INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-03-29 18:36 UTC

Legitimate Microsoft Infrastructure Weaponized in Callback Phishing Campaign Bypassing All Email Authentication

THREAT CAMPAIGN | **HIGH** | CVSS 7.5

SCC Item ID	SCC-CAM-2026-0071
Type	Threat Campaign
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Microsoft Azure Monitor alert notification system; Microsoft 365 email recipients; azure-noreply@microsoft.com sending infrastructure
Published	2026-03-21

Executive Summary

Threat actors are abusing Microsoft Azure Monitor's legitimate alert notification system to send phishing emails from the genuine azure-noreply@microsoft.com address, allowing messages to pass SPF, DKIM, and DMARC authentication checks. Corporate users are the primary targets; lures impersonate Microsoft billing notices citing a fabricated $389.90 Windows Defender charge and direct recipients to call attacker-controlled numbers, a technique designed to bypass email security controls that rely on authentication validation alone. If successful, follow-on objectives include credential theft, payment fraud, and remote access tool deployment, posing direct financial and operational risk to affected organizations.

Technical Analysis

This campaign exploits a legitimate service feature rather than a software vulnerability, no CVE is assigned. Threat actors configure Azure Monitor alert rules to trigger notification emails sent by Microsoft's own azure-noreply@microsoft.com infrastructure, causing the messages to pass SPF, DKIM, and DMARC checks natively. The attack technique is classified as TOAD (Telephone-Oriented Attack Delivery), specifically callback phishing (MITRE T1566.004, T1598.003). Recipients are socially engineered via impersonation of Microsoft billing communications (T1656) with fabricated $389.90 Windows Defender subscription charges, then instructed to call attacker-controlled phone numbers. Assessed follow-on objectives include credential harvesting (T1078), RAT deployment via malicious user execution (T1204.001), and financial fraud (T1657). Because email originates from authentic Microsoft infrastructure, perimeter controls validating only sender authentication are ineffective. Relevant weaknesses: CWE-345 (Insufficient Verification of Data Authenticity),

CWE-20 (Improper Input Validation). No patch is available; this is a service abuse scenario. Attribution: unknown as of reporting date.

Action Checklist

1. Step 1, Immediate: Alert security awareness teams and issue an internal advisory warning corporate users of phishing emails originating from azure-noreply@microsoft.com referencing Windows Defender billing charges; instruct recipients to not call any phone numbers in unsolicited Microsoft billing emails.
2. Step 2, Detection: Search email gateway and Microsoft 365 Defender logs for inbound messages from azure-noreply@microsoft.com containing billing-related subject lines or body text referencing Windows Defender subscription charges and phone numbers; flag and quarantine pending review.
3. Step 3, Detection: Review Azure Monitor alert notification configurations in your tenant for alert rules created by unrecognized identities or containing external email recipients not associated with your organization's domains.
4. Step 4, Assessment: Identify any users who may have received and interacted with these messages; prioritize accounts that may have called attacker-provided numbers, submitted credentials, or installed software at caller instruction for immediate incident triage.
5. Step 5, Communication: Notify affected users and escalate confirmed or suspected callback interactions to incident response; report abuse to Microsoft via abuse@microsoft.com and the Microsoft 365 abuse reporting channel, or submit directly in Microsoft 365 Defender under Email > Collaboration > Submitted for review.
6. Step 6, Long-term: Evaluate supplementary email filtering rules that flag messages from Microsoft infrastructure containing phone numbers and billing language not matching known Microsoft billing patterns; review security awareness training to explicitly cover TOAD/callback phishing scenarios.

IR / Forensic Enrichment

Triage Priority	IMMEDIATE
Escalation Criteria	Escalate to CISO and external IR firm immediately if any user confirms callback interaction, credential submission, or software installation at attacker instruction, or if Azure Monitor abuse indicates persistent attacker access to alert rule creation (suggests compromised service principal or overly permissive RBAC).
Recovery Notes	Post-containment: (1) Reset passwords and review MFA settings for all users who interacted with phishing or called attacker numbers; audit their Azure AD sign-in logs and privileged role assignments for 30 days to detect lateral movement. (2) Review and lock down Azure Monitor alert rule creation permissions — restrict to security-approved service principals and require approval workflow for email action groups; audit all alert rules created by non-privileged accounts in last 120 days. (3) Implement supplementary email filtering rules and conduct security awareness refresh; measure awareness response rate and track repeat reporting of similar lures to validate training effectiveness.

Forensic Artifacts	Email gateway transaction logs and Message Trace exports (full headers, recipient lists, delivery status, 30+ day window) | Azure Activity Log (alert rule creation/modification, service principal activity, role assignment changes, 90-120 day window) | Azure AD Sign-in logs and Audit logs (suspicious account activity, failed logins from callback-era timeframe, service principal credential usage) | Help desk ticketing system (callback-related support requests, billing inquiries, password reset spikes correlating to Message Trace send dates) | User endpoint forensics (Windows Event Log 4625/4648/4688, VPN/MFA logs, browser history for credential entry or software downloads from attacker-provided links)

Per-Action IR Details

Step 1 — Immediate: Alert security awareness teams and issue an internal advisory warning corporate users of phishing emails originating from azure-noreply@microsoft.com referencing Windows Defender billing charges; instruct recipients to not call any phone numbers in unsolicited Microsoft billing emails.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2.1 (Preparation phase — user awareness and training)

Controls: NIST 800-53 AT-2 (Security Awareness Training), NIST 800-53 AT-3 (Role-Based Security Training), CIS 6.3 (Address unauthorized software)

Compensating: Draft alert email from CISO/security lead; distribute via internal email list, Slack, Teams, or intranet banner. Include: sender address (azure-noreply@microsoft.com), fake charge amount ($389.90), and instruction to forward suspicious messages to security@[yourdomain] instead of calling. If no email distribution tool exists, contact department heads directly with talking points for their teams.

Evidence: Capture the date/time of advisory distribution, distribution method (email headers, message ID, delivery receipts), and any prior phishing reports to this address before advisory went out (check helpdesk tickets, email security logs, or user-reported mailbox items). Document baseline to measure response velocity post-advisory.

Step 2 — Detection: Search email gateway and Microsoft 365 Defender logs for inbound messages from azure-noreply@microsoft.com containing billing-related subject lines or body text referencing Windows Defender subscription charges and phone numbers; flag and quarantine pending review.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2.4 (Detection and Analysis — source identification and log review)

Controls: NIST 800-53 AU-6 (Audit Review, Analysis, and Reporting), NIST 800-53 SI-4 (Information System Monitoring), CIS 8.1 (Collect and analyze logs)

Compensating: If no SIEM: (1) Export email gateway logs (Postfix/Sendmail: grep 'azure-noreply' /var/log/mail.log or equivalent; Exchange: Export-Message from PowerShell with -Sender filter). (2) Search O365 Message Trace via admin portal: FilterBy Sender = azure-noreply@microsoft.com, last 30 days. (3) Manually grep message bodies for strings: '\$389.90', 'Windows Defender', 'subscription', 'call', phone number regex (\d{3}-\d{3}-\d{4}). (4) Create CSV export with To, From, Subject, Received date, message ID. Quarantine by creating mail flow rule: if From = azure-noreply AND (Subject contains 'billing' OR Body contains 'Windows Defender') → move to quarantine folder or disable delivery.

Evidence: Before running detection query, preserve: (1) Complete email gateway transaction logs (24-48 hours before search start date to establish baseline of legitimate azure-noreply traffic). (2) Full message headers and body content of any matched messages (do not delete; store in isolated folder). (3) Recipient list (To: field) for each flagged message — this becomes your user notification list. (4) Email gateway filter/rule configuration snapshots (time-stamped) showing what rules were active during the campaign window. (5) Message trace timestamps and delivery status (delivered, bounced, quarantined) to understand scope.

Step 3 — Detection: Review Azure Monitor alert notification configurations in your tenant for alert rules created by unrecognized identities or containing external email recipients not associated with your organization's domains.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2.3 (Detection and Analysis — scope determination and cloud configuration review)

Controls: NIST 800-53 AC-2 (Account Management), NIST 800-53 CA-7 (Continuous Monitoring), CIS 1.5 (Account creation and configuration control)

Compensating: If no Azure portal access/limited permissions: (1) Request Azure Activity Log export from cloud admin via PowerShell: Get-AzActivityLog -ResourceGroupName [RG] -StartTime (Get-Date).AddDays(-90) | Export-CSV. (2) Filter for 'Microsoft.Insights/alertrules/write' operations and 'Microsoft.AlertsManagement' provider events. (3) For each alert rule, query: Get-AzMetricAlertV2 | Select Name, AlertRuleResourceId, Description, Enabled, Actions (which contains action group/email list). (4) Cross-reference email recipients against corporate domain allowlist (user @yourdomain.com accounts only); flag any external email addresses or distribution lists not in approved contact list. (5) If no direct Azure access, request tenant security team to run report and provide CSV of all alert rules + action group recipients created/modified in last 120 days.

Evidence: Before querying Azure Monitor: (1) Snapshot current Azure AD/Entra ID account inventory and role assignments (Privileged Identity Management reports — who has contributor/owner on Azure subscriptions, last modified). (2) Export full Azure Activity Log covering 90-120 days prior (includes Create/Modify alerts, account logins, role changes). (3) Document baseline: list of known service principals, managed identities, and user accounts that legitimately create alert rules. (4) Capture alert rule definitions (JSON payload) for any rules created/modified by unfamiliar service accounts — these become evidence of lateral movement or privilege escalation. (5) Timestamp Azure Monitor diagnostic settings to verify who modified log retention or forwarding rules.

Step 4 — Assessment: Identify any users who may have received and interacted with these messages; prioritize accounts that may have called attacker-provided numbers, submitted credentials, or installed software at caller instruction for immediate incident triage.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2.6 (Detection and Analysis — impact assessment and user interaction forensics)

Controls: NIST 800-53 IR-4 (Incident Handling), NIST 800-53 AU-12 (Audit Generation — log collection from user endpoints), CIS 8.7 (Detect credential access attempts)

Compensating: If no EDR/MDE: (1) Distribute incident response questionnaire to all Message Trace recipients: 'Did you receive a billing email from azure-noreply@microsoft.com? Did you click links, call a phone number, or enter credentials?' Use Google Forms or email survey for response tracking. (2) For callback triage, cross-reference Message Trace recipient list against (a) help desk tickets mentioning 'Microsoft billing', 'Windows Defender charge', or 'account verification'; (b) VPN/MFA logs showing failed/unusual authentication attempts (last 7-14 days, especially from recipients' IP addresses post-email date); (c) Windows Event Log 4648 (explicit credential use) or 4625 (failed login) on recipient workstations (query local logs via remote PowerShell if available). (3) If software installation suspected: query Application event logs for MSI installations or program execution logs (PowerShell Constrained Language Policy logs or manually requested from users). (4) Prioritize triage: accounts with VPN logins from non-corporate IPs, failed auth spikes, or help desk calls matching Message Trace send dates.

Evidence: Before user outreach: (1) Preserve complete Message Trace data (list of all recipients) — this is your interview target list. (2) Capture help desk ticket system queries (last 30 days) searching for keywords: 'billing', 'charge', 'Windows Defender', 'Microsoft', 'phone number', 'credit card', 'account verification'. (3) Export authentication logs (Azure AD sign-in logs, VPN access logs, MFA challenge logs) for all recipients covering 7 days post-first email detection to establish interaction baseline. (4) Snapshot endpoint file system and process execution (if EDR available: process creation events, file modifications, network connections for recipients' machines; if manual: request users run Get-Process and Get-EventLog cmdlets). (5) Document call center records or phone bill metadata if available (inbound/outbound call logs to attacker phone numbers — correlate with recipient account IDs).

Step 5 — Communication: Notify affected users and escalate confirmed or suspected callback interactions to incident response; report abuse to Microsoft via msrc.microsoft.com and the Microsoft 365 abuse reporting channel.

NIST Phase: Containment

Reference: NIST 800-61r3 §3.3 (Containment) and §4 (Post-Incident Activities — reporting and stakeholder communication)

Controls: NIST 800-53 IR-6 (Incident Reporting), NIST 800-53 IR-7 (Incident Response Assistance), CIS 5.2 (Incident notification and escalation)

Compensating: If no formal IR communication plan: (1) Draft user notification email from CISO/Security Officer (not IT Help Desk — adds authority): include list of Message Trace recipients, time window, description of phishing lure, and action they should take (password reset, MFA review, monitor account for unauthorized activity). (2) For callback interactions: create incident ticket per user (link to Message Trace record and help desk ticket if applicable); assign to IR lead for triage interview. (3) Microsoft abuse reporting: forward sample phishing message headers + body to abuse@microsoft.com AND use Microsoft 365 admin center > Threat Management > Review > Threat Explorer to mark messages as phishing (automated reporting). Also report to MSRC via https://msrc.microsoft.com/report/vulnerability — document Azure Monitor configuration abuse case (attacker used legitimate Azure notification service as vector). (4) Document all communications (dates, recipients, content snapshots, response tracking) for post-incident review.

Evidence: Before user notification: (1) Finalize list of affected users (Message Trace recipients). (2) Prepare evidence package: sample phishing email (MIME headers + body, sanitized of any user PII), Message Trace output, alert rule definitions from Azure Monitor showing attacker-created rules. (3) Screenshot Microsoft 365 Defender detections or email gateway quarantine records (if available). (4) Prepare internal incident ticket template with links to Message Trace, Azure Activity Log findings, and help desk correlations. (5) Capture initial Microsoft abuse report ticket number (for follow-up reference and SLA tracking).

Step 6 — Long-term: Evaluate supplementary email filtering rules that flag messages from Microsoft infrastructure containing phone numbers and billing language not matching known Microsoft billing patterns; review security awareness training to explicitly cover TOAD/callback phishing scenarios.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4.2 (Post-Incident Activities — recommendations and lessons learned); NIST 800-53 IR-5 (Improvement)

Controls: NIST 800-53 SI-4 (Information System Monitoring — rule tuning), NIST 800-53 AT-2 (Security Awareness and Training — curriculum updates), CIS 6.3 (Application control and phishing training), CIS 8.5 (Alert tuning to reduce false positives)

Compensating: If no SIEM/email gateway rule engine: (1) Create mail flow rules in Exchange or O365: Rule 1 — If From = 'microsoft.com' OR 'microsoft-noreply' domain AND (Body contains phone number regex OR Subject contains 'billing' OR 'charge') AND NOT (From = 'billing@microsoft.com' OR other whitelisted billing senders) → Add disclaimer header 'UNVERIFIED: Microsoft emails do not request payment by phone call.' Rule 2 — If From contains 'noreply' AND Body contains phone number regex → flag for manual review. (2) For awareness training: create 1-page scenario card (sample phishing email + correct response steps), distribute to all staff quarterly. Include: 'Legitimate Microsoft never sends unsolicited billing emails with phone numbers; legitimate Microsoft support requests go through Service Health Dashboard, not email callbacks.' (3) Conduct tabletop with security team, helpdesk, and finance to establish 'know your Microsoft billing patterns' checklist (e.g., actual billing senders, legitimate subject lines, expected charge categories). (4) Add to annual security training: 'Callback phishing / TOAD (Toll Fraud, Overpayment, Abuse Dispatch) scenario quiz.'

Evidence: Before implementing rules: (1) Establish baseline of legitimate Microsoft infrastructure emails your organization actually receives (billing senders, alert senders, notification senders) — export sample emails over 60-90 days to build whitelist. (2) Audit existing email filtering rules to identify conflicts or overlaps (rules that might block legitimate Microsoft email). (3) Document current awareness training curriculum — identify gaps in callback phishing / TOAD coverage. (4) Prepare metrics: number of flagged messages, false positive rate (legitimate Microsoft email caught by new rule). Set acceptable false positive threshold (e.g., < 0.5%) before deploying broadly. (5) Coordinate with service owners (finance, helpdesk, Azure admins) to notify them of new filtering rules so they understand why some Microsoft notifications might be delayed.

Detection Guidance

Email gateway and Microsoft 365 Defender: Query inbound mail where sender domain is microsoft.com or azure.microsoft.com, subject or body contains terms such as 'Windows Defender', 'subscription', 'renewal', or dollar amounts, and body contains phone number patterns (e.g., regex: \+?1?[\s.-]?\(?\d{3}\)?[\s.-]?\d{3}[\s.-]?\d{4}). Note: Phone number regex is tuned for US formats; adjust for your region or broaden to catch international patterns (+country codes, variable digit counts). Authentic Microsoft billing communications do not instruct recipients to call a phone number. Azure portal, audit logs: Review Microsoft.Insights/alertRules Create and Update events in Azure Activity Log for rules configured with external action groups targeting email addresses outside your tenant. Look for alert rules created by service principals or users with no prior administrative history. Behavioral indicator: Any user who received one of these emails and subsequently made an outbound call to an unrecognized number, visited an unfamiliar URL, or installed software should be treated as a potential incident. Endpoint telemetry: Monitor for remote access tool (RAT) installation or execution following any reported callback interaction, common tooling in TOAD follow-on stages includes AnyDesk, TeamViewer, and similar remote access software (T1219). IOC confidence is low given no specific attacker infrastructure has been publicly confirmed as of reporting date; treat phone numbers in any unsolicited azure-noreply billing email as malicious by default.

Indicators of Compromise

Type	Value	Context	Confidence
DOMAIN	`azure-noreply@microsoft.com (sending address)`	Legitimate Microsoft sending address abused as delivery vector; presence alone does not confirm malicious activity — context of message content is required for triage	LOW
URL	`No confirmed attacker-controlled URLs published as of reporting date`	Campaign relies on phone callback rather than embedded URLs; no specific attacker infrastructure URLs confirmed in available sources	LOW
DOMAIN	`No confirmed attacker-controlled domains published as of reporting date`	Attribution and infrastructure remain unconfirmed; monitor threat intelligence feeds for updates	LOW

Framework Mappings

MITRE-ATTACK

- **T1566.004** — Spearphishing Voice
- **T1598.003** — Spearphishing Link
- **T1204.001** — Malicious Link
- **T1598** — Phishing for Information
- **T1656** — Impersonation

- **T1657** — Financial Theft
- **T1199** — Trusted Relationship
- **T1219** — Remote Access Tools
- **T1078** — Valid Accounts

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **SI-7** — Software, Firmware, and Information Integrity
- **SI-10** — Information Input Validation
- **AT-2** — Literacy Training and Awareness
- **SI-4** — System Monitoring

OWASP-TOP10-2021

- **A08:2021** — Software and Data Integrity Failures
- **A03:2021** — Injection

CIS-V8

- **2.5**
- **16.10**
- **6.3** — Require MFA for Externally-Exposed Applications
- **14.2** — Train Workforce Members to Recognize Social Engineering Attacks
- **8.2** — Collect Audit Logs

ISO-27001-2022

- **A.8.26** — Application security requirements
- **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

HIPAA-SECURITY

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

SOC2-TSC

- **CC6.1** — Logical access security software, infrastructure, and architectures

NIST-CSF-2

- **DE.CM-01** — Networks and network services are monitored

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1566.004	Spearphishing Voice	Initial-Access
T1598.003	Spearphishing Link	Reconnaissance
T1204.001	Malicious Link	Execution
T1598	Phishing for Information	Reconnaissance
T1656	Impersonation	Defense-Evasion
T1657	Financial Theft	Impact
T1199	Trusted Relationship	Initial-Access
T1219	Remote Access Tools	Command-And-Control
T1078	Valid Accounts	Defense-Evasion

Sources

Source	URL	Tier
Security News	https://www.bleepingcomputer.com/news/security/microsoft-azure-moni...	T3
is an email from Microsoft Azure legit?	https://learn.microsoft.com/en-us/answers/questions/5790345/is-an-e...	T1
Possible phishing from Microsoft Azure and Microsoft Cloud.	https://learn.microsoft.com/en-us/answers/questions/5790477/possibl...	T1
I received an email communication from Microsoft Azure and would ...	https://learn.microsoft.com/en-us/answers/questions/5815880/i-recei...	T1
Microsoft Azure email - not sure if phishing - Reddit	https://www.reddit.com/r/phishing/comments/1rg6wu0/microsoft_azure_...	T3

DISCLAIMER

This intelligence report is produced by Tech Jacks Solutions Security Command Center (SCC) for informational purposes only. It does not constitute professional security advice, legal counsel, or an incident response engagement. The information herein is derived from publicly available sources and AI-assisted analysis; while every effort is made to ensure accuracy, Tech Jacks Solutions makes no warranties regarding completeness or timeliness. Organizations should conduct their own validation and consult qualified security professionals before taking action based on this report. Tech Jacks Solutions is not liable for any damages resulting from the use of this information.

Generated 2026-03-29 18:36 UTC by TJS Security Command Center