Material

STOP SECURING EMAIL.

Start securing the cloud workspace.

STOP SECURING EMAIL.

Email security is locked in a permanent stalemate. It's foundational, but not enough in the age of AI. Attackers generate unlimited bespoke attempts at near-zero cost, defenders catch most of them, and some predictable fraction always gets through.

This isn't because your filters are bad. Two decades of secure email gateways, DMARC, and awareness training have made the inbox the most scrutinized square inch in security, and Verizon's 2026 Data Breach Investigations Report notes that phishing's share of initial access "has barely moved over the past few years."[1] Credential abuse touches 39% of breach progressions, more than any other vectors.[2] The most defended door in security keeps losing the house.

Here's why: the inbox was the perimeter when email was where work lived. Work moved into shared drives, always-on sessions, and a mesh of apps and agents authorized with a click. Attackers moved with it. AI made their opening move free, credentials became the currency, and third-party tokens became the side door. The attack is everything that happens after: in the accounts, the sessions, the files, and the sprawling web of connected apps that make up your cloud workspace.

If you run security for a company built on Google Workspace or Microsoft 365, and especially if "run security" means you and maybe two other people, this paper is a playbook about where to spend your scarcest resource: attention.

Stop optimizing the inbox. Start securing the workspace.

Table of Contents

Four reasons to stop securing email

Modern, AI-driven attacks deliberately span email, identities, data, and connected apps. Defending any single layer in isolation leaves the rest exposed.

01. Phishing catch rate doesn't decide the outcome, the blast radius after delivery does.

02. Inbound email security tools have nothing to say about sessions, mail rules, OAuth grants, or data sitting at rest.

03. The risk has moved to files, accounts, and connected apps: Breaches like Salesloft Drift, Composio, and Vercel touched the inbox only after using a compromised token to gain workspace access.

04. AI agents work through the workspace, not email alone.

The inbox was the perimeter when email was where work lived. Work moved into shared drives, always-on sessions, and a mesh of apps and agents authorized with a click. Attackers moved with it. AI made their opening move free, credentials became the currency, and third-party tokens became the side door. The attack is everything that happens after: in the accounts, the sessions, the files, and the sprawling web of connected apps that make up your cloud workspace.

13x

increase in phishing post-AI[3]

25%

of OAuth connected apps hold restricted Google scopes, most commonly full Gmail and Drive access together: read every email, send as the employee, reach every file.[4]

45%

of security professionals are neglecting OAuth app governance, while 33% rely on manual processes. Only a small minority have implemented dedicated tools or automated controls.[5]

1,100%

Sensitive files in shared drives grew 1,100% in eight months, 3x the rate of files overall. Over a quarter of all Drive files now contain sensitive data.[6]

How AI broke the economics of the inbox

For years, the implicit bet behind email-centric security was that filtering could outpace attacker volume. Generative AI ended that bet, not by making phishing smarter than filters, but by making infinite, polished variation effectively free.

> By April 2025, at least 51% of spam and 14% of business email compromise attempts were LLM-generated, with BEC nearly doubling in a single year.[7]

Those figures come from an academic study of 481,558 real malicious emails collected across thousands of organizations, presented at the 2025 ACM Internet Measurement Conference, and they are the conservative estimate.[7] The same study found LLM-written attacks are measurably more formal, more grammatically correct, and more linguistically sophisticated than their human-written counterparts.[8] Every heuristic your users were trained on is gone: the typos, the broken English, the clumsy urgency. The DBIR authors put it more bluntly, joking that detection guidance should change from "does it contain many typos" to "does it contain em dashes."[9]

The study also caught attackers doing something more interesting than writing better emails: using LLMs to mass-produce rewritten variants of the same message to slip past signature- and volume-based filters. In the largest campaign clusters analyzed, 79% and 52% of messages were LLM-generated rewrites, up to ten times the baseline rate.[10] Meanwhile, Verizon's analysis of real attacker usage of AI found phishing was the single most common AI-assisted initial access technique, at 44%.[11]

None of this means your email filter is failing. It means the war of attrition is permanent, and it is no longer the war that decides outcomes. The metric that matters isn't your catch rate; it's the blast radius of the phish that inevitably lands. And that is a question about your workspace, not your gateway.

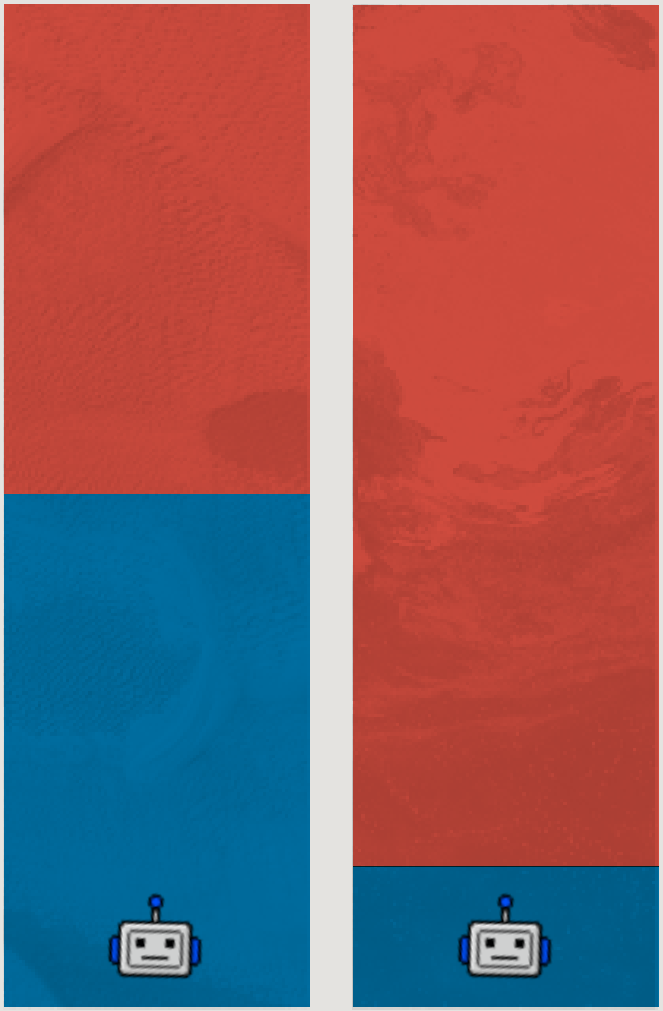

The breach happens after the click *(or with no click at all)*

Follow the DBIR's data past the inbox and a consistent picture emerges: modern breaches are account and access stories, not message stories.

Stolen credentials lead every other attack action, showing up in 39% of breaches at some point in the attack chain.[2] The pipeline feeding it runs quietly in the background: among ransomware victims with a prior credential or infostealer leak, half were hit within 95 days of that leak.[12] Small organizations, the ones least likely to be watching, experienced a median of seven credential-leak events in a year.[12]

> Among ransomware victims with a prior credential leak, 50% were compromised within 95 days of the leak. Small orgs saw a median of seven leak events per year.[12]

Then there is the path that never touches your inbox at all. Third-party involvement in breaches jumped 60% year over year, reaching 48% of all breaches.[13] The emblematic case: attackers compromised OAuth tokens from the Salesloft Drift integration and used them to pivot into the Salesforce environments of companies including Google, Zscaler, and Cisco.[14] No phishing email. No malware. Just a trusted app token, quietly authorized long ago, doing exactly what tokens do. Email security had nothing to say about any of it.

This is the blind spot in the point-solution model. A tool built to watch the inbox stops looking at the moment of delivery, precisely when the consequential part begins: the session that gets hijacked, the OAuth grant that gets abused, the years of sensitive mail sitting in a compromised mailbox, the Drive folder shared one permission too wide. Nobody is monitoring downstream of email, and downstream of email is where the losses are.

Your workspace has already sprawled past the inbox

The workspace isn't just where attacks land; it's an attack surface that has been quietly compounding on its own. Shared files, OAuth apps, app-specific passwords and magic links, and more: the cloud workspace is a broader and deeper threat landscape than inbound attacks alone.

Google Drive often contains far more sensitive data than organizations realize. In Material's investigations, we found that well over a quarter of Google Drive files contain sensitive data. And it continues to grow over time. In an eight-month sample window, total files grew by nearly 400%, and total sensitive files grew by over 1,100%.

That is, of course, the whole point of shared files: to create documents and share them. It's maintaining visibility into exactly what is shared and with who that becomes critical at scale.[6]

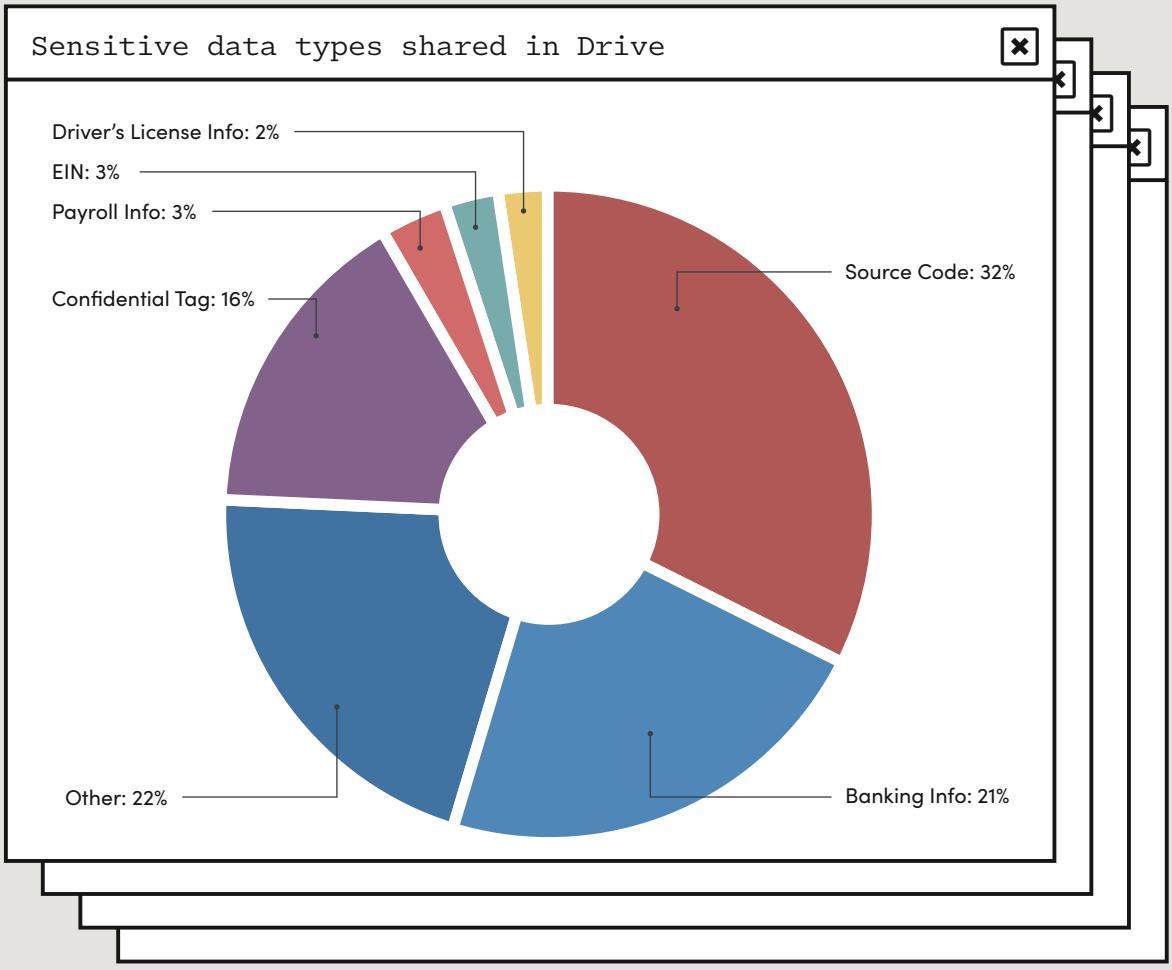

"When you revoke a user's access or offboard a vendor, you're thinking about accounts and credentials — but OAuth grants operate on a different lifecycle. These are zombie connections — technically authorized, practically abandoned, and invisible to most of the controls we rely on."

Chaim Sanders
CISO, Lyft

Connected apps can serve as another unmonitored attack surface. Material Security's analysis of 22,332 OAuth-connected applications across 21 enterprise Google Workspace environments found a layer of standing, unmonitored access that most teams have never inventoried:[4]

Nearly half of all connected apps (47%) showed no active usage in 90 days, yet their authorizations remain fully intact. More than 1,000 apps had zero active users but live tokens; 43.5% of those "zombie" apps hold sensitive or restricted scopes. And 24.5% of all apps hold restricted Google scopes, most commonly Gmail and Drive together, which the report calls "the highest-consequence permission profile available": read all of an employee's email, send on their behalf, and access every file in their Drive.[4]

Lyft CISO Chaim Sanders calls them zombie connections: "technically authorized, practically abandoned, and invisible to most of the controls we rely on."[15]

AI adoption is pouring accelerant on this. Of the AI apps found connected to corporate workspaces, 91% appeared within the last 16 months, more than half hold sensitive or restricted scopes, and most were never formally approved: someone clicked authorize, and that was it.[16] Verizon sees the same thing from the user side: 67% of users access AI services from non-corporate accounts on corporate devices, and 45% of employees are now regular AI users at work, up from 15% a year earlier.[17] Unlike last decade's SaaS sprawl, these apps don't just read data and make suggestions. AI agents act, and what they do is determined at execution time, not at the moment someone granted them access.[4]

> 91% of AI apps connected to corporate workspaces appeared in the last 16 months. More than half hold sensitive or restricted scopes. Most were never formally approved.[16]

The scale of what's exposed tends to surprise even the teams who suspect it. In one recent evaluation, Material surfaced more than 20,000 sensitive documents shared externally, a corpus the organization had no visibility into before looking.[18] None of this describes compromise. It describes, in the words of the OAuth report, "the conditions that make compromise harder to detect and easier to sustain."[4]

Secure the workspace, not the message

"Stop securing email" doesn't mean turn off your filters. It means stop treating the inbox as the perimeter and start defending the workspace as one interconnected system: mail, files, accounts, sessions, and connected apps.

In practice, five principles:

01

Email isn't the only path to an account takeover.

Grade your program on what an attacker can do once they gain control of an account, not on phishing catch rate alone. A phish will land or a token will be stolen; the question is whether the attack goes anywhere after initial access.

02

Protect data where it lives.

Years of PII, credentials, and contracts sit in mailboxes and Drive at rest. A compromised account should yield a locked vault, not an open archive.

03

Govern the grants

Inventory OAuth connections, revoke on offboarding, set a dormancy threshold (90 days is a reasonable default), and build a sanctioned path for AI tools before the next wave arrives.[4]

04

Contain the account

Account takeover is a when, not an if. Detection and response inside the workspace (sessions, mail rules, app grants, sharing changes) is what turns an incident into a non-event.

05

Demand one view.

Email, files, and identity aren't three products' worth of problems; they're one attack path. Point tools that each watch one door leave the hallways between them dark.

For a lean team, this is the practical appeal: one platform watching one interconnected surface is less work than five tools watching five fragments of it. This isn't a maturity model that assumes a SOC you don't have. It's a reallocation of attention toward where the data says breaches actually unfold.

Security leaders have already moved past the inbox

CISOs, incident responders, and the press have reached the same conclusion: the inbox is where the attack starts, not where it ends.

"Expanding beyond email to secure the entire cloud workspace is a game changer. By safeguarding not just email, but also documents, accounts, and configurations, they're paving the way for smarter automated security solutions that reduce risks and save time. Very impressive to see a security platform that meets the reality of how work gets done today."

Erik Wille
CISO, Cabinetworks Group

"The team at Material Security has built a really interesting product that completely bypasses how everybody has thought about email security historically."

Troy Wilkinson
CISO, Interpublic Group (IPG)

"They're solving one of the most important problems in security: how to survive even after an attacker breaks into a mailbox."

Feross Aboukhadijeh
Founder, Socket

"A really strong team taking a novel approach of protecting the sensitive data, rather than trying to detect if an email is good/bad."

Dan Wendlandt
Co-founder & CEO, Isovalent

"In a past life, I was an Incident Responder helping many different victim companies. Material Security would have been a difference maker in stymying the attackers."

Ernest Liu
CISO, UTA

"As a small team, I needed a tool that worked great out of the box... but still offered customization when needed. Material delivered."

Frank Wang
Lead Security Engineer, Surge AI

THE BOTTOM LINE

Email security fought its war to a permanent draw. *The war moved.*

Secure the cloud workspace, the email, files, and accounts your company actually runs on, or accept that the most important part of your attack surface is the part nobody's watching.

Material

www.material.security

Ready to stop securing email?

Show the world with your very own "Don't Email Me" shirt. Screenshot this page, post it on LinkedIn and tag Material Security.

Sources

1. Verizon, 2026 DBIR, p. 26 (phishing's share of initial access; AI-assisted techniques).
2. Verizon, 2026 DBIR, p. 16 (credential abuse at any point in breach progression: 39%).
3. SlashNext, The State of Phishing 2023 (1,265% increase in malicious phishing emails since the launch of ChatGPT, reported Oct. 2023, dataset Q4 2022–Q3 2023).
4. Material Security, OAuth & Google Workspace Risk Report (22,332 apps across 21 enterprise environments; 47% dormant 90+ days; 1,064 zombie-token apps, 43.5% with sensitive/restricted scopes; 24.5% with restricted scopes; Gmail+Drive as highest-consequence profile; pp. 2–11).
5. Material Security, "Automating OAuth Grant Management: Material's Research Shows the Growing Gap between Awareness and Action" (April 2026; survey of 45 security leaders — 45% neglecting OAuth app governance, 33% relying on manual processes).
6. Material Security, "Exploring Sensitive Data Trends: What We've Learned from Protecting Email and Files" (Jan. 2025; Drive files grew ~400% and sensitive files grew over 1,100% in an eight-month sample window — nearly 3x the overall growth rate; over a quarter of Drive files contain sensitive data).
7. Hao, W., et al., "Do Spammers Dream of Electric Sheep? Characterizing the Prevalence of LLM-Generated Malicious Emails," Proc. ACM Internet Measurement Conference (IMC '25), 2025. Abstract & §4.3: ≥51% of spam and 14% of BEC LLM-generated as of April 2025 (vs. 16.2% and 7.6% in April 2024); dataset of 481,558 malicious emails, Feb 2022–Apr 2025.
8. Hao et al., IMC '25, §5.2, Table 3 (LLM-generated emails significantly more formal, fewer grammar errors, more sophisticated language; $p < 0.001$).
9. Verizon, 2026 DBIR, p.16, footnote 7 (em-dash detection joke).
10. Hao et al., IMC '25, §5.3 (78.9% and 52.1% LLM prevalence in largest rewrite clusters vs. 7.8% average).
11. Verizon, 2026 DBIR, p. 26, Figure 27 (phishing = 44% of GenAI-assisted initial access techniques, n=837).
12. Verizon, 2026 DBIR, pp. 44–45 (50% of ransomware victims with prior credential/infostealer leak compromised within 95 days; small orgs: median 7 credential-leak events/year).
13. Verizon, 2026 DBIR, p. 11 (third-party involvement up 60% year over year, reaching 48% of breaches).
14. Verizon, 2026 DBIR, pp. 21, 109 (Salesloft Drift OAuth token compromise used to access Salesforce instances incl. Google, Zscaler, Cisco).
15. Chaim Sanders, CISO, Lyft, quoted in Material Security, OAuth & Google Workspace Risk Report, p. 9.
16. Material Security, OAuth & Google Workspace Risk Report, pp. 3–5 (91% of AI apps appeared in last 16 months; >50% hold sensitive or restricted scopes; most never formally approved).
17. Verizon, 2026 DBIR, p. 13 (Shadow AI: 67% of users on non-corporate AI accounts from corporate devices; 45% regular AI users, up from 15%).
18. Material Security customer evaluation data (proof-of-concept discovery of 20,000+ externally shared sensitive documents).

The companies that run on Material

The companies defining the AI era secure their cloud workspace with Material.

OpenAI reddit PagerDuty gusto POSTMAN

ANTHROP\C Figma Globe Quora

Amplitude MARS lyft NAVAN SIERRA

Material

See what's already connected to your workspace.

Material Security provides detection and response across email, files, connected apps, and accounts in Google Workspace and Microsoft 365, deployed via API in minutes with no MX record changes.

Get a demo ↗