



POST-INCIDENT REPORT

Business Email Compromise via Adversary-in-the-Middle Phishing

Compromise of [REDACTED]@[REDACTED], persistent token replay, and mass phishing campaign of September 16, 2026

Prepared for

[REDACTED]

Attn: [REDACTED]

[REDACTED]

Prepared by

IT Partner LLC - Microsoft Cloud Solution Provider

Mike Mackey, CEO | mym@itpartner365.com | +1-480-535-0365

Security operations: support@o365hq.com | <https://o365hq.com>

Incident reference

IT Partner tickets 16435, 16704, 16727

Microsoft Defender XDR incidents 25188 and 25347

Microsoft 365 tenant [REDACTED]

Incident window

August 19, 2026 (initial compromise) - September 16, 2026 (mass mailing and containment)

Report date

September 16, 2026

Version

1.0 - Public sample (personal and customer identifiers redacted)

Classification

Public sample - all personal data replaced with redaction blocks

Distribution

May be published as an example of IT Partner incident-response deliverables

This report was prepared from Microsoft Entra ID, Microsoft Defender XDR, Exchange Online and IT Partner service-desk records collected on September 16, 2026. All times are UTC unless stated otherwise. Findings reflect the evidence available at the time of writing; open items are listed in Section 7.

Contents

Incident at a glance	3
1. Executive Summary	4
2. Scope, Sources and Method	4
3. What Happened: Attack Narrative	5
3.1 Precursors (August 17-19).....	5
3.2 Initial access: adversary-in-the-middle, August 19	5
3.3 Reconnaissance and the missed opportunity, August 19	6
3.4 Persistence: four weeks of silent token refresh.....	6
3.5 Action on objective: mass phishing from the mailbox, September 16	6
3.6 Detection and containment, September 16	7
3.7 Secondary compromise: internal click and malware execution	7
4. Detailed Timeline of Events	8
5. Why It Happened: Root Cause and Contributing Factors	11
6. Impact Assessment.....	12
7. Containment and Recovery Status	12
8. Recommendations	13
Appendix A. Indicators of Compromise.....	15
Appendix B. Recipients of the September 16 campaign.....	16
Appendix C. Baseline: the user's legitimate networks and devices	16

Incident at a glance

Incident	Adversary-in-the-Middle (AiTM) credential/session theft followed by 4 weeks of persistent mailbox access, directory reconnaissance targeting payroll/HR/finance staff, and a mass outbound phishing campaign sent from the compromised mailbox.
Affected account	██████████, ██████@██████████ (Object ID ████████████████████). Assistant to ██████████. Holds the Helpdesk Administrator and Entra Joined Device Local Administrator roles.
Initial compromise	Tuesday, August 19, 2026 at 17:47 UTC (12:47 PM CDT) via AiTM phishing proxy from 50.232.246.45 (Lititz, PA). MFA push approved by the user through the proxy.
Persistence	August 19 through September 16, 2026: stolen "Microsoft Outlook" refresh token replayed roughly three times a day from 88 different US residential-proxy IP addresses (95 successful token refreshes).
Mass mailing	Wednesday, September 16, 2026, 13:31 and 13:38 UTC (8:31 / 8:38 AM CDT): two batches, 202 unique recipients (187 external, 15 internal), subject "From ██████████ - Please check it out", link to hxxps://novix[.]vu/adobereader delivering Trojan:Script/Sabsik.EN.A!ml (VBS downloader).
Secondary impact	Internal user ██████████ (██████@██████████) clicked the link at 16:34 UTC on device ██████████, downloaded and ran the VBS four times; Defender Antivirus quarantined it (Incident 25347, still Active/Unassigned).
Containment	September 16, 13:41-15:23 UTC: sessions revoked, password reset, all MFA methods and Windows Hello credentials removed, user re-registered MFA. Account remains enabled. No malicious inbox rules, forwarding, or app consents were found.
Overall status	Contained for the primary account. Open items: Defender incidents 25188 and 25347 unresolved; phishing messages still sitting in 13 internal inboxes; endpoint ██████████ not isolated; external recipients not yet notified; ██████████ risk not remediated.

1. Executive Summary

On the morning of September 16, 2026, roughly 200 contacts of ██████ received an email from her legitimate mailbox titled "From ██████ - Please check it out", inviting them to open a document shared "via OneDrive". The link pointed to a malware download. The message was not spoofed: it was sent by an attacker who was signed in to ██████ mailbox through Outlook on the web from a hosting-provider IP address in New Jersey.

The investigation of Entra ID sign-in and audit logs, Identity Protection, Defender XDR advanced hunting, Exchange message trace, and IT Partner's ticket history shows that this was not a same-day event. The account was first compromised four weeks earlier, on August 19, 2026, when ██████ credentials and MFA approval were captured by an adversary-in-the-middle phishing proxy. Microsoft Entra ID Protection detected the AiTM sign-in in real time and Defender XDR raised a High-severity incident (25188) the same day, including an alert that the attacker was using her token to search the directory for payroll, HR, human resources, finance, support and info accounts. Because the sign-in had "passed" MFA the user risk auto-remediated to Low, and at 19:24 UTC that day the remaining risk was dismissed by ██████@██████. The Defender incident was never assigned or closed. From that point the attacker quietly refreshed the stolen Outlook token about three times a day from a rotating pool of 88 US residential IP addresses, keeping mailbox access alive without triggering any further Identity Protection detections.

On September 16 the attacker cashed in. At 12:22 UTC they opened an Outlook on the web session from 74.50.68.42, harvested ██████ contacts, and at 13:31 and 13:38 UTC sent two ~100-recipient batches of the malicious message. ██████ noticed within five minutes and emailed ██████ ("Spam Emails") at 13:36 UTC; the tenant admin (██████@██████) revoked sessions at 13:41 UTC and reset the password at 13:42 UTC, about ten minutes after the first batch left. Because Outlook on the web was still holding a valid access token, and Continuous Access Evaluation was not in effect for that session, the second batch was already gone by then. Recipients included ██████ staff, the organization's IT and telecom vendors (CDW, Spectrum/Charter, Avaya, Granite, Cloudscale365, Forte, IT Partner), and ██████ own family and personal addresses. One internal recipient (██████) clicked the link three hours later and executed the VBS payload on a managed laptop; Defender Antivirus quarantined it, but that incident is also still open.

Root cause: a phishing-resistant sign-in was not required, so a proxied MFA push approval was enough to mint a long-lived session. Contributing causes: High-severity alerts on August 19 were dismissed rather than remediated, the daily token-refresh pattern from unfamiliar residential IPs was not monitored, and a routine end-user account held the Helpdesk Administrator role. Detailed timeline, evidence, gaps and recommendations follow.

2. Scope, Sources and Method

This report covers activity for the account ██████@██████ from August 17 to September 16, 2026 (the 30-day retention window of the sources used). The following sources in the ██████ tenant were reviewed on September 16, 2026:

- Microsoft Entra ID sign-in logs, interactive (236 events, 30 days) and non-interactive (1,530 aggregated rows, 30 days), including per-event authentication details, device and Conditional Access results.
- Microsoft Entra ID audit logs for the user (360 events, 30 days) and the user's directory role assignments.

- Microsoft Entra ID Protection risk detections and risky-user history for the account.
- Microsoft Defender XDR: incidents 25188 and 25347, alert details, and Advanced Hunting queries against CloudAppEvents, EmailEvents, EmailUrlInfo, UrlClickEvents, AlertInfo/AlertEvidence, DeviceProcessEvents, DeviceFileEvents, DeviceNetworkEvents and DeviceEvents.
- Exchange Online message trace for messages sent by the mailbox, September 6-16 (675 rows), mailbox forwarding configuration, and the Restricted Entities list.
- IT Partner Zoho Desk tickets 16435, 16438, 16443, 16444, 16704 and 16727.

All times in this report are given in UTC with Central Daylight Time (CDT, UTC-5) in parentheses where useful. Note that the Entra and Defender portals were displaying America/Phoenix time (UTC-7) during collection and the Exchange message-trace grid labelled its column "UTC-05:00" but its values were one hour ahead of the Defender EmailEvents record for the same messages; the EmailEvents timestamps were used as authoritative for the mail-flow times.

Limitations: mailbox-level audit events (MailItemsAccessed, Send, folder access) were only sparsely surfaced in CloudAppEvents, so the exact set of emails and attachments the attacker read cannot be enumerated from the sources used; a Purview Audit search for the attacker IP addresses is recommended to complete that picture. The original phishing lure that produced the August 19 AiTM sign-in was not found in the 30-day inbound mail data and may have arrived earlier, by SMS, or on the user's phone (the proxy presented an iPhone Safari user agent).

3. What Happened: Attack Narrative

3.1 Precursors (August 17-19)

The user's contact cache contained at least 16 look-alike phishing sender domains (for example secure-micros0ft365.com, ad0besign-secure.com and secure-sharpeoint-docs.com both impersonating [REDACTED]@, cl0udsca1e365.com, hellos1gn-mail.com, [REDACTED], cloudsca1e365.comn). These only surface because the September 16 campaign mailed every cached contact; they show the account had been a repeated phishing target well before the compromise. Three "Request to Reset Microsoft Online Services Password" self-service-reset notifications arrived on August 18 at 18:51 and 22:20 UTC and August 19 at 11:23 UTC. On August 18 at 22:42 UTC [REDACTED] ([REDACTED]@) reset [REDACTED] password as an admin and at 22:54 UTC she changed it again herself from her home network in [REDACTED], TX. IT Partner's risky-user monitor also raised tickets for several other [REDACTED] users on August 18-20 ([REDACTED], [REDACTED], [REDACTED], [REDACTED]), which is consistent with a campaign hitting the organization at that time.

3.2 Initial access: adversary-in-the-middle, August 19

At 17:47:47 UTC on August 19 a sign-in for [REDACTED] to the "OfficeHome" application arrived from 50.232.246.45 (Lititz, Pennsylvania). The correct password was supplied. The user agent claimed to be an iPhone running Safari 26.5, but Entra recorded the operating system as Windows 10, a classic artifact of an AiTM proxy relaying the victim's browser session. Twenty-one seconds later (17:48:08 UTC) the sign-in was interrupted at the "Stay signed in?" prompt with the record showing "Mobile app notification: MFA completed in Azure AD", meaning [REDACTED] approved a Microsoft Authenticator push that was in fact generated by the attacker's proxy. At 17:48:49 UTC the session completed successfully with "MFA requirement satisfied by claim in the token". Entra ID Protection flagged this exact event in real time as "Successful user sign-in due to Adversary-in-The-Middle activity" (High severity, Defender alert adcfdaa5...) and Defender XDR opened incident 25188 at 18:44 UTC.

At 18:57:45 UTC the attacker attempted to redeem the stolen session for a "Microsoft Outlook" (app ID 5d661950-3475-41cd-a2c3-d671a3162bc1) token from 24.106.195.254 (Oxford, NC) with a Firefox/Windows user agent; that first attempt failed with error 50199 (user confirmation required) but the retry six seconds later from 98.114.220.125 (Philadelphia, PA) succeeded, again with first factor and MFA "previously satisfied by claim in the token". This gave the attacker a refresh token for the Outlook family of applications scoped to Exchange Online and Microsoft Graph, which is what they used for the next four weeks.

3.3 Reconnaissance and the missed opportunity, August 19

At 19:24:10 UTC on August 19, about 90 minutes after the AiTM sign-in and 40 minutes after Defender opened the incident, [REDACTED]@[REDACTED] dismissed the user risk on [REDACTED] in Identity Protection (audit activity "DismissUser"). Identity Protection had already down-rated the four atypical-travel detections to "Remediated / Low / User passed multifactor authentication", because from Entra's perspective MFA had been satisfied. No password reset, session revocation, or MFA re-registration was performed on the account that day (the only resets in the audit log on August 19 were ones [REDACTED] herself performed for other users as a helpdesk admin).

At 21:55:55 UTC a second High-severity Entra threat-intelligence alert fired: "Activity linked to a compromised user session or stolen token(s)". Eleven seconds later Defender XDR raised "Suspicious Exchange Online Graph Reconnaissance Activity" (High): using [REDACTED] Outlook token, the attacker called GET https://graph.microsoft.com/v1.0/users?\$search= with an OR-list of displayName, givenName, surname, jobTitle, mail and userPrincipalName matches for "payroll", "pay", "hr", "human", "resources", "support", "info" and "finance". This is textbook targeting for payroll-diversion and invoice-fraud BEC. Mailbox item and attachment access through Graph was logged at 22:44 UTC. All three alerts and the incident remain in status New / Active / Unassigned as of this report.

3.4 Persistence: four weeks of silent token refresh (August 19 - September 16)

From the evening of August 19 until the day of the mass mailing the non-interactive sign-in log shows the "Microsoft Outlook" application obtaining tokens for Office 365 Exchange Online for [REDACTED] approximately three times per day, every day, each time from a different IP address: 95 successful sign-ins from 88 unique addresses in Comcast, Spectrum, Frontier, AT&T, Verizon, T-Mobile, Cox and other consumer ranges spanning at least a dozen US states. Not one address repeated, and none belonged to [REDACTED] known networks ([REDACTED] [REDACTED] office, [REDACTED] [REDACTED] branch, [REDACTED] and [REDACTED] at home, T-Mobile and AT&T mobile ranges on her phone). This is the signature of a stolen refresh token being replayed through a residential proxy network. Because every refresh presented a valid MFA claim, Identity Protection generated no detections for any of them, and the "GeoFence" and "Traveling Geofence" Conditional Access block policies never applied. On September 14 at 21:48 UTC the attacker tried an interactive Outlook sign-in from 23.27.93.176 (Orem, Utah) that was interrupted with error 50199; more successful refreshes followed on September 14-15 from 23.27.93.40, 23.116.73.113, 50.88.194.86 and 216.75.151.162.

3.5 Action on objective: mass phishing from the mailbox, September 16

At 12:22:17 UTC on September 16 the attacker opened an Outlook on the web ("One Outlook Web") session from 74.50.68.42, an address in AS19318 (a New Jersey hosting provider), on an unmanaged, non-compliant Windows 10 device running Chrome 152. The sign-in record again shows first factor and MFA "previously satisfied by claim in the token", the "MFA enforce" Conditional Access policy reporting Success, Token Protection "Unbound", and Continuous Access Evaluation "No". Over the following hour the same IP generated 15 non-interactive token requests for OWA sub-services (attachment download,

URL preview, presence, Teams/Graph app catalog and the "WeveAgave" Exchange service), consistent with browsing the mailbox and exporting the address book.

Between 13:31:20 and 13:31:45 UTC (8:31 AM CDT) the first batch of "From ██████████ - Please check it out" was submitted from 74.50.68.42 to about 102 recipients, and between 13:38:23 and 13:39:01 UTC a second batch went to a further 101. The body claimed a document had been shared via OneDrive because it exceeded the attachment limit and carried a single "VIEW DOCUMENT" link to `hxxps://novix[.]vu/adobereader`; the message was signed with ██████████ real signature block, added by the Exclaimer Cloud signature connector as the mail left the tenant (each message therefore appears twice in message trace). Defender for Office 365 assigned no threat verdict at send time. Of 202 unique recipients, 15 were internal ██████████ addresses (13 delivered to the Inbox and still there at the time of writing, 2 failed), 41 addresses failed (mostly dead look-alike domains and stale contacts), and the remainder were delivered externally to vendors, partners, personal contacts and IT Partner's own addresses (support@, sales@, ██████████; sales@ and ██████████).

3.6 Detection and containment, September 16

██████████ recognised the outbreak almost immediately: at 13:36 UTC she emailed ██████████ "Spam Emails" and at 13:37 UTC forwarded it to ██████████. At 13:41:13 UTC the admin account ██████████@██████████ invalidated all refresh tokens (Update StsRefreshTokenValidFrom Timestamp) and at 13:42:08 UTC reset the password. ██████████ changed her password again at 13:44 and 13:51 UTC and updated her MFA registration. Between 14:46 and 15:22 UTC ██████████ deleted every registered security-info method on the account (100 "Admin deleted security info" events) and 27 Windows Hello for Business credentials, and at 15:57-15:58 UTC ██████████ re-registered MFA from the ██████████ office. A recall was issued at 16:03 UTC, and IT Partner ticket 16704 (a 9/15 risky-event ticket) was updated with the hack at 14:15 UTC. The second batch had already left the tenant two to three minutes before the revocation; the ten-minute gap between the first batch and containment is a good response time, but the attacker's OWA session had been open for 69 minutes before any mail was sent.

3.7 Secondary compromise: internal click and malware execution

At 16:34:39 UTC ██████████ (██████████@██████████) clicked the novix[.]vu link from ██████████ (the ██████████ office); Safe Links recorded "ClickAllowed" because the URL had no verdict. At 16:34:53 UTC Microsoft Edge on device ██████████ saved AdobeReader-Setup03638x.vbs to the Downloads folder and at 16:35:01 UTC wscript.exe executed it. Defender Antivirus detected the file as Trojan:Script/Sabsik.EN.A!ml at 16:35:59 and 16:36:24 UTC and reported it remediated, yet wscript.exe launched the script again at 16:36:35 and 16:38:13 UTC (a second download named "AdobeReader-Setup03638x (1).vbs" is referenced in the ticket). No child processes or outbound network connections from wscript.exe were recorded, which suggests the payload was blocked before it could stage, but this has not been confirmed by a forensic review. Defender XDR created incident 25347 "Script dropped and launched" (High) at 16:37:57 UTC; Identity Protection flagged ██████████ as risky (aiCompoundAccountRisk) at 16:43 UTC, generating IT Partner ticket 16727. That ticket was answered at 17:19 UTC with a recommendation to run an antivirus scan, treated the file as "Adobe Reader-related", and was closed at 18:08 UTC after ██████████ said she would look into it. The connection between this event and the phishing campaign sent from her own mailbox three hours earlier was not made at the time.

4. Detailed Timeline of Events

Source abbreviations: SI = Entra interactive sign-in log; NI = Entra non-interactive sign-in log; AU = Entra audit log; IDP = Identity Protection; XDR = Defender XDR alert/incident; AH = Advanced Hunting; MT = Exchange message trace / EmailEvents; ZD = IT Partner Zoho Desk.

Time (UTC)	Event	Source
Aug 18 18:51	SSPR "Request to Reset Microsoft Online Services Password" notification delivered to ██████ (first of three).	AH
Aug 18 22:20	Second SSPR reset notification delivered.	AH
Aug 18 22:42	██████@██████ resets ██████ password (admin reset) and invalidates tokens.	AU
Aug 18 22:50-22:54	██████ updates MFA registration and changes password (self-service) from home network ██████, ██████.	AU/SI
Aug 19 11:23	Third SSPR reset notification delivered.	AH
Aug 19 13:47	Four "Atypical travel" detections (Low, real-time) as ██████ signs in from new branch IP ██████ (██████). Risk auto-remediated ("User passed MFA").	IDP
Aug 19 14:00-14:10	IT Partner ticket 16435 "4 risky events were detected for user ██████" opened by the risk monitor and closed 10 minutes later (Texas location exclusion).	ZD
Aug 19 17:47:47	AtTM: correct password submitted for ██████ from 50.232.246.45 (Lititz, PA), app OfficeHome, UA iPhone/Safari but OS recorded as Windows 10. Error 399218 (user confirmation required).	SI
Aug 19 17:48:08	Same source: "Mobile app notification - MFA completed" (user approved push through the proxy); interrupted at Keep-Me-Signed-In (50140).	SI
Aug 19 17:48:49	Session established from 50.232.246.45; MFA satisfied by claim in token. IDP real-time detection "Successful user sign-in due to Adversary-in-The-Middle activity". Defender alert (High, Persistence).	SI/IDP/XDR
Aug 19 18:00:56	21 UpdateInboxRules events from ██████ via OUTLOOK.EXE (user's own client syncing benign rules "Spectrum Business Experience Team", "Amazon.com"). No attacker rules.	AH
Aug 19 18:44:12	Defender XDR incident 25188 "Multi-stage incident involving one user reported by multiple sources" created (High).	XDR
Aug 19 18:57:45	Token replay attempt for app "Microsoft Outlook" (5d661950) from 24.106.195.254 (Oxford, NC), Firefox 151/Win10, fails 50199.	SI
Aug 19 18:57:51	Retry from 98.114.220.125 (Philadelphia, PA) succeeds: refresh token for Exchange Online/Graph issued to the attacker.	SI
Aug 19 19:24:10	██████@██████ dismisses user risk for ██████ (IdentityProtection DismissUser). No password reset or token revocation performed.	AU
Aug 19 21:55:55	IDP threat-intelligence alert "Activity linked to a compromised user session or stolen token(s)" (High).	XDR
Aug 19 21:56:06	Defender alert "Suspicious Exchange Online Graph Reconnaissance Activity" (High): Graph /users?\$search for payroll, pay, hr, human, resources, support, info, finance using ██████'s token.	XDR
Aug 19 22:44:26	MailItemsAccessed and AttachmentAccess via Graph (Microsoft service IP 2603:10b6:408:275::15).	AH

Time (UTC)	Event	Source
Aug 19 - Sep 15	Persistence: 95 successful non-interactive "Microsoft Outlook" token refreshes for Exchange Online from 88 unique residential IPs, ~3 per day, no repeats (see Appendix A). Zero IDP detections.	NI
Aug 20-25	IT Partner tickets 16443/16444/16454/16470 for other █████ users (█████, █████, █████) flagged risky; ticket 16438 for █████. Incident 25188 not actioned.	ZD
Sep 2 14:52	█████ registers new security info from a T-Mobile mobile IP (Dallas). Benign.	AU
Sep 11 03:53	Two "Self-service password reset flow activity progress" events for █████ (late evening local). Origin not visible; should be validated with the user.	AU
Sep 14 21:48:17	Interactive token replay for "Microsoft Outlook" from 23.27.93.176 (Orem, UT), Firefox 153/Win10, error 50199. Further successful refreshes from 23.27.93.40 and 23.116.73.113.	SI/NI
Sep 15 13:15	IT Partner ticket 16704 "A risky event was detected for user █████" opened (█████ as contact).	ZD
Sep 15 (day)	Token refreshes from 50.88.194.86 (success) and 216.75.151.162 (fail).	NI
Sep 16 12:22:17	Attacker opens Outlook on the web from 74.50.68.42 (Englewood Cliffs NJ, AS19318), Chrome 152/Win10, unmanaged device; MFA claim in token; CA "MFA enforce" Success; GeoFence not applied; CAE No; Token Protection Unbound.	SI
Sep 16 12:22-13:30	15 OWA sub-service token requests from 74.50.68.42 (OwaDownloadAttachment, URL preview, presence, Teams/Graph app catalog, WeveAgave/Exchange).	NI
Sep 16 13:31:20-13:31:45	Batch 1: "From █████ - Please check it out" sent from 74.50.68.42 to ~102 recipients (link hxxps://novix[.]vu/adobereader). 13 internal copies delivered to Inbox.	MT
Sep 16 13:36	█████ emails █████ "Spam Emails"; 13:37 forwards to █████.	MT
Sep 16 13:38:23-13:39:01	Batch 2 to ~101 further recipients. 41 addresses overall fail delivery.	MT
Sep 16 13:41:13	█████@█████ revokes all sessions (StsRefreshTokenValidFrom) for █████.	AU
Sep 16 13:42:08	█████ resets █████ password (admin).	AU
Sep 16 13:44-13:51	█████ changes password (self-service) and updates MFA from █████.	AU
Sep 16 13:54-14:00	"Email Hack?" thread with █████ and █████; "FW: Document" to Forte.	MT
Sep 16 14:15-14:16	█████ replies on IT Partner ticket 16704 reporting the risky event/hack.	MT/ZD
Sep 16 14:46-15:22	█████ deletes all security info (100 events) and 27 Windows Hello for Business credentials; device registration updates.	AU
Sep 16 15:57-15:58	█████ re-registers MFA (Microsoft Account Controls V2 / My Profile) from █████ office.	SI/AU
Sep 16 16:03	Recall attempt for the phishing message.	MT
Sep 16 16:34:39	█████@█████ clicks novix[.]vu link from █████; Safe Links "ClickAllowed".	AH
Sep 16 16:34:53	msedge.exe writes C:\Users\█████\Downloads\AdobeReader-Setup03638x.vbs on device █████.	AH
Sep 16 16:35:01	wscript.exe executes the VBS (parent msedge.exe).	AH

Time (UTC)	Event	Source
Sep 16 16:35:59 / 16:36:24	Defender AV detects Trojan:Script/Sabsik.EN.A!ml, action quarantine, WasRemediated=true.	AH
Sep 16 16:36:35 / 16:38:13	wscript.exe launched twice more (second download "(1).vbs"). No child processes or network connections observed.	AH
Sep 16 16:37:57	Defender XDR incident 25347 "Script dropped and launched" created (High, Active, Unassigned).	XDR
Sep 16 16:43:24	IDP flags █████ (aiCompoundAccountRisk, Low, atRisk). IT Partner ticket 16727 created 16:45.	IDP/ZD
Sep 16 17:19	IT Partner responds on 16727: "looks like an Adobe Reader-related process", recommends AV scan and contacting internal IT.	ZD
Sep 16 17:53	█████ replies to █████ and █████ on the phishing thread (warning recipients).	MT
Sep 16 18:04-18:08	█████ replies on 16727 ("I will look into on my end"); ticket closed.	ZD
Sep 16 (report time)	█████ account enabled; forwarding off; no inbox rules; incidents 25188/25347 Active; 13 phishing copies in internal inboxes; █████ not isolated.	All

5. Why It Happened: Root Cause and Contributing Factors

5.1 Root cause

The account was protected by push-notification MFA enforced through the "MFA enforce" Conditional Access policy, but nothing required the authentication to be phishing-resistant or bound to a trusted device. An AiTM proxy sat between [REDACTED] browser and login.microsoftonline.com, relayed her password, triggered a genuine Authenticator push that she approved, and captured the resulting session cookie. From Entra's point of view every later use of that session and the refresh tokens derived from it legitimately carried an MFA claim, so the MFA policy reported "Success" on the attacker's sign-ins from Pennsylvania, Utah and New Jersey.

5.2 Contributing factors

- **Alerts were dismissed instead of remediated.** Identity Protection raised a real-time AiTM detection and Defender XDR opened a High incident within an hour of the compromise. Because the risk had auto-remediated to Low, it was dismissed by the customer administrator at 19:24 UTC and the incident was never assigned. Two more High alerts at 21:55-21:56 UTC (stolen token, Graph reconnaissance for payroll/HR/finance accounts) went unactioned for 28 days.
- **The risky-user monitor closed the wrong ticket.** Ticket 16435 covered the harmless 13:47 UTC atypical-travel events from a [REDACTED] branch and was closed on the "Texas" exclusion. The AiTM detection four hours later did not generate a new ticket because the user's risk state was already "Remediated"/dismissed, and the monitor does not ingest Defender XDR incidents.
- **Token refresh from unfamiliar networks was invisible.** Eighty-eight different residential IPs refreshing an Outlook token three times a day is a strong signal, but non-interactive sign-ins do not feed Identity Protection risk, and no Conditional Access policy, named-location alert or hunting rule looked at them. The two geofence block policies existed but did not apply to token-based sign-ins.
- **No token protection or continuous access evaluation on the OWA session.** The September 16 OWA sign-in shows Token Protection "Unbound" and CAE "No". Revoking sessions at 13:41 UTC therefore did not tear down the access token the attacker already held; the second batch was sent at 13:38 UTC and the OWA session could have continued for up to an hour.
- **An end-user account held Helpdesk Administrator.** [REDACTED] performs password resets for staff (12 resets for other users between August 19 and September 14 are in the audit log, all consistent with her normal working hours and office IPs). A compromised Helpdesk Administrator can reset non-admin users' passwords and revoke their sessions through Graph. No evidence was found that the attacker used the role, but the exposure was real and the Graph search for payroll and HR staff suggests it was being considered.
- **The account was a known repeat phishing target.** At least 16 look-alike domains, several impersonating [REDACTED], were already in her contact cache, and three SSPR notifications hit the mailbox in the 24 hours before the compromise. There was no elevated monitoring or user coaching in response.
- **Outbound mail controls did not slow the burst.** Two batches of ~100 external recipients within seven minutes did not trip an outbound spam or recipient-rate restriction (the Restricted Entities list is empty), the URL had no reputation and Safe Links let the internal click through, and Zero-hour Auto Purge has not removed the 13 internal copies.

6. Impact Assessment

- **Mailbox confidentiality:** the attacker held read access to ██████'s mailbox (assistant to the CEO-level executive ██████) for 28 days and accessed mail items and attachments through Graph on August 19. Content read after that date cannot be enumerated from the sources used; assume all mail, calendar and contacts through September 16 are exposed. The mailbox contains vendor quotes, invoices (I.T. Department Invoices), telecom account correspondence, verification codes and internal HR/payroll adjacent threads.
- **Directory exposure:** the Graph search enumerated users matching payroll, HR, finance, support and info; names, titles and addresses of those staff should be treated as known to the attacker and those users warned about payroll-diversion attempts.
- **Reputational / third-party:** 187 external addresses received malware from a legitimate ██████ mailbox, including CDW, Spectrum/Charter, Avaya, Granite, Cloudscale365, Forte Technologies, UT Dallas, NTBHA, IT Partner and personal family contacts. Any of those recipients may have executed the VBS.
- **Internal endpoint:** device ██████ executed the Sabsik VBS downloader four times; AV reports remediation but the device has not been isolated or forensically reviewed, and ██████'s identity is flagged at risk.
- **Other tenant users:** risky-user tickets for ██████, ██████, ██████, ██████, ██████, ██████ and the ██████ admin account (with sign-ins recorded from Cape Town/Kwadabeka South Africa and New Jersey) over the same window indicate the campaign and/or the token-replay infrastructure may have touched more than one account.

7. Containment and Recovery Status

Action	Status	Evidence
Revoke refresh tokens / sessions	Done 9/16 13:41 UTC	AU: Update StsRefreshTokenValidFrom by ██████
Reset password	Done 9/16 13:42 UTC (admin) and 13:44/13:51 (user)	AU
Remove all MFA methods and re-register	Done 9/16 14:46-15:58 UTC	AU: 100 "Admin deleted security info", 27 WHfB deletions, user re-registration
Mailbox forwarding	Verified OFF	EAC mailbox properties
Malicious inbox rules	None found	AH CloudAppEvents: only benign UpdateInboxRules from user IPs
OAuth app consents / new devices by attacker	None found	AU: device registrations 9/3 and 9/14 from office IPs (user)
Disable account / block sign-in	Not done (account Enabled)	Entra user overview
Purge phishing from internal mailboxes	Not done (13 copies in Inbox)	AH EmailEvents LatestDeliveryLocation
Isolate / scan endpoint ██████	Not done	XDR incident 25347 Active
Resolve Defender incidents 25188 / 25347	Not done (Active, Unassigned)	XDR
Remediate ██████ risky user	Not done	IDP / ticket 16727 closed

Action	Status	Evidence
Block IOCs (novix[.]vu, sender IPs)	Not done	No tenant allow/block entry observed
Notify external recipients	Partial (user replied to 2 recipients; Spectrum and Forte aware)	MT

8. Recommendations

8.1 Immediate (today)

- Isolate device ██████████ in Defender for Endpoint, run a full offline scan, collect the investigation package, and review for persistence (scheduled tasks, Run keys, additional downloads) before returning it to service. Reset ██████████'s password, revoke sessions and confirm the user risk, then resolve incident 25347.
- Use Threat Explorer or a Purview content search to soft-delete every copy of "From ██████████ - Please check it out" from the 13 internal inboxes, and add novix[.]vu and the URL to the Tenant Allow/Block List as Block.
- Assign, investigate and resolve Defender incident 25188 as True Positive / Compromised account so the AiTM, stolen-token and Graph-reconnaissance alerts are closed with the right classification, and confirm ██████████ as compromised in Identity Protection rather than dismissed.
- Run a Purview Audit search on MailItemsAccessed, Send, SendAs, MoveToDeletedItems and HardDelete for ██████████ filtered to the attacker IP addresses in Appendix A to establish exactly what was read, sent or deleted, and check the Deleted Items and Sent Items folders for evidence of attacker replies or invoice/banking conversations.
- Send a notification to all 202 recipients (Appendix B) from a different mailbox, on ██████████ letterhead, stating the message was sent by an attacker, not to open the link or any AdobeReader-Setup*.vbs file, and to contact their IT if they did. Prioritise vendors handling payments (CDW, Avaya, Spectrum, Granite, Cloudscale365) and warn finance that payment-instruction changes must be verified by phone.
- Remove the Helpdesk Administrator and Entra Joined Device Local Administrator roles from ██████████ (or move them to a separate cloud-only admin account protected by phishing-resistant MFA and PIM), and review the ██████████ admin account's South-African and New-Jersey sign-ins.
- Interview ██████████ about the Authenticator prompt she approved around 12:48 PM CDT on August 19, and about the September 11 evening self-service-reset attempt, to identify the lure (email, text or QR code) so it can be blocked and other recipients found.

8.2 Short term (next 30 days)

- Require phishing-resistant MFA (Authenticator passkeys / FIDO2 / Windows Hello for Business) via an authentication-strength Conditional Access policy, starting with all admins, executives and executive assistants, finance, HR and payroll. Until then enable number matching, additional context and Authenticator app-lock, and disable SMS/voice.
- Enable Token Protection (sign-in session binding) for Exchange Online and SharePoint on Windows devices, and make Continuous Access Evaluation strict-location enforcement part of the policy so revocation takes effect within minutes rather than up to an hour.

- Require compliant or hybrid-joined devices for Outlook on the web and desktop Outlook, or at minimum for all admin roles; the September 16 OWA session came from an unmanaged Windows device.
- Turn on the Microsoft-managed / custom risk-based policies: block on High sign-in risk, require secure password change on High user risk, and require MFA re-registration on Medium+; and stop dismissing risk as a first response, dismissal must be reserved for confirmed benign events documented in the ticket.
- Fix the alert pipeline so that Defender XDR incidents of High severity, Identity Protection "AiTM" and "compromised session / stolen token" detections, and Graph reconnaissance alerts create tickets independent of the risky-user state, with a playbook that mandates password reset plus session revocation plus MFA re-registration (not "run an AV scan") for any script-execution or AiTM alert.
- Add a Defender/Sentinel custom detection for the persistence pattern seen here: non-interactive Exchange token refreshes for one user from more than N distinct IPs or ASNs in 24 hours, or from residential ISPs outside the user's named locations.
- Tighten outbound mail: set an outbound spam policy recipient limit (for example 200/day external, 50/hour) with an admin notification, enable Zero-hour Auto Purge for phishing and malware, enable Safe Links "do not allow click-through", and confirm Safe Attachments Dynamic Delivery is on.
- Ensure Defender for Endpoint attack-surface-reduction rules "Block JavaScript or VBScript from launching downloaded executable content" and "Block execution of potentially obfuscated scripts" are in Block mode, and associate .vbs/.js/.jse/.wsf with Notepad on user workstations.

8.3 Longer term

- Adopt Privileged Identity Management with just-in-time activation for every Entra role, and separate admin identities from daily-driver mailboxes.
- Deploy Defender for Cloud Apps session policies and the Exchange mailbox auditing (Audit Premium) needed to answer "what did the attacker read" quickly, and extend log retention beyond 30 days (Sentinel or Purview 1-year audit).
- Run targeted phishing-simulation and AiTM-awareness training for executive assistants and finance, specifically on approving MFA prompts they did not initiate and on look-alike domains impersonating leadership.
- Register DMARC with p=reject and monitor the look-alike domains observed here; several (ad0besign-secure.com, secure-sharpeoint-docs.com) impersonate the CEO to his own assistant.

Appendix A. Indicators of Compromise

Indicator	Type	Context
hxxps://novix[.]vu/adobereader	URL	Payload link in the 9/16 phishing message (VIEW DOCUMENT)
AdobeReader-Setup03638x.vbs / AdobeReader-Setup03638x (1).vbs	File	Downloaded payload, Trojan:Script/Sabsik.EN.A!ml
74.50.68.42	IPv4 (AS19318, Englewood Cliffs NJ)	9/16 OWA session that sent the campaign
50.232.246.45	IPv4 (Lititz PA)	8/19 AiTM proxy sign-in
24.106.195.254	IPv4 (Oxford NC)	8/19 token replay attempt
98.114.220.125	IPv4 (Philadelphia PA)	8/19 Outlook refresh token issued
23.27.93.176, 23.27.93.40, 23.116.73.113, 50.88.194.86, 216.75.151.162	IPv4	9/14-9/15 token replay
74.59.59.81	IPv4	Associated with 8/19 21:55 UTC stolen-token alert (incident 25188 graph)
Residential proxy pool (88 IPs), e.g. 68.100.131.167, 136.62.208.50, 73.172.236.13, 75.139.187.49, 47.222.174.192, 73.122.151.65, 107.198.102.48, 149.18.9.126, 188.221.107.198, 71.43.52.230, 67.146.202.18, 73.43.70.36, 170.124.109.148, 76.17.122.70, 73.248.81.60, 75.112.184.118, 71.143.145.59, 169.128.192.95, 67.243.80.62, 162.204.8.178, 108.31.25.68, 67.83.39.48, 24.170.63.79, 99.31.227.159, 73.135.117.92, 107.213.239.236, 24.8.121.227, 70.8.165.1, 73.63.229.110, 100.37.233.24, 192.184.132.171, 35.136.129.130, 70.40.124.9, 47.149.7.24, 99.180.90.222, 51.194.83.237, 172.56.162.60, 35.144.159.35, 99.89.81.147, 32.218.45.115, 73.208.239.138, 108.86.11.195, 73.56.124.74, 166.199.152.133, 50.114.117.97, 99.116.245.173, 153.66.147.122, 75.80.146.47, 23.26.94.164, 76.33.183.64, 173.189.188.63, 97.202.209.25, 99.97.143.79, 73.7.106.229, 172.56.232.87, 172.56.36.75, 24.128.223.32, 69.143.49.244, 149.18.10.65, 64.52.31.133, 76.32.0.131, 76.20.194.214, 47.184.47.194, 73.126.130.167, 73.128.189.45, 47.205.179.53, 169.128.193.248, 172.58.130.144, 208.161.128.169, 142.111.239.119, 169.128.192.167, 108.30.134.195, 68.62.120.81, 149.18.10.254, 47.146.212.90, 75.15.148.64, 73.201.159.28, 172.59.215.103	IPv4	Non-interactive "Microsoft Outlook" (5d661950-3475-41cd-a2c3-d671a3162bc1) token refreshes 8/19-9/15, one use each; block-listing individual IPs is low value, the pattern is the indicator
User agents	UA	Mozilla/5.0 (Windows NT 10.0; Win64; x64; rv:151.0) Gecko/20100101 Firefox/151.0; ...Firefox/153.0; Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/152.0.0.0 Safari/537.36; iPhone OS 18_7 Safari 26.5.2 with OS reported as Windows10
Look-alike sender domains in contact cache	Domain	secure-micros0ft365.com, ad0besign-secure.com, secure-sharpeoint-docs.com, hellos1gn-mail.com, cl0udsca365.com, abnorma1platform.com, secure-b0x.com, smartsheef-app.com, trel1o-projects.com, mondy-comms.com, [REDACTED], cloudsca365.comn, cloudsca365.co, secure-hashicorpvaatl.com, taxfiling-notices.org, pagerduty-ops.com
Graph request	Query	GET /v1.0/users?\$search="displayName:payroll" OR ... "jobTitle:finance" (payroll, pay, hr, human, resources, support, info, finance) at 2026-08-19 21:56 UTC

Appendix B. Recipients of the September 16 campaign (202 unique)

(F) = delivery failed. Internal [REDACTED] recipients are listed first.

The list of 202 recipient addresses (15 internal, 187 external) is withheld from this sample report.

Recipient domain concentration (message-trace rows): granitenet.com 44, avaya.com 24, cloudscales365.com 22, charter.com 20, gmail.com 20, cdw.com 18, spectrum.com 14, o365hq.com 10, fortetechnologies.com 10, hightideco.com 8.

Appendix C. Baseline: the user's legitimate networks and devices

Network	Observed as	Interactive sign-ins (30d)
[REDACTED]	[REDACTED] - [REDACTED] office	140
[REDACTED]	[REDACTED] - [REDACTED] branch (first seen 8/19)	67
[REDACTED]	[REDACTED] (home or mobile)	3
[REDACTED]	AT&T residential, [REDACTED] (home)	14
2607:fb90/fb91::, 2600:387/383::, 172.56-59.x	T-Mobile / AT&T mobile (iPhone: Apple Internet Accounts, Authenticator)	non-interactive only

Every interactive sign-in from these networks used the expected client mix (Windows Sign In, Office365 Shell, SharePoint, Teams, Authentication Broker) and, where MFA applied, satisfied it with a real Authenticator interaction. All attacker sign-ins are distinguishable by (a) MFA "previously satisfied by claim in the token", (b) unmanaged/non-compliant device, and (c) a US location outside Texas.