 'location' and 'deviceDetail' fields for inconsistencies.

Evidence: Capture the full Unified Audit Log export for 'ConsentGranted' and 'Add delegated permission grant' operations spanning 90 days — these records contain the OAuth scope strings (e.g., 'Mail.Read', 'Files.ReadWrite.All', 'Contacts.Read') that directly map to the data the attacker exfiltrated post-compromise. Preserve Entra ID sign-in log entries where 'tokenIssuanceType' equals 'AzureAD' and 'authenticationDetails' shows 'Non-interactive' token refresh events from IP addresses not matching the user's historical login geolocations — these non-interactive refresh token uses are the operational fingerprint of the attacker reusing harvested tokens without re-authenticating. Document all unique 'AppId' and 'ServicePrincipalId' values observed in these logs as your malicious application indicator set.

Step 3: Eradication — Remove unauthorized OAuth application registrations from Entra ID. Revoke delegated permission grants using PowerShell (Remove-MgOauth2PermissionGrant) or the Entra admin portal. Enforce admin consent requirements for all OAuth application authorizations by enabling the 'Require admin consent for new application registrations' policy in Entra ID > User Settings > Enterprise Applications. This eliminates the user-consent path attackers rely on. Reference: Microsoft Entra ID application consent configuration documentation.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST SI-2 (Flaw Remediation), NIST CM-7 (Least Functionality), NIST AC-2 (Account Management), CIS 4.6 (Securely Manage Enterprise Assets and Software), CIS 2.3 (Address Unauthorized Software)

Compensating: For teams without Entra ID P2 governance features, execute eradication entirely via Microsoft Graph PowerShell (free, no license required): enumerate all user-consented grants with 'Get-MgOauth2PermissionGrant -Filter "consentType eq 'Principal'" | Export-Csv user_consented_grants.csv', then for each confirmed malicious grant run 'Remove-MgOauth2PermissionGrant -OAuth2PermissionGrantId '. To block the user-consent path without a GUI policy change, use 'Update-MgPolicyAuthorizationPolicy -AllowInvitesFrom adminsAndGuestInviters -PermissionGrantPolicyIdsAssignedToDefaultUserRole @()' to strip default user consent permissions entirely — this is the free equivalent of the admin-consent-only policy and closes the attack vector this campaign exploits.

Evidence: Before executing 'Remove-MgOauth2PermissionGrant', record the full permission grant object for each malicious app including 'resourceId' (the Graph/Exchange Online service principal), 'scope' (the delegated permission strings), 'principalId' (victim user object ID), and 'startTime' (grant creation timestamp) — this preserves the legal and forensic record of what data access the attacker's application was authorized to perform. Separately, capture the App Registration object for any attacker-controlled apps found in 'App Registrations' (not just Enterprise Applications), including the 'replyUrls' (redirect URIs) field, which will contain the attacker's token-harvesting infrastructure URL and is critical evidence for threat intelligence and potential law enforcement referral.

Step 4: Recovery — After revoking compromised token grants and removing malicious app registrations, validate by re-querying the OAuth consent audit logs to confirm no residual grants exist for the identified malicious application IDs. Re-enable affected accounts only after token revocation is confirmed. Monitor the re-enabled accounts via Entra ID risky sign-ins and Defender for Cloud Apps anomaly detection for at least 30 days post-remediation. Validate that Conditional Access policies requiring compliant devices and MFA are enforced for all cloud account access.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST IR-4 (Incident Handling), NIST SI-4 (System Monitoring), NIST AC-17 (Remote Access), NIST AU-6 (Audit Record Review, Analysis, and Reporting), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 6.5 (Require MFA for Administrative Access)

Compensating: Without Defender for Cloud Apps, configure Microsoft Entra ID Identity Protection (included in Entra ID P1) alert policies for 'Unfamiliar sign-in properties' and 'Anomalous Token' detections on the re-enabled accounts — the 'Anomalous Token' detection specifically fires on refresh token reuse anomalies consistent with this campaign's post-compromise behavior. For teams without any P1 licensing, schedule a daily PowerShell job: 'Get-MgUserOauth2PermissionGrant -UserId ' on all formerly compromised accounts and diff the output against the post-eradication baseline CSV, alerting on any new user-consented grants — this provides a manual tripwire for reinfection within a two-person team's operational capacity.

Evidence: Post-recovery validation evidence: re-run 'Get-MgOauth2PermissionGrant -Filter "clientId eq '" and confirm zero results for each of the identified malicious application IDs — screenshot or log the empty result set as the official eradication confirmation artifact. Query Exchange Online audit logs ('Search-UnifiedAuditLog -Operations MailItemsAccessed -UserIds ') for any 'MailItemsAccessed' events attributed to the malicious AppId post-revocation — any hits indicate the token revocation did not propagate fully to Exchange Online and the session is still active. Retain Conditional Access policy export (JSON via Graph API 'GET /identity/conditionalAccess/policies') as the baseline configuration record for the 30-day monitoring window.

Step 5: Post-Incident — This campaign exposes a control gap in OAuth application consent governance. Implement a formal OAuth application allowlist and enforce admin-only consent as a standing policy. Evaluate deployment of Microsoft Entra ID's App Governance add-on for continuous OAuth permission monitoring. Review phishing awareness training to include OAuth consent phishing scenarios, which bypass credential-focused training. Map control improvements to NIST SP 800-53 AC-2 (Account Management), AC-17 (Remote Access), and SI-4 (Information System Monitoring).

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity

Controls: NIST IR-4 (Incident Handling), NIST IR-8 (Incident Response Plan), NIST AC-2 (Account Management), NIST AC-17 (Remote Access), NIST SI-4 (System Monitoring), NIST SI-5 (Security Alerts, Advisories, and Directives), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 5.1 (Establish and Maintain an Inventory of Accounts)

Compensating: For organizations without budget for the Entra App Governance add-on, implement a free continuous monitoring equivalent using a scheduled PowerShell script that runs nightly: export all OAuth2PermissionGrants and ServicePrincipalCreation audit events via Graph API, diff against a known-good baseline, and email a delta report to the security team — this provides the core detection capability (new user-consented grants and new app registrations) that App Governance delivers commercially. For phishing awareness training specific to OAuth consent phishing, use the Microsoft Attack Simulator (included in Microsoft 365 E3/E5) 'OAuth consent grant' attack simulation template, which replicates this campaign's exact delivery mechanism without requiring third-party budget. Publish the malicious AppId values and redirect URI patterns identified in this incident as internal threat intelligence to your email security gateway's URL reputation feed.

Evidence: Compile the post-incident evidence package: (1) the full timeline of consent grant events correlated with the AI-generated phishing email delivery timestamps from Exchange Online Message Trace logs to establish the attacker's operational tempo; (2) the complete list of Microsoft Graph API calls made by the malicious application during its authorized access window, extracted from 'resourceId' and 'scope' fields in sign-in logs, documenting what data was actually accessed (mail, files, contacts) across all 344 affected tenant-equivalent accounts in your organization; (3) the before-and-after Conditional Access policy exports demonstrating the control gap that permitted OAuth flows without MFA step-up — this document package supports both the internal lessons-learned review and any regulatory breach notification obligations under HIPAA (if healthcare sector) or state PII notification laws.

Detection Guidance

Primary detection surface is Microsoft Entra ID audit logs and sign-in logs. Query for the event 'Consent to application' (AuditLogs where ActivityDisplayName == 'Consent to application') and filter for user-initiated consents granted to applications not in your approved inventory. In Microsoft Sentinel, the query table 'AuditLogs' with OperationName 'Add OAuth2PermissionGrant' surfaces delegated permission grants. Cross-reference application client IDs against your known-good application registry. Secondary indicator: Defender for Cloud Apps will alert on OAuth app policy violations if app governance policies are configured; look for alerts under 'OAuth App Policies' for apps requesting high-privilege delegated scopes (Mail.Read, Calendars.ReadWrite, Files.ReadWrite.All). Behavioral IOC: user account showing resource access activity (Exchange, SharePoint, OneDrive) via an application client ID the user has never previously used, particularly from a cloud-only IP with no browser session correlation. No public IOC list (domains, IPs, hashes) has been released as of the item's source date (2026-03-02/03); check the Microsoft Security Blog post directly for any updates to published indicators.

Indicators of Compromise

Type	Value	Context	Confidence
URL	`No specific IOCs published as of 2026-03-03`	The Microsoft Security Blog and The Register coverage did not release specific attacker-controlled domains, IPs, or application IDs as of the item source dates. Monitor the Microsoft Security Blog post (2026-03-02) for IOC updates. Do not treat absence of listed IOCs as absence of compromise — detection relies on behavioral and consent-log analysis.	LOW

Framework Mappings

MITRE-ATTACK

- **T1566.002** — Spearphishing Link
- **T1566** — Phishing
- **T1528** — Steal Application Access Token
- **T1539** — Steal Web Session Cookie
- **T1556.006** — Multi-Factor Authentication
- **T1078** — Valid Accounts
- **T1550.001** — Application Access Token
- **T1078.004** — Cloud Accounts

NIST-800-53R5

- **AT-2** — Literacy Training and Awareness
- **SC-7** — Boundary Protection
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **SI-8** — Spam Protection
- **CA-7** — Continuous Monitoring
- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **SC-23** — Session Authenticity
- **SI-10** — Information Input Validation
- **IA-8** — Identification and Authentication (Non-Organizational Users)

OWASP-TOP10-2021

- **A01:2021** — Broken Access Control
- **A07:2021** — Identification and Authentication Failures

CIS-V8

- **16.10**
- **6.3**
- **6.4**
- **6.5**
- **14.2** — Train Workforce Members to Recognize Social Engineering Attacks

SOC2-TSC

- **CC6.1** — The entity implements logical access security software, infrastructure, and architectures over protected information assets

HIPAA-SECURITY

- **164.312(d)** — Person or Entity Authentication
- **164.308(a)(5)(i)** — Security Awareness and Training

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.34** — Privacy and protection of personal information
- **A.5.23** — Information security for use of cloud services

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1566.002	Spearphishing Link	Initial-Access
T1566	Phishing	Initial-Access
T1528	Steal Application Access Token	Credential-Access
T1539	Steal Web Session Cookie	Credential-Access
T1556.006	Multi-Factor Authentication	Credential-Access
T1078	Valid Accounts	Defense-Evasion
T1550.001	Application Access Token	Defense-Evasion
T1078.004	Cloud Accounts	Defense-Evasion

Sources

Source	URL	Tier
OAuth redirection abuse enables phishing and malware delivery	https://www.microsoft.com/en-us/security/blog/2026/03/02/oauth-redi...	T1
OAuth Phishing Attacks Exploit Microsoft 365 Accounts: Key Risks	https://mojoauth.com/news/oauth-phishing-attacks-exploit-microsoft-...	T3
Microsoft OAuth scams abuse redirects for malware delivery	https://www.theregister.com/2026/03/03/microsoft_oauth_scams/	T3
OAuth Vulnerabilities Every Security Team Should Know	https://www.obsidiansecurity.com/blog/oauth-vulnerabilities-securit...	T3
Hackers target compromised Microsoft Entra accounts in campaigns ...	https://www.reddit.com/r/technews/comments/1r9l2yn/hackers_target_c...	T3

Generated 2026-03-29 18:33 UTC by TJS Security Command Center