Material

Already Inside:

AI Agents Are Accessing Sensitive Data Across Your Google Workspace

Visibility is table stakes. Control is the point.

ACROSS THE BOARD:

Sensitive Data Exposure at Enterprise Scale

AI agents are no longer a future consideration for enterprise IT. They are operating inside your Google Workspace today. And they are reading your data.

The pace of AI adoption has outrun the policy frameworks meant to govern it. Employees connect AI tools, coding assistants, document analyzers, note-takers, enterprise search products, directly to their Google accounts through OAuth. No ticket. No IT approval. No visibility into what the tool will access.

This creates a question that most organizations have not yet asked: **if AI agents are going to operate inside your cloud workspace, do you know what they are operating on?**

53.4%
public apps have restricted permissions

The majority of applications users are connecting to have deep access into files and email.

373M+
sensitive emails

Over 2.3 million emails per organization with usernames, passwords, employee numbers, financial data, and more.

475M+
sensitive Drive files

Nearly 3 million sensitive files per organization with sensitive, confidential, or regulated data.

In August 2026, we analyzed 159 organizations to answer that question. What follows is what we found.

FINDING 01:

AI App Adoption Is Already at Scale, Across Every Organization

In order to understand the scale of sensitive data exposure, we need to first understand the scale of AI adoption. Most companies today are adopting a range of sanctioned enterprise AI tools–but the approved tools are only the tip of the iceberg.

Employees are not waiting for IT to approve AI tools. They are connecting through OAuth, the same mechanism behind every "Sign in with Google" button on the internet. The tool gets access to Drive, to Gmail, to calendars, to everything the employee can reach.

Across the organizations we analyzed, 978 distinct applications with active OAuth connections have been connected this way. Not deployed by IT. Not inventoried by security. Authorized by individual employees making individual decisions about tools they wanted to use. This count reflects apps with observed grant activity - dormant or historically connected apps are not included, meaning the actual footprint is larger.

The top application alone has been authorized by **152,085 unique users** aggregated across all 159 organizations, with 655,430 total OAuth authorization events. This is a cross-customer total: individual organizations see a fraction of this number in their own OAuth reports. One app. More user connections than most enterprise software sees in its entire deployment lifetime. On average, that is nearly 1,000 users per organization connecting to a single unidentified application.

978
active OAuth apps

152K
users, top app

655K
grant events, top app

978 Active Apps. Roughly 6 Per Organization. Most of Them Unnamed.

Of those 978 applications, a significant portion were never public products to begin with. They are private OAuth projects built by employees or teams within the organization itself. What appears in connection logs is a technical project number, not a vendor name, not a product, not something a security team can look up in any policy document. The only people who knew what these apps did were the ones who built them. If they have left, nobody knows what the app accesses, why it exists, or whether it is still needed. The app still holds its OAuth grant. It still has access. There is simply no way to evaluate whether that access is appropriate.

The Most Widespread App Has No Name

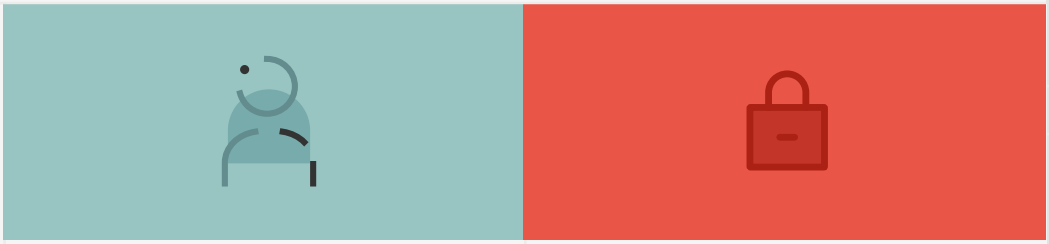

	Users	OAuth Grants
Unidentified These numbers compile **many unknown apps** across all organizations. This represents the most widespread unidentified app. It is a **compilation of many private or internal apps** that appear in OAuth connections but have no name or vendor.	152,085	655,430
OpenAI	9,725	23,982

What OAuth Access Means in Practice

When an employee authorizes an AI tool via OAuth, the tool reads with that employee's eyes. It sees every file they can open. Every email in their inbox. Every document shared with them. There is no separate permission set for AI. There is no approval step between the employee's access and the tool's access.

An employee with access to payroll files, source code, and financial reports authorizes a productivity AI. That AI now has access to payroll files, source code, and financial reports. No additional review. No visibility to IT. No audit trail that distinguishes the AI's activity from the employee's.

Across 159 organizations, with **373 million sensitive emails and 475 million sensitive Drive files** already in scope, this is not a theoretical risk. The data exists. The tools are connected. The access has already been granted.

FINDING 02:

Every Organization Has Sensitive Data in Email, at Scale

Of the organizations we studied with email monitoring in place, every single one of them had sensitive data in email. Not most of them: all of them. There is no such thing as a clean inbox at enterprise scale.

Of the 159 organizations in this dataset, 158 had email monitoring enabled. Across those 158, we counted more than **373 million** emails containing sensitive content. Bank account details. Employee credentials. Financial records. Confidential business communications. Moving in and out of inboxes every day, in every organization, without anyone tracking where it goes.

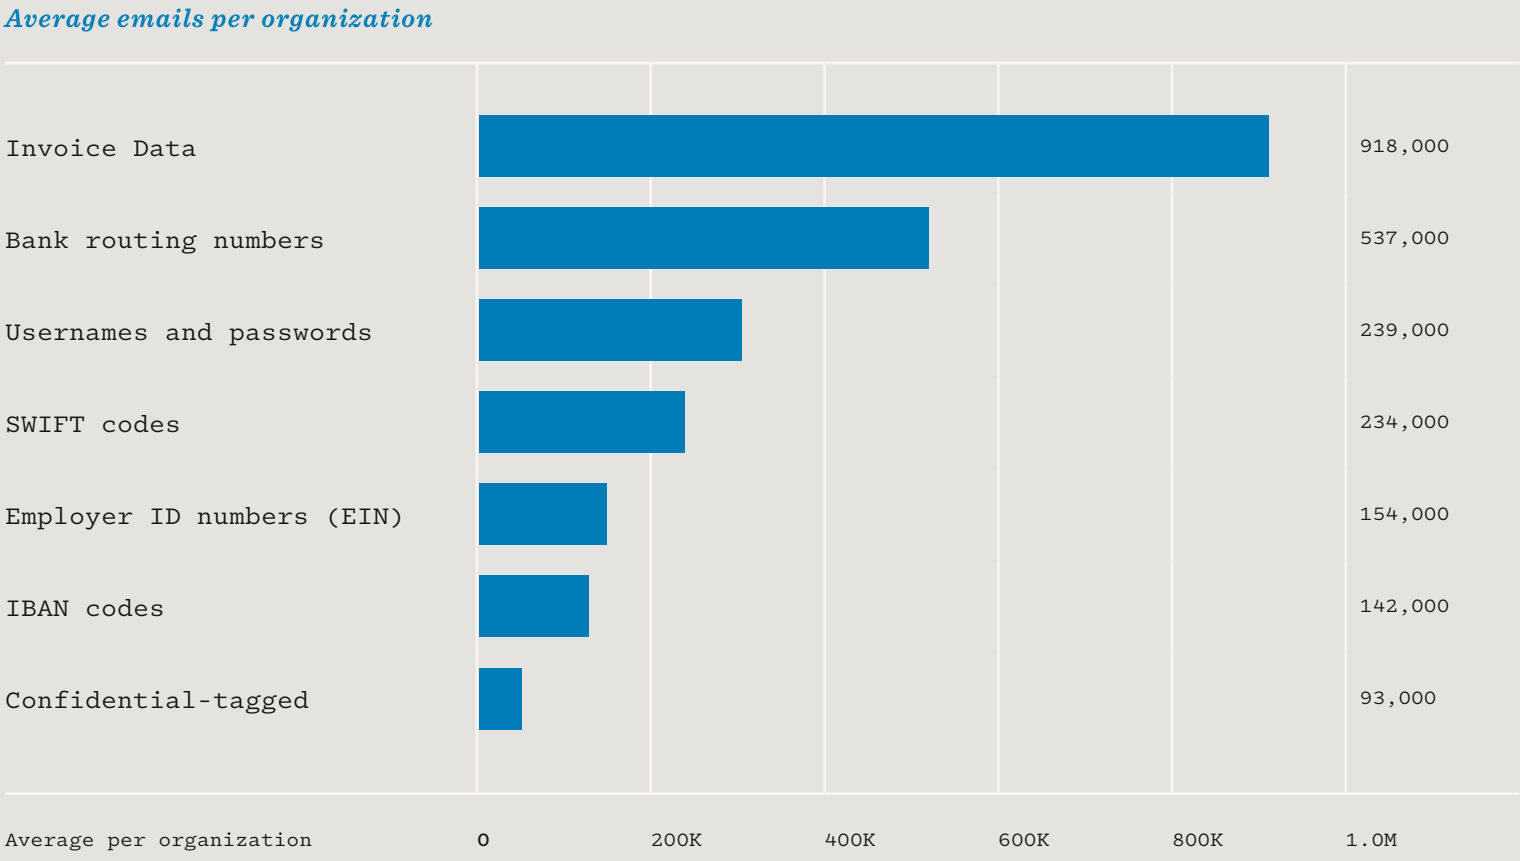

46 Million Emails Contain Passwords

To put the scale in context: the average organization in this dataset had 918,000 emails containing invoice data, 537,000 with bank routing numbers, and 293,000 with usernames and passwords. These are not outlier organizations. This is the baseline.

Forty-six million emails containing usernames and passwords. Across those 158 organizations with email monitoring–that's an average of nearly 300,000 emails containing usernames and passwords per organization. Not because someone made a mistake. Because that is how enterprise systems are built. Onboarding workflows send credentials to new users. Vendor portals email login details to account managers. IT systems confirm password resets in plaintext. Nobody decided this was a good idea. It just became the default.

The threat model is straightforward: an AI tool with inbox access does not need to hack anything. It reads what is already there. And what is already there, in the inboxes of every organization into which we had visibility, is sensitive data that most security teams have never fully measured.

FINDING 03:

Sensitive Files in Drive Are a Universal Condition

Drive is where organizations keep their work. Contracts. Payroll records. Source code. Financial models. Engineering specs. It is also where that work sits indefinitely, accessible to anyone with the right permissions, long after the project it was created for has ended.

Of 159 organizations with Drive monitoring, 100 have files flagged for sensitive content by Google's DLP (Data Loss Prevention) classifiers. That is 63 percent. And it reflects only what Google's native, general-purpose classifiers can surface - broad pattern-matching rules built for common data types, not tuned to any individual organization's data. The actual number is higher.

375M
source code files

42M
financial reports

17M
bank routing files

16.7M
files with SSN data

12M
password files

7.8M
payroll files

375 Million Source Code Files

Source code is the largest single sensitive data category in Drive across all 159 organizations, averaging 2.4 million source code files per organization. That number is not surprising once you think about how engineers actually work. They store code in shared drives. That code contains API keys, database credentials, deployment configurations, and proprietary logic. It gets shared across teams, across contractors, across time. Nobody deletes it when the project ends.

Average sensitive Drive files per organization

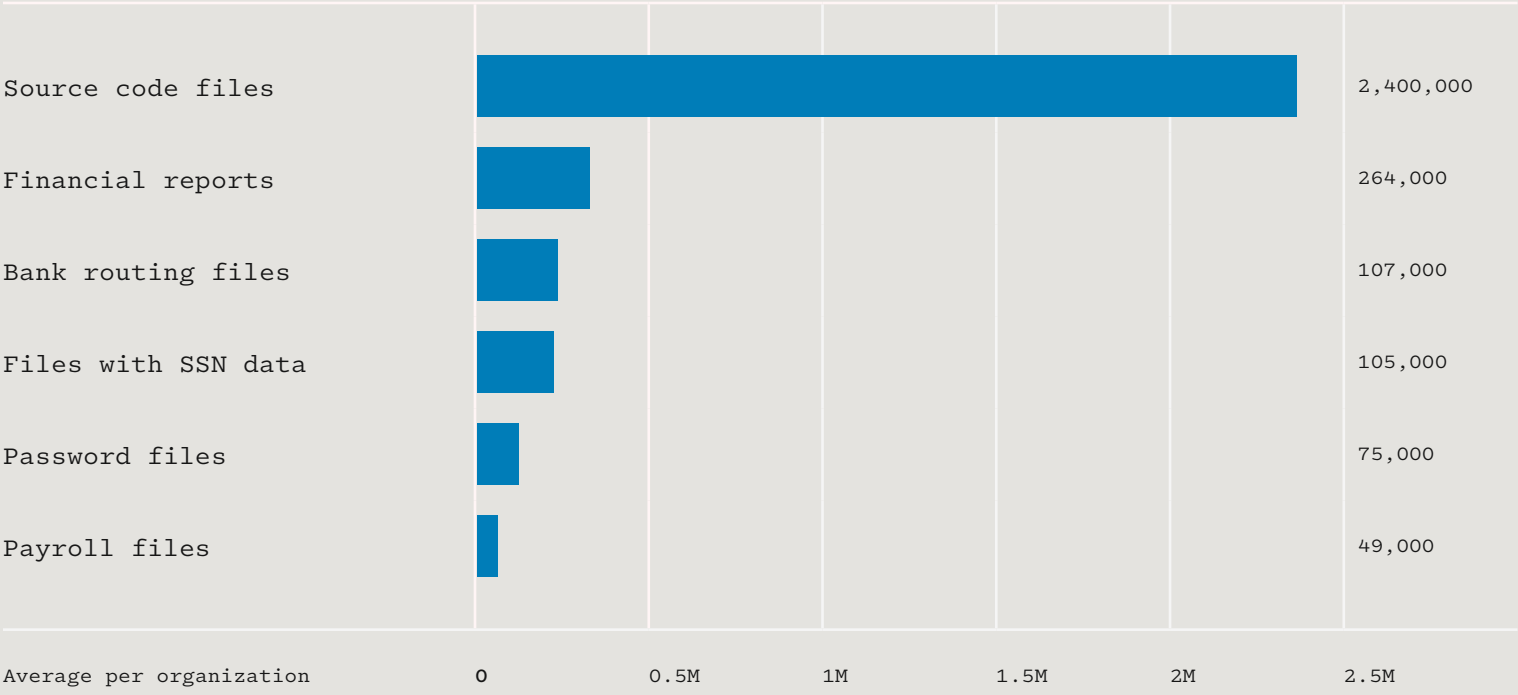

A single source code file can contain more exploitable information than an entire financial folder. API keys. Database connection strings. Encryption keys. These are not accidents. They are the daily output of engineering teams doing exactly what they are supposed to do.

Get A Demo →

17 Million Files With Bank Routing Numbers

Seventeen million files containing bank routing numbers, an average of 107,000 per organization. 16.7 million with Social Security Number data, averaging 105,000 per organization. 7.8 million payroll files, an average of 49,000 per organization. This is not the result of careless behavior. This is what happens when HR runs payroll through Drive, when finance shares bank details to set up vendor payments, when legal stores employee agreements in shared folders. Normal operations. Invisible risk.

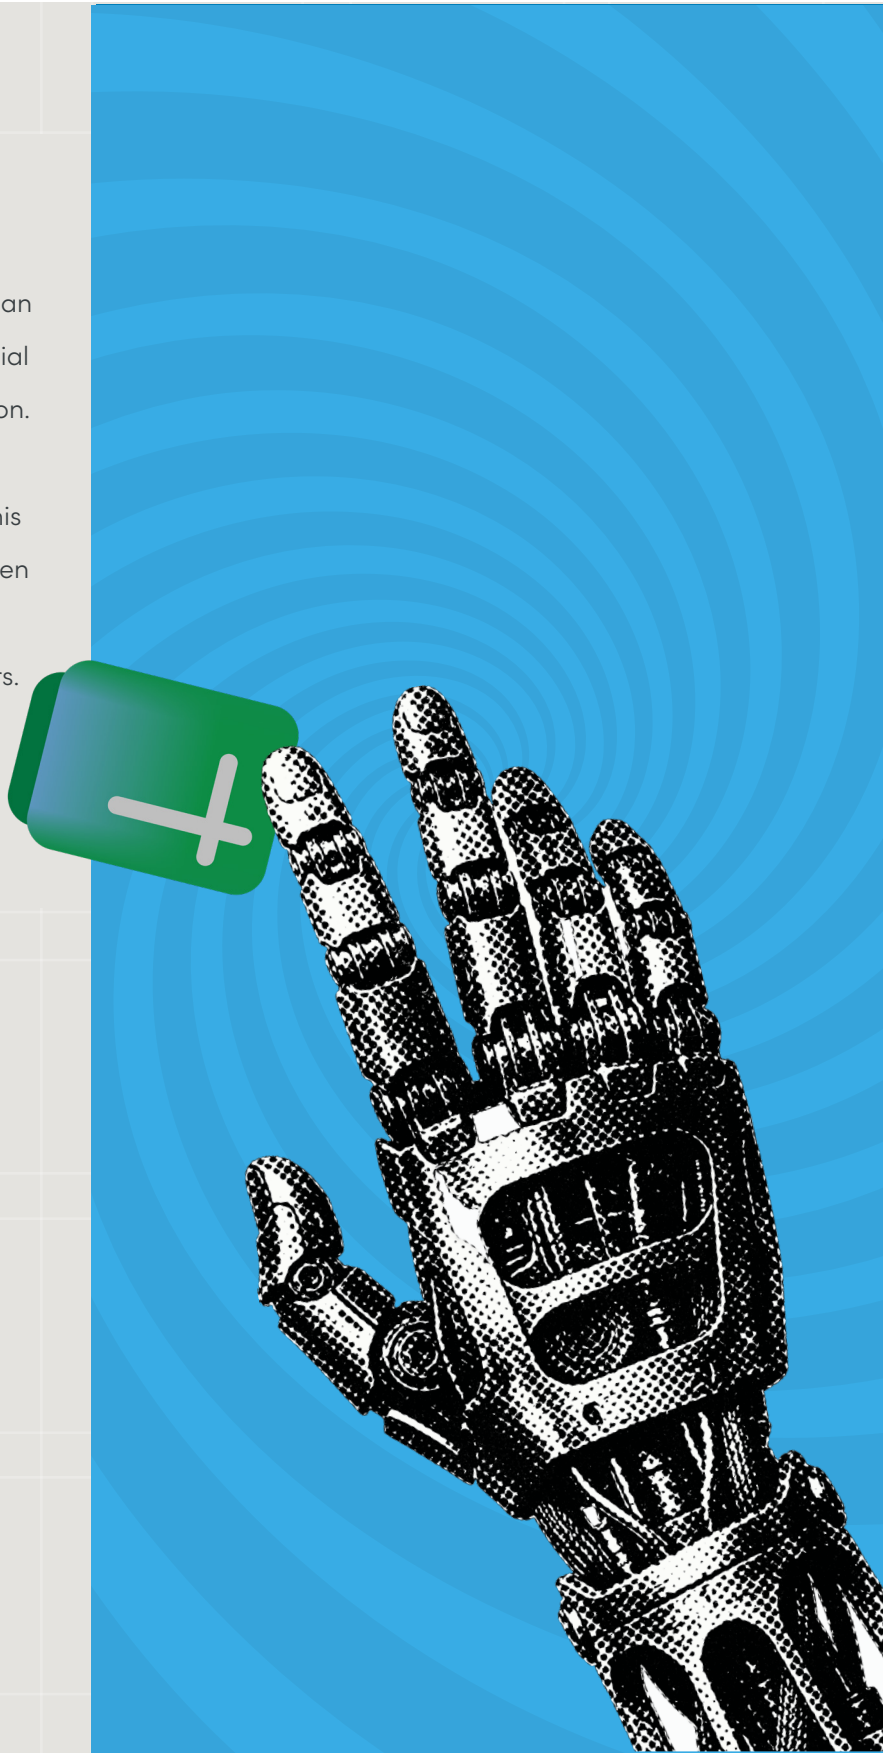

FINDING 04:

Organizations Are Starting to Course-Correct, But the Sharing Surface Keeps Growing

There is a pattern emerging across the organizations we analyzed. The pace of restricting access to Drive files is accelerating. In mid-2025, restricting events accounted for roughly 14 percent of all access-change activity. By spring 2026, that figure had risen to over 44 percent. Organizations that have visibility into their data exposure are acting on it.

But this story has two sides. At the same time restricting activity is accelerating, the absolute volume of sharing is still growing. More files are being shared every month than the month before. The course-correction is real. It just has not caught up to the pace of expansion.

Month	*Broadening Events*	*Restricting Events*	*Restricting Share*
June 2025	**181 million**	26 million	14%
Sept 2025	**248 million**	40 million	16%
Dec 2025	**406 million**	98 million	24%
March 2026	**550 million**	241 million	44%
June 2026	**633 million**	269 million	43%

Restricting events grew more than 10x between June 2025 and spring 2026, nearly 270 million per month at peak. This is not normal operational noise. It reflects what becomes possible when security teams can see what they have. Organizations without that visibility have no signal to act on.

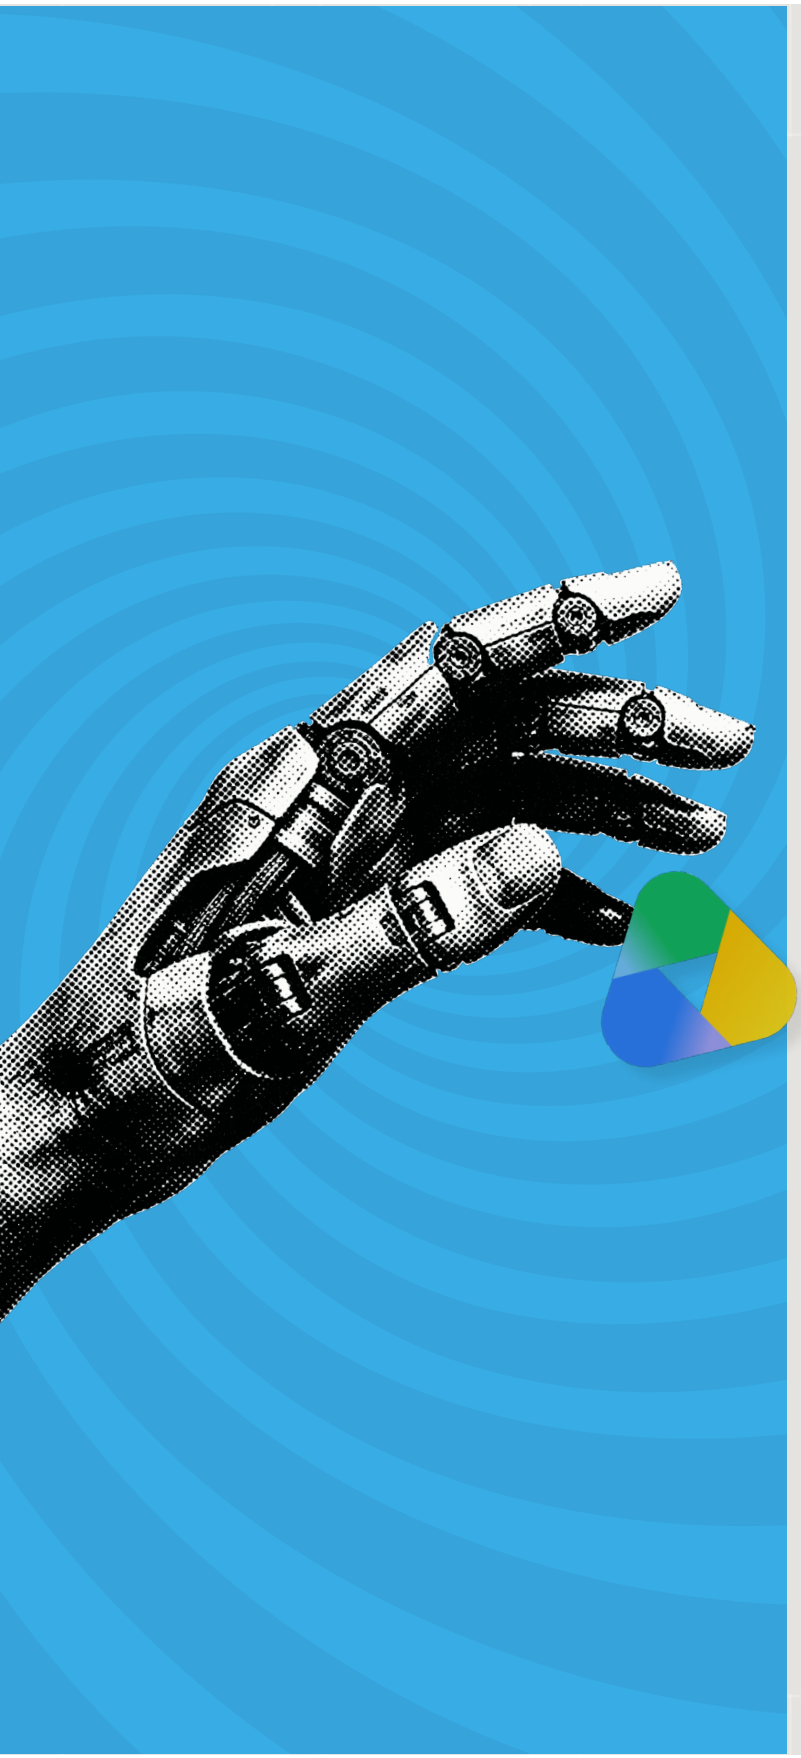

The Catch: AI Doesn't Care When Access Was Granted

Here is the problem with the course-correction story: it addresses sharing controls. AI tools operate on user permissions. The distinction matters.

When an employee connects an AI tool via OAuth, that tool inherits everything the employee can access at that moment, across email and Drive simultaneously. Files shared last year that were never cleaned up. Emails from three years ago that were never archived. Documents from departed employees that were never deleted. None of this shows up in a restricting event. It just sits there, in scope, readable.

The organizations tightening their Drive permissions are doing the right thing. But the window of exposure that already exists, 373 million sensitive emails and 475 million sensitive Drive files accumulated over years, does not shrink when a sharing policy changes. It only shrinks when someone measures it, maps it, and removes it.

What This Means for AI Adoption

Enterprises are not deciding whether to adopt AI. That decision has already been made, in most cases by the employees who started connecting tools to their Google accounts months or years ago. The actual decision in front of IT and security leadership today is different: whether to adopt AI with visibility into what it can reach, or without it.

Responsible AI adoption at the enterprise level starts with three questions:

1. What sensitive data exists in the environments AI agents will access?
2. Which AI apps are connected, and what are they actually reading?
3. Is there a way to detect and respond when access patterns change?

The data needed to answer all three questions already exists inside every Google Workspace. The organizations that measure it before deploying AI broadly are the ones that can scale AI confidently, because they know exactly what they are scaling into.

Across 159 organizations, the pattern is the same regardless of industry, size, or geography. Sensitive data accumulates in email and Drive faster than anyone measures it. AI tools arrive faster than any policy can track them. And the workspace AI operates in reflects years of history, not just last week's files. AI adoption is not the risk. Adopting AI without knowing what AI can read: that is the risk.

You can't govern what you can't see.

AUTHOR :

Belem Relegado

Staff Threat Research Engineer, Intelligence Reporting and the Material Security Threat Research Team

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 · X · 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 →