AegisAI

SEVEN ATTACK TECHNIQUES · OCT-NOV 2026

# Threat Watch Report: Oct-Nov 2026

See what the control saw, what the person saw, and the recommended steps.

MICROSOFT 365 | GOOGLE WORKSPACE

THE VERDICT EACH CONTROL GAVE

| | |
|---|---|
| 01 TRUSTED SENDER R | 05 NO PAYLOAD FOUND I |
| 02 CLEAN I | 06 ORDINARY SITE I |
| 03 GENUINE SIGN-IN I | 07 DELIVERED R |
| 04 MFA ALREADY SATISFIED R | |

R = REPORTED BY RESEARCHERS  I = INFERRED

CONTENTS

# Inside this report


| | | |
|---|---|---|
| INTRO | Introduction | 03 |
| PART 1 | The seven attack techniques | 04 |
| | 01 · Genuine sender, attacker's words | 05 |
| | 02 · Scanner-gated pages | 06 |
| | 03 · Device code sign-in | 07 |
| | 04 · Replayed session tokens | 08 |
| | 05 · Reply-first lures | 09 |
| | 06 · Browser-in-the-browser windows | 10 |
| | 07 · Inbox flood, then Teams | 11 |
| END | Conclusion | 12 |
| SOURCES | Sources | 13 |
| ABOUT | About AegisAI | 14 |

INTRODUCTION

# Seven attacks built on a sender, sign-in or tool your controls already trust

Attackers typed "Helpline" and a phone number into the Ship To field of a Dell order. Dell's own system sent the confirmation, and it passed DKIM and DMARC, the checks that confirm which domain sent a message. Neither check looks at what the message asks the reader to do.[1]

The other six work the same way. YAM Store kits supported fake CAPTCHA gates to obstruct security analysis.[2] Mir0Auth led victims to Microsoft's own sign-in.[3] Stolen session tokens, the proof Microsoft issues after a sign-in, had already passed MFA, the second check after the password.[4] TA419 sent its link only after a reply.[5] Fake AI ad tools drew a Google window inside the page.[6] In "a significant number of cases" seen by Arctic Wolf, an inbox flood came before or alongside an "IT" chat on Teams.[7]

Three of the seven verdicts are reported by the researchers. Four are inferred from how a standard control works, and those pages say what the inference rests on. None of the seven cases was replayed through AegisAI for this report. Small numbers link to the sources on page 13.

TRUST HIJACKING

## An attack delivered through a sender, sign-in page, session or workplace tool that your people and your security controls already accept as genuine.

Sender authentication (DKIM and DMARC), link scanning and MFA each test one thing: which domain sent a message, where a link leads, whether a second proof was given. In the seven cases here that one thing was real or looked ordinary, so the check passed or had nothing to flag. What the checks did not test is what the message was trying to do. That is the question AegisAI's agents investigate on every inbound email.

PART 1

# The seven attack techniques

| | ATTACK TECHNIQUES | FEATURED CASE | ENTRY PATH | WHAT THE CONTROL SAID | PAGE |
|---|---|---|---|---|---|
| 01 | **Genuine sender, attacker's words** | Dell and Staples order emails[1] | Email | TRUSTED SENDER REPORTED | 05 |
| 02 | **Scanner-gated pages** | YAM Store network phishing kits[2] | Email > Web | CLEAN INFERRED | 06 |
| 03 | **Device code sign-in** | Mir0Auth kit[3] | Email | GENUINE SIGN-IN INFERRED | 07 |
| 04 | **Replayed session tokens** | Token spraying[4] | Not reported | MFA ALREADY SATISFIED REPORTED | 08 |
| 05 | **Reply-first lures** | TA419 committee invitations[5] | Email | NO PAYLOAD FOUND INFERRED | 09 |
| 06 | **Browser-in-the-browser windows** | Fake AI ad tools[6] | Email, no detail | ORDINARY SITE INFERRED | 10 |
| 07 | **Inbox flood, then Teams** | Fake Teams help desk[7] | Email > Teams | DELIVERED REPORTED | 11 |

### ENTRY PATH

Entry path is how the attack reached the person, as far as researchers reported it.

### WHAT THE CONTROL SAID

Reported: the researcher states what the control returned.
Inferred: what a standard control would return, with the basis given on the page.

### NOTES

Attackers combine techniques. TA419 used a reply-first lure (05), then a drawn window (06).[5] Kits tied to YAM Store used gates (02), and many of their messages used device code phishing (03).[2]

MFA is multi-factor authentication, a second proof on top of the password.

01 GENUINE SENDER, ATTACKER'S WORDS

ENTRY PATH: EMAIL

# Dell's own order system sent the attacker's phone number.

Across two campaigns of roughly 1,400 messages, attackers got fake support numbers printed in the billing and shipping fields of Staples and Dell order confirmations.[1]

VERDICT GIVEN — REPORTED

## TRUSTED SENDER

The Dell email passed DKIM and DMARC. Those checks "say nothing about what the sending system was told to say."[1]

| LURE | HIDDEN BY | WHAT THE CONTROL SAW | VERDICT | WHAT THE PERSON SAW | OUTCOME |
|---|---|---|---|---|---|
| A real Dell or Staples order confirmation. | Nothing. No link, no file. | Dell's real domain and template. | TRUSTED SENDER | An order they never placed, and a helpline. | The aim was a phone call. |

### HOW IT GOT IN

01 On the Dell order, the Ship To field held the word "Helpline" and a phone number instead of an address.[1]

02 Dell's own system added the number to its real template and sent the email.[1]

03 The attacker wanted a phone call, which takes the victim off email. This is callback phishing.[1]

### WHERE TO FIND IT

MICROSOFT 365

Email & collaboration > Policies & rules > Threat policies: Tenant Allow/Block Lists; Anti-spam allowed senders.[8,10] Exchange admin center > Mail flow > Rules: rules that bypass spam filtering.[9]

GOOGLE WORKSPACE

Apps > Google Workspace > Gmail > Spam, Phishing and Malware: Email allowlist.[11]

### RECOMMENDED STEPS

CHANGE Delete every retailer, shipping and SaaS domain from Tenant Allow/Block Lists and anti-spam allowed senders (Google Workspace: Gmail Email allowlist). Mail from an allowed domain skips spam filtering, whatever the message asks the reader to do.[8,10,11]

CHANGE In Exchange admin center > Mail flow > Rules, delete any rule that bypasses spam filtering for a sender domain, unless someone owns it and can say why it is needed. These rules let every message from that domain through, the same as an allowed sender.[9]

STAFF Add a fake order confirmation with a support number and no link to your next phishing simulation, and count how many staff report it. Then remind everyone that a receipt for an order they did not place goes to security, and the number in it is never called.[1]

### WHERE AEGISAI LOOKS

AegisAI's agents investigate what each inbound email is trying to do, even with no link or file to scan. Here the request is a phone call to a number typed into a shipping field. Proteus, AegisAI's own model, weighs it against evidence the sender does not control. DKIM and DMARC confirm the domain only.

02 SCANNER-GATED PAGES

ENTRY PATH: EMAIL > WEB

# Kits hid their pages behind fake CAPTCHA gates.

Investigators traced an ecosystem that facilitated more than a billion malicious messages, targeting over 75,000 organizations, to the YAM Store marketplace and the phishing kits around it.[2]

VERDICT GIVEN — INFERRED

## CLEAN

A link-scanning control stopped at a CAPTCHA or an IP filter would never reach the phishing page, so it would have nothing to flag.

| LURE | HIDDEN BY | WHAT THE CONTROL SAW | VERDICT | WHAT THE PERSON SAW | OUTCOME |
|---|---|---|---|---|---|
| Unpaid tolls, deliveries, account checks. | Geofencing, fake CAPTCHAs, IP filters. | A gate, and no phishing page behind it. | **CLEAN** | A toll, delivery or sign-in page. | Credentials or account tokens. |

### HOW IT GOT IN

01 The kits offered lures such as unpaid tolls, traffic penalties, delivery problems and account verification requests, sent by email and SMS text.[2]

02 The kits "supported geofencing, fake CAPTCHA gates, IP filtering, and other measures to obstruct security analysis." Geofencing shows the page only in chosen places.[2]

03 Over 750 million phishing messages were sent to Microsoft 365 customers this year, many using device code phishing to obtain tokens.[2]

### WHERE TO FIND IT

MICROSOFT 365

Email & collaboration > Policies & rules > Threat policies > Safe Links: a policy for all users, or the Standard or Strict preset.[12]

GOOGLE WORKSPACE

Apps > Google Workspace > Gmail > Safety > Links and external images: turn on all four settings.[13]

### RECOMMENDED STEPS

REVIEW Compare the users, groups and domains each Safe Links policy covers with all licensed users and shared mailboxes (Google Workspace: confirm all four link settings are on in every organizational unit). Add anyone missing, since they get no link checking when they click.[12,13]

ASK VENDOR What verdict does a link get when your scan ends at a CAPTCHA, a block page or a timeout? If the answer is "clean" or "no verdict", gated phishing pages reach the inbox.[2]

REVIEW Take five phishing emails staff reported last month and look up the verdict each link got at delivery in your email tool's message log. Send any marked clean to your vendor with the message ID and ask why.[2]

### WHERE AEGISAI LOOKS

AegisAI opens every link, file and QR code in a sandbox, an isolated browser, and follows it to the final page. A CAPTCHA or a block page on the way is treated as evidence, and a page it could not reach is never scored clean. Once your admin grants write access, a convicted message is removed from every mailbox it reached.

03 DEVICE CODE SIGN-IN

ENTRY PATH: EMAIL

# A fake DocuSign page handed victims to Microsoft's real sign-in.

Mir0Auth begins with phishing emails carrying PDF attachments. The PDF leads to a fake DocuSign page, then into Microsoft's legitimate device code sign-in.[3]

VERDICT GIVEN — INFERRED

## GENUINE SIGN-IN

The victim signs in on Microsoft's real page, so the log would show a success, and MFA alone cannot prevent it.[3]

| LURE | HIDDEN BY | WHAT THE CONTROL SAW | VERDICT | WHAT THE PERSON SAW | OUTCOME |
|---|---|---|---|---|---|
| A phishing email with a PDF attached. | A fake DocuSign page before Microsoft's. | A PDF, then a sign-in the user completed. | GENUINE SIGN-IN | A document to sign, then a code to enter. | Access tokens, no password stolen. |

### HOW IT GOT IN

01 Device code sign-in lets a device with no keyboard sign in with a code typed on another screen. That device then gets an access token, a key to the account.

02 Hidden endpoints such as /task_queue.cs, /context.go and /window.rs request the device codes and follow each sign-in session.[3]

03 When a victim finishes signing in, the attacker can get access tokens without the password. ANY.RUN saw activity mostly against US organizations.[3]

### WHERE TO FIND IT

MICROSOFT 365

Entra ID > Conditional Access > Policies > New policy. Conditions > Authentication flows > Device code flow. Grant > Block access. Start in Report-only.[14]

GOOGLE WORKSPACE

No equivalent. Nearest: Security > Access and data control > API controls.[16]

### RECOMMENDED STEPS

REVIEW In Entra ID > Conditional Access, confirm the policy that blocks device code flow is On, not Report-only, and give each exception an owner. In Report-only, device code sign-ins are logged but not stopped.[14]

ASK VENDOR Is a PDF whose link ends at Microsoft's device code page flagged as phishing? If not, the email step of this attack reaches the inbox, and the sign-in that follows happens on Microsoft's real page.[3]

REVIEW Check that your account-compromise runbook includes Revoke sessions in Entra ID among its first steps, not only a password reset (Google Workspace: reset the user's sign-in cookies). Add it if it is missing, so tokens already issued stop working.[3,15,17]

### WHERE AEGISAI LOOKS

AegisAI covers the email step of this attack. It opens the PDF and its link in a sandbox, follows each hop and records the chain in the verdict. The sign-in step is outside email. Blocking device code sign-in is your administrator's setting, shown beside this box.

04 REPLAYED SESSION TOKENS

ENTRY PATH: NOT REPORTED

# Replayed session tokens needed no password and no MFA.

Huntress responded to incidents at 79 organizations where attackers replayed stolen Microsoft 365 session tokens, the proof Microsoft issues after a sign-in.[4]

VERDICT GIVEN — REPORTED

## MFA ALREADY SATISFIED

Huntress reports that the token "is minted after MFA", so whoever holds it does not need MFA again.[4]

| LURE | HIDDEN BY | WHAT THE CONTROL SAW | VERDICT | WHAT THE PERSON SAW | OUTCOME |
|---|---|---|---|---|---|
| Not published by Huntress. | Relays, device code phishing, live logins. | A valid token from a rented host. | MFA ALREADY SATISFIED | Their own successful login. | Hidden inbox rules and a payroll email. |

### HOW IT GOT IN

01 Tokens came from relays behind fake sign-in pages (which pass the real login through), device code phishing (technique 03) and social engineering during a real login.[4]

02 Between 14 July 2026 and 22 September 2026, at least 497 sign-ins succeeded against more than 200,000 failed attempts, from rented hosts in one block of IP addresses.[4]

03 Attackers registered at least 48 fake computers, replaced MFA methods and hid inbox rules. One marked all email from a CPA's address as read.[4]

### WHERE TO FIND IT

MICROSOFT 365

Entra ID > Conditional Access > New policy. Users: All users. Target resources: All resources. Grant: Require MFA.[4] After a compromise: Users > the user > Revoke sessions.[15]

GOOGLE WORKSPACE

Directory > Users > the user > Security > Sign-in cookies: Reset.[17]

### RECOMMENDED STEPS

REVIEW Export the exclusions from your all-users MFA Conditional Access policy, and give each one a named owner and a reason, or remove it. Every excluded user or app can sign in without MFA. Emergency access (break-glass) accounts stay excluded.[4,28]

REVIEW Filter Entra ID sign-in logs for error code 53003 over the last 30 days, then reset the password and revoke sessions for every account that appears. Huntress calls this error "a sign of a valid credential, not a save."[4]

CHANGE After any reset, remove the devices and authentication methods the attacker added to the account in Entra ID, then delete inbox rules the user did not create. Huntress saw attackers register fake devices, replace MFA methods and hide rules that marked mail as read.[4]

### WHERE AEGISAI LOOKS

AegisAI's email investigation does not see this technique: a replayed token never passes through the inbox. Its account takeover monitoring, in early access and review-only, reads sign-ins and mailbox activity, inbox rule changes included. Your team acts on each finding in Entra ID, Okta or Google Workspace.

05 REPLY-FIRST LURES

ENTRY PATH: EMAIL

# The link arrived only after the target wrote back.

TA419's emails invited US AI policy experts to join an AI policy advisory committee, in the names of real people in their field.[5]

VERDICT GIVEN — INFERRED

## NO PAYLOAD FOUND

Inferred from the first messages, which asked for no password and pointed to no sign-in page.[5]

| LURE | HIDDEN BY | WHAT THE CONTROL SAW | VERDICT | WHAT THE PERSON SAW | OUTCOME |
|---|---|---|---|---|---|
| An invitation in a real person's name. | A reply. The link came only after it. | A message with nothing to open. | NO PAYLOAD FOUND | An invitation in their own field. | A reply, then a link to a drawn window. |

### HOW IT GOT IN

01 Beginning 8 July 2026, TA419 sent emails spoofing Lynne Edwards Parker, a former White House technology policy official, then economist Heidi Crebo-Rediker.[18]

02 After a target replied, the group sent a shortened link. It redirected through several websites to a false OneDrive sign-in page.[5]

03 CyberScoop says the report "did not identify victims or state whether any accounts were compromised." It does not tie the activity to the Chinese government directly.[5]

### WHERE TO FIND IT

MICROSOFT 365

Set-ExternalInOutlook -Enabled $true (External tag).[19] Threat policies > Anti-phishing (Defender for Office 365): users to protect, outside contacts included.[20]

GOOGLE WORKSPACE

Apps > Google Workspace > Gmail > Safety: Protect against spoofing of employee names.[13]

### RECOMMENDED STEPS

REVIEW Confirm the External tag is on with Set-ExternalInOutlook, then send a test from a personal address and check that Outlook marks it External. Allow up to 48 hours after turning it on before you test.[19]

CHANGE Add the outside experts, advisers and board members your staff write to into Users to protect in your Defender anti-phishing policy (Google Workspace covers employee names only: turn on Protect against spoofing of employee names). Mail using their names from an unknown address is then flagged.[13,20]

ASK VENDOR Is a link that arrives in a reply, inside an existing thread, scanned like a link in a first message? If not, reply-first lures deliver their link after the check has already passed.[5]

### WHERE AEGISAI LOOKS

AegisAI does not need a link or a file to open an investigation. Proteus, AegisAI's own model, checks whether the two parties have written before, the question a first-contact invitation raises. When the link arrives after the reply, it is opened in a sandbox and followed to the end.

06 BROWSER-IN-THE-BROWSER WINDOWS

ENTRY PATH: EMAIL, NO DETAIL

# A drawn Google window asked for the ad account password.

Fake AI ad tools asked marketers to connect their ad accounts. Connect did not open Google. The page drew a Google window inside itself, with a fake address bar.[6]

VERDICT GIVEN — INFERRED

## ORDINARY SITE

Island's capture shows the browser never left the phishing domain, so a link-scanning control would see no sign-in hop.[6]

| LURE | HIDDEN BY | WHAT THE CONTROL SAW | VERDICT | WHAT THE PERSON SAW | OUTCOME |
|---|---|---|---|---|---|
| A new AI ad tool with a Connect button. | A live operator picking each screen. | One marketing site, no sign-in hop. | ORDINARY SITE | A Google or Okta sign-in window. | Passwords and codes, read live. |

### HOW IT GOT IN

01 The fake tools posed as ChatGPT, Gemini, Claude and Meta's Muse. Invitation emails brought victims in, but researchers give no detail on them.[6,21]

02 The fake address bar showed accounts.google.com or an Okta tenant. Island's screenshot after Connect still shows the phishing domain in the real address bar.[6]

03 A human operator saw each submission and chose the next screen: a rejected password, a Google approval prompt or an Okta push.[6]

### RECOMMENDED STEPS

REVIEW Ask marketing and finance to list every ad, payment and vendor account they sign in to, and mark each one outside company single sign-on. That list shows where a drawn sign-in window can still collect a password.[6]

CHANGE Move the owners of those accounts to passkeys or hardware security keys, and require Phishing-resistant MFA for them in Conditional Access (Google Workspace: 2-Step Verification methods set to Only security key). These work only on the real site, so a drawn window gets nothing it can use.[6,22,24]

STAFF Teach staff to read the address in the browser's own bar, not the one drawn inside the page. A drawn window can show any address, including accounts.google.com, while the real bar still shows the phishing site.[6]

### WHERE TO FIND IT

MICROSOFT 365

Entra ID > Conditional Access > New policy. Grant > Require authentication strength > Phishing-resistant MFA.[22] Put ad and vendor apps behind single sign-on (one company login for every app) so the policy covers them.[23]

GOOGLE WORKSPACE

Security > Authentication > 2-step verification. Methods: Only security key (also allows passkeys).[24]

### WHERE AEGISAI LOOKS

AegisAI's part depends on the link arriving by email, and researchers give no detail on the invitation emails. Where it does, AegisAI opens the link in a sandbox and follows it to the page with the Connect button. The verdict records every hop.

07 INBOX FLOOD, THEN TEAMS

ENTRY PATH: EMAIL > TEAMS

# Hundreds of sign-up emails, then "IT Support" on Teams.

Attackers posing as IT support messaged staff on Teams from outside organizations. In "a significant number of cases," customers reported an inbox flood before or alongside the chat.[7]

VERDICT GIVEN — REPORTED

## DELIVERED

Customers reported hundreds of newsletter or subscription confirmations within minutes, then the Teams message moments later.[7]

| LURE | HIDDEN BY | WHAT THE CONTROL SAW | VERDICT | WHAT THE PERSON SAW | OUTCOME |
|---|---|---|---|---|---|
| Hundreds of newsletter sign-ups. | Nothing in the email step. | Sign-up emails, one after another. | DELIVERED | An outside "IT Support" chat offering help. | Attack code running in under 15 minutes. |

### HOW IT GOT IN

01 Attackers appended "(Internal)" to Teams display names such as "IT Support". Teams still labeled the sender External.[7]

02 Quick Assist, the remote help tool built into Windows, was the tool used most often.[7,25]

03 After a victim engaged, customers saw the attacker's program (the payload) run in under 15 minutes.[7]

### WHERE TO FIND IT

MICROSOFT 365

Teams admin center > Users > External access: Allow only specific external domains.[26]

GOOGLE WORKSPACE

Apps > Google Workspace > Google Chat > External Chat Settings: Only allow this for allowlisted domains.[27]

### RECOMMENDED STEPS

CHANGE In Teams admin center > Users > External access, choose Allow only specific external domains and add only the partners staff chat with (Google Workspace: Google Chat external chat for allowlisted domains only). The fake "IT Support" chats came from outside organizations.[7,26,27]

CHANGE Uninstall Quick Assist on devices where IT uses a different remote tool. Arctic Wolf found it was the remote tool attackers used most often.[7,25]

STAFF Tell staff the one channel real IT uses to contact them, in onboarding and in the next all-staff note. An outside "IT Support" chat, especially after an inbox flood, gets reported, not answered.[7]

### WHERE AEGISAI LOOKS

AegisAI sees the flood, not the chat. Its agents investigate every inbound email, each sign-up message included. The decisive step happens in Teams, outside email. The two settings beside this box cover that step.

CONCLUSION

# Every control in this report passed what it checks.

In all seven techniques, the control that ran got a real or ordinary answer: a genuine sender, a clean link, a real Microsoft sign-in, a token that had already passed MFA. None of those checks asks what the message wants the reader to do.

Start with the recommended steps on pages 5 to 11. Each technique has three recommended steps from settings to review or change, to questions for your email security vendor, and points to brief staff on.

Six of the seven techniques have an email step. AegisAI's agents investigate every inbound email for what it asks the reader to do, with or without a link or file to scan. That covers the order confirmation from technique #1, the gated link in technique #2, the PDF in technique #3, the invitation in technique #5, the drawn sign-in window in technique #6 when its link arrives by email, and the inbox flood in technique #7.

Two steps happen outside email. Replayed session tokens (technique #4) never pass through the inbox. AegisAI's account takeover monitoring, in early access and review-only, reads sign-ins and mailbox activity and flags findings for your team to act on. The Teams chat in technique #7 is controlled by your Teams external access setting, not by email security.

## See what your controls let through.

An admin approves a read-only connection through the Microsoft 365 or Google Workspace API in about five minutes. There is no MX change, so mail routing stays as it is. The scan runs for five days beside your current tools and looks back over the last 14 days of delivered mail. Nothing is moved or deleted. At the end you review each finding with its evidence, and the report is yours to keep.

Book a demo →

SOURCES

# Sources

Written by the AegisAI Threat Watch team. Published November 2026.


01 **Help Net Security**, attackers abuse legitimate platforms to deliver phishing emails that pass authentication checks, 22 October 2026

02 **BleepingComputer**, takedown of the network behind 1 billion phishing messages, 14 October 2026

03 **ANY.RUN**, major cyber attacks in October 2026 (Mir0Auth section), 3 November 2026

04 **Huntress**, token replay and spraying, 22 October 2026

05 **CyberScoop**, TA419 targets US AI policy experts, 1 October 2026

06 **Island**, the fake AI ads campaign, 6 October 2026

07 **Arctic Wolf Labs**, Teams external chat for initial access, 26 October 2026

08 **Microsoft Learn**, configure anti-spam policies

09 **Microsoft Learn**, manage mail flow rules in Exchange Online

10 **Microsoft Learn**, the Tenant Allow/Block List

11 **Google Workspace Admin Help**, spam filter allowlists

12 **Microsoft Learn**, Safe Links policies

13 **Google Workspace Admin Help**, advanced phishing and malware protection

14 **Microsoft Learn**, block authentication flows

15 **Microsoft Learn**, revoke user access

16 **Google Workspace Admin Help**, API controls

17 **Google Workspace Admin Help**, a user's security settings

18 **The Register**, TA419 AI policy phishing, 1 October 2026

19 **Microsoft Learn**, Set-ExternalInOutlook

20 **Microsoft Learn**, anti-phishing policies

21 **The Hacker News**, fake AI ad tools, 6 October 2026

22 **Microsoft Learn**, require phishing-resistant MFA

23 **Microsoft Learn**, enable single sign-on for an enterprise app

24 **Google Workspace Admin Help**, deploy 2-Step Verification

25 **Microsoft Learn**, use Quick Assist to help users

26 **Microsoft Learn**, Teams external access

27 **Google Workspace Admin Help**, external chat settings

28 **Microsoft Learn**, manage emergency access admin accounts

AegisAI

ABOUT AEGISAI

# The first AI lab for email security

## Every email, investigated.

AegisAI puts an investigation on every email. Its agents check who really sent a message, follow every link to the final page and read what the message is asking someone to do. Proteus, AegisAI's own model, weighs every claim against the evidence, and every verdict ships with its reason.

AegisAI connects to Microsoft 365 and Google Workspace by API in about five minutes, with no MX change, and runs beside your current tools. It is built by the team behind Google Safe Browsing and reCAPTCHA, and backed by Battery Ventures, Accel and Foundation Capital.

ABOUT 5 MINUTES TO CONNECT | 5-DAY READ-ONLY SCAN | NO MX CHANGE | SOC 2 TYPE II

Book a demo →

© 2026 AegisAI, Inc.