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

Threat Actors Weaponize Azure Monitor Alert Descriptions to Deliver Authenticated Callback Phishing

THREAT CAMPAIGN | **HIGH** | CVSS 7.5

SCC Item ID	SCC-CAM-2026-0075
Type	Threat Campaign
Severity	HIGH
CVSS Base Score	7.5
Affected Products	Microsoft Azure Monitor, Microsoft 365, azure-noreply@microsoft.com email infrastructure
Published	2026-03-21

Executive Summary

Threat actors are injecting malicious content into legitimate Microsoft Azure Monitor alert emails sent from azure-noreply@microsoft.com, causing them to pass all standard email authentication checks (SPF, DKIM, DMARC). Recipients receive what appears to be an official Microsoft Security Team message containing a fabricated $389.90 Windows Defender billing charge and instructions to call an attacker-controlled phone number. Any organization using Azure Monitor alert notifications is potentially exposed; the primary business risks are credential theft, unauthorized network access, and financial fraud via social engineering.

Technical Analysis

Attack vector: Threat actors create or modify Azure Monitor alert rules, inserting arbitrary content into the alert description field. When Azure Monitor fires the alert, Microsoft's own notification pipeline delivers the crafted message from azure-noreply@microsoft.com with valid SPF, DKIM, and DMARC authentication, bypassing email security gateways and anti-phishing controls that rely on sender authentication. The lure impersonates the Microsoft Security Team and presents a fabricated billing charge to coerce recipients into calling attacker-controlled numbers (callback phishing / TOAD, Telephone-Oriented Attack Delivery). Over the phone, attackers likely pursue credential harvesting, remote access tool installation (T1219), or direct financial fraud. No CVE is assigned; the behavior exploits a legitimate platform feature, not a vulnerability in the traditional sense. Relevant CWEs: CWE-284 (Improper Access Control, alert field accepts arbitrary content), CWE-20 (Improper Input Validation, no sanitization of alert description content before inclusion in outbound email), CWE-940 (Improper Verification of Source of a Communication Channel, recipients cannot distinguish legitimate from

weaponized alert emails by sender authentication alone). MITRE ATT&CK coverage: T1566 (Phishing), T1566.004 (Spearphishing Voice), T1036 / T1036.005 (Masquerading), T1598 (Phishing for Information), T1204.002 (User Execution: Malicious File, applicable if phone call leads to remote tool delivery), T1219 (Remote Access Software), T1078 (Valid Accounts), T1059 (Command and Scripting Interpreter, post-access), T1199 (Trusted Relationship, abuse of Microsoft's trusted sending infrastructure). No patch is currently available; Microsoft has not issued a platform-level fix as of reporting. Mitigation is defensive and procedural.

Action Checklist

1. Step 1, Immediate: Alert all Azure administrators and help desk staff that azure-noreply@microsoft.com alert emails may contain attacker-injected content; instruct staff not to call any phone number presented in an Azure Monitor alert email without out-of-band verification.
2. Step 2, Immediate: Review Azure Monitor alert rules across all subscriptions for unauthorized or unfamiliar alert configurations, paying particular attention to alert descriptions containing phone numbers, billing language, or impersonation of Microsoft support teams.
3. Step 3, Detection: Query Microsoft 365 mail flow logs and email security platform logs for inbound messages from azure-noreply@microsoft.com containing keywords: 'Windows Defender', 'billing', 'charge', '$', 'call', '1-8', or phone number patterns, flag any matches for analyst review.
4. Step 4, Assessment: Audit Azure Monitor alert rule creation and modification logs (Azure Activity Log, resource type: microsoft.insights/alertrules) for changes made by unfamiliar identities or service principals over the past 90 days; cross-reference with Azure AD sign-in logs for anomalous access.
5. Step 5, Communication: Notify end users and executives via a verified internal channel that this campaign is active; provide examples of the lure message and clear instructions for reporting suspected callback phishing attempts to the security team.
6. Step 6, Long-term: Evaluate Azure Monitor alert notification configurations to restrict who can create or modify alert rules with broad distribution; consider implementing Azure Policy to enforce alert description content standards or limit external email targets for alert notifications.
7. Step 7, Long-term: Update security awareness training to include TOAD / callback phishing as a distinct threat pattern, with specific guidance on Microsoft billing communications (Microsoft does not issue billing charges via Azure Monitor alert emails).

IR / Forensic Enrichment

Triage Priority	IMMEDIATE
Escalation Criteria	Escalate to IR management and notify CISO if more than 5 users report receiving the phishing email, if any user reports calling the attacker number, or if unauthorized alert rules are confirmed in production subscriptions; consider external IR firm engagement if alert rules are discovered on critical infrastructure or healthcare subscriptions.

Recovery Notes	After unauthorized alert rules are removed and email filters are deployed, conduct a full review of Azure Monitor alert rule backups and compare against the current state to identify any rules created during the attack window that were subsequently deleted. Audit all Azure services that triggered alerts during the 90-day lookback period to ensure alert content was not maliciously modified. Provide affected users with verification guidance: direct them to log into the Azure portal, navigate to Monitor > Alerts, and compare alert rule descriptions against any emails they received to confirm legitimacy.
Forensic Artifacts	Azure Activity Log entries for resource type 'microsoft.insights/alertrules' (90-day lookback): operations CreateOrUpdateMetricAlertRule, DeleteMetricAlertRule, and ModifyAlertRule with Caller identity and timestamp | Azure AD sign-in logs for all identities that created or modified alert rules, including IP address, user agent, device compliance, and MFA status | Exchange Online mail flow logs and message headers for all inbound messages from azure-noreply@microsoft.com (30-day lookback minimum) with body content containing keywords: 'billing', 'charge', 'Defender', 'call', phone number patterns | Azure Monitor alert rule export (JSON or PowerShell Get-AzMetricAlertRuleV2 output) showing rule name, description, creator, creation/modification timestamps, and notification targets for all subscriptions | Email security platform threat logs (if available) showing sender reputation, DKIM/SPF/DMARC pass/fail status, URL rewrite events, and phishing confidence scores for messages from azure-noreply@microsoft.com

Per-Action IR Details

Step 1 — Immediate: Alert all Azure administrators and help desk staff that azure-noreply@microsoft.com alert emails may contain attacker-injected content; instruct staff not to call any phone number presented in an Azure Monitor alert email without out-of-band verification.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2.1 (organization preparation) and §2.3 (tools and resources)

Controls: NIST IR-1 (incident response policy), NIST AT-2 (security awareness training), CIS 17.1 (communication templates for security incidents)

Compensating: Create and distribute a one-page alert via email, Slack, or internal wiki with a screenshot of the lure and the rule: 'Do not call any phone number in an Azure Monitor alert email. Instead, verify the alert by logging into your Azure portal directly and checking the alert rule.' Document distribution in a shared spreadsheet with timestamps and recipient confirmation.

Evidence: Capture the original phishing email headers (Message-ID, X-Originating-IP, X-Mailer, Received chain) before forwarding to security team. Preserve the alert email in a dedicated folder for later forensic analysis. Document the timestamp when staff were notified to establish timeline of awareness.

Step 2 — Immediate: Review Azure Monitor alert rules across all subscriptions for unauthorized or unfamiliar alert configurations, paying particular attention to alert descriptions containing phone numbers, billing language, or impersonation of Microsoft support teams.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2.2 (initial analysis and hypothesis) and §3.2.3 (event documentation)

Controls: NIST CA-7 (continuous monitoring), NIST SI-4 (information system monitoring), CIS 8.2 (configuration baseline and integrity checking)

Compensating: Export all Azure Monitor alert rules via Azure PowerShell: `Get-AzMetricAlertRuleV2 -ResourceGroupName '*' | Select-Object Name, Description, CreatedTime, LastModifiedTime | Export-Csv alerts.csv`. Manually review the Description field for keywords: phone numbers (regex: \d{3}-\d{3}-\d{4}), 'billing', 'charge', 'Defender', 'support team', 'confirm', 'verify payment'. Cross-reference alert creators against known administrators in a spreadsheet.

Evidence: Capture Azure Activity Log entries for resource type 'microsoft.insights/alertrules' with operations 'CreateOrUpdateMetricAlertRule' and 'DeleteMetricAlertRule' from the past 90 days. Export via: `Get-AzActivityLog -ResourceGroupName '*' -ResourceType 'microsoft.insights/alertrules' -StartTime (Get-Date).AddDays(-90)`. Document the exact alert rule ID, description text, and creator identity for each suspicious rule.

Step 3 — Detection: Query Microsoft 365 mail flow logs and email security platform logs for inbound messages from azure-noreply@microsoft.com containing keywords: 'Windows Defender', 'billing', 'charge', '$', 'call', '1-8', or phone number patterns — flag any matches for analyst review.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2.1 (detection and analysis overview) and §3.2.4 (event correlation)

Controls: NIST SI-4(1) (system monitoring with organization-defined monitoring tools), NIST AU-6(9) (audit record review with data analytics), CIS 8.8 (e-mail and web browser protections)

Compensating: Query Exchange Online Mail Flow logs using: `$logs = Get-MessageTrace -SenderAddress 'azure-noreply@microsoft.com' -StartDate (Get-Date).AddDays(-30) -EndDate (Get-Date); $logs | Select-Object * | Export-Csv azure_alert_emails.csv`. Use a text editor or grep to search the exported body/subject for patterns: `grep -iE '(Windows Defender|billing|charge|\$|call|1-8[0-9]{2})' azure_alert_emails.csv`. Document all matches with recipient, timestamp, and subject line.

Evidence: Preserve the full email headers (Internet Message ID, received timestamps, DKIM/SPF/DMARC results) and message body for each flagged email. Export message headers using: `Get-Message -Identity '' | Export-CliXml message_headers.xml`. If email security platform is present, export threat logs with sender, recipient, keywords detected, and timestamp. Create a list of all recipient addresses and forward to notification step.

Step 4 — Assessment: Audit Azure Monitor alert rule creation and modification logs (Azure Activity Log, resource type: microsoft.insights/alertrules) for changes made by unfamiliar identities or service principals over the past 90 days; cross-reference with Azure AD sign-in logs for anomalous access.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2.5 (impact analysis) and §3.4.1 (evidence gathering and handling)

Controls: NIST AU-2 (audit events), NIST AU-12 (audit generation), NIST IA-4 (identifier management), CIS 6.2 (managed logging and log alerting)

Compensating: Export Azure Activity Log: `Get-AzActivityLog -ResourceType 'microsoft.insights/alertrules' -StartTime (Get-Date).AddDays(-90) | Select-Object EventTimestamp, Caller, OperationName, ResourceId, StatusCode, Properties | Export-Csv activity_log.csv`. Export Azure AD sign-in logs: `Export-AzureADSignInLog -StartDate (Get-Date).AddDays(-90) -EndDate (Get-Date) | Export-Csv signin_log.csv`. Cross-reference 'Caller' identity from activity log against a list of known administrators and approved service principals. Flag any entries where Caller is unfamiliar or sign-in shows unusual time, location, or IP address.

Evidence: Preserve the full activity log entry including EventTimestamp, Caller (user principal name and object ID), OperationName, Resource ID, and Properties JSON. For each flagged Caller, capture Azure AD sign-in logs showing all login attempts from that identity during the 90-day window, including IP address, device, browser, and MFA status. Document any gaps in audit logs (retention limits, deletions).

Step 5 — Communication: Notify end users and executives via a verified internal channel that this campaign is active; provide examples of the lure message and clear instructions for reporting suspected callback phishing attempts to the security team.

NIST Phase: Containment

Reference: NIST 800-61r3 §2.3.4 (providing user awareness training) and §3.3.1 (containment planning)

Controls: NIST AT-2 (security awareness and training), NIST IR-4 (incident handling), CIS 17.8 (incident response process)

Compensating: Draft a templated security alert email with: threat name, a sanitized screenshot of the lure (phone number redacted), key indicators (billing language, request to call), and a reporting mechanism (email to security-team@company.com or phone number for verified security desk). Distribute via email from the CIO/Security

Officer with confirmation of read receipts. Post the same alert in a pinned message in Slack #general and on the internal security wiki. Log distribution timestamp and recipient count.

Evidence: Capture and archive the original phishing email and any variant samples discovered during Step 3. Document the exact lure text, phone number(s), and alert rule descriptions. Record the timestamp and audience of each notification sent. If end users report similar emails, capture those reports with headers and timestamps to assess campaign scope and evolution.

Step 6 — Long-term: Evaluate Azure Monitor alert notification configurations to restrict who can create or modify alert rules with broad distribution; consider implementing Azure Policy to enforce alert description content standards or limit external email targets for alert notifications.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4.1 (eradication) and §3.4.2 (eradication strategy)

Controls: NIST AC-2 (account management), NIST AC-3 (access enforcement), NIST CA-9 (internal system connections), CIS 5.3 (access control list management)

Compensating: Manually audit Azure Monitor alert rule creators using Azure portal: IAM > Roles > Custom roles > filter by operations containing 'alertrules'. Create a custom Azure role that denies 'microsoft.insights/alertrules/write' except for a named security team group. Assign this role at the subscription level. For alert descriptions, create a PowerShell script to scan all alert rules monthly and flag any containing regex patterns (phone numbers, billing keywords, 'call'). Store the script in Azure Automation Runbook and log findings to a CSV for manual review. Document the control in a runbook wiki page.

Evidence: Capture the current RBAC assignments for Azure Monitor Alert rule creator permissions. Document the custom role definition (JSON) and assignment scope. Record the baseline of all alert rules before and after the control is applied, including rule ID, creator, description, and modification history. Establish a baseline for future compliance audits.

Step 7 — Long-term: Update security awareness training to include TOAD / callback phishing as a distinct threat pattern, with specific guidance on Microsoft billing communications (Microsoft does not issue billing charges via Azure Monitor alert emails).

NIST Phase: Post Incident

Reference: NIST 800-61r3 §3.5.1 (post-incident activities) and NIST SP 800-50 (security awareness and training)

Controls: NIST AT-2 (security awareness and training), NIST AT-3 (role-based security training), CIS 17.7 (security awareness training program)

Compensating: Add a 5-minute module to annual security training covering: what TOAD (Toll Fraud and Authenticated Callback Phishing) is, how this campaign weaponizes legitimate Azure Monitor emails, the red flags (billing language, phone numbers, urgency), and the correct response (verify in Azure portal, do not call). Distribute via LMS or video link. Create a one-page quick reference card for help desk and post to the security intranet. Track completion rates and quiz scores. Update the company security policy to state: 'Microsoft never sends billing charges via alert emails. Always verify security alerts by logging into the Microsoft portal directly.'

Evidence: Document training module creation date, audience, completion metrics, and assessment scores. Capture a sample of the training material and quick reference card. Track any phishing simulation results related to this threat pattern before and after training deployment to measure awareness improvement.

Detection Guidance

Primary detection surface is Azure Activity Log and Microsoft 365 mail flow. Key indicators and queries: (1) Azure Activity Log, filter for operationName 'microsoft.insights/alertrules/write' or 'microsoft.insights/scheduledqueryrules/write' by identities outside your expected administrator set; look for creation or modification events from unfamiliar service principals or guest accounts. (2) Microsoft 365 / Exchange mail flow logs, search for messages where SenderAddress = 'azure-noreply@microsoft.com' AND

body or subject contains phone number patterns (e.g., regex: \b(1[-.]?\(?8[0-9]{2}\)?[-.]?[0-9]{3}[-.]?[0-9]{4})\b) or billing terms ('charge', 'invoice', 'Windows Defender', '$'). (3) Defender for Office 365 or equivalent, build a custom detection rule triggering on authenticated mail from azure-noreply@microsoft.com where body content matches social engineering keyword clusters. (4) Endpoint telemetry, if a user has already called a number and interacted with an attacker, look for: new remote access tool installations (AnyDesk, TeamViewer, ScreenConnect), LOLBin execution chains (mshta, wscript, cscript, powershell) spawned from browser or email client processes, and new scheduled tasks or services created within 30 minutes of an anomalous phone call. Behavioral indicator: any user-reported phone call to a 'Microsoft billing' number received via an Azure email alert is a confirmed TOAD attempt and warrants immediate endpoint investigation.

Indicators of Compromise

Type	Value	Context	Confidence
DOMAIN	`azure-noreply@microsoft.com (sending address)`	Legitimate Microsoft sending address weaponized by injecting attacker content into Azure Monitor alert descriptions; presence alone is not malicious — inspect email body for phone numbers and billing lures	HIGH
URL	`No attacker-controlled URLs reported in available sources`	Campaign relies on phone callback rather than embedded URLs; standard URL-based IOCs are not applicable for this vector	HIGH
PATTERN	`Fabricated charge: $389.90 Windows Defender billing`	Lure amount reported in BleepingComputer coverage; may vary across campaign variants — treat any Azure alert email referencing Defender billing as suspicious regardless of amount	MEDIUM

Framework Mappings

MITRE-ATTACK

- **T1199** — Trusted Relationship
- **T1566** — Phishing
- **T1204.002** — Malicious File
- **T1036** — Masquerading
- **T1219** — Remote Access Tools
- **T1566.004** — Spearphishing Voice
- **T1078** — Valid Accounts
- **T1598** — Phishing for Information
- **T1059** — Command and Scripting Interpreter
- **T1036.005** — Match Legitimate Resource Name or Location

NIST-800-53R5

- **AT-2** — Literacy Training and Awareness
- **CA-7** — Continuous Monitoring
- **SC-7** — Boundary Protection
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **SI-8** — Spam Protection
- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **CM-7** — Least Functionality
- **SI-7** — Software, Firmware, and Information Integrity
- **AC-3** — Access Enforcement
- **SI-10** — Information Input Validation

OWASP-TOP10-2021

- **A01:2021** — Broken Access Control
- **A03:2021** — Injection

CIS-V8

- **6.1**
- **6.2**
- **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

SOC2-TSC

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

HIPAA-SECURITY

- **164.312(a)(1)** — Access Control
- **164.312(d)** — Person or Entity Authentication
- **164.308(a)(5)(i)** — Security Awareness and Training

ISO-27001-2022

- **A.8.26** — Application security requirements
- **A.5.34** — Privacy and protection of personal information
- **A.5.23** — Information security for use of cloud services

NIST-CSF-2

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

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1199	Trusted Relationship	Initial-Access
T1566	Phishing	Initial-Access
T1204.002	Malicious File	Execution
T1036	Masquerading	Defense-Evasion
T1219	Remote Access Tools	Command-And-Control
T1566.004	Spearphishing Voice	Initial-Access
T1078	Valid Accounts	Defense-Evasion
T1598	Phishing for Information	Reconnaissance
T1059	Command and Scripting Interpreter	Execution
T1036.005	Match Legitimate Resource Name or Location	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
is an email from Microsoft Azure legit	https://learn.microsoft.com/en-us/answers/questions/5815890/is-an-e...	T1

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