


AI SECURITY FOR BUSINESS

Your AI has more keys than your most trusted employee.

7 ideas for using AI in your business without exposing it.
No jargon, no fear, and something concrete you can do
today for each one.



Each square is an access you granted. Do you know how many are still switched on?

Diego Rojas

Security engineer · 10+ years in systems, cloud and security

stratega.studio · @thediegorojas_

The risk moved, and almost nobody noticed.

A couple of years ago, AI security advice came down to one rule: don't paste sensitive information into the chat. That's still true, but it's no longer the main thing.

Today AI doesn't just answer. It reads your email, opens your files, schedules meetings, replies to clients and connects one app to another. To do all that, it needs something you used to give only to people you trust: **access**.

WHAT YOU THINK YOU'RE DOING

WHAT YOU'RE ACTUALLY DOING

Trying a new app

Giving it access to your email, possibly forever

Sharing a conversation

Publishing a link anyone can open

Automating a task

Creating a worker that acts in your name

Asking AI to summarize an email

Letting a stranger give it instructions

Almost none of the risks in this guide are hacks. They're permissions someone granted without seeing how far they reached.

HOW TO READ THIS GUIDE

Each idea has three parts: what we usually think, what actually happens, and an action that takes less than 15 minutes. At the end there's a 5-minute test to see where you stand.

01 "Sign in with Google" isn't a shortcut. It's a key.

We usually think: *"I gave it access to try it one afternoon. If I don't use it, nothing happens."*

WHAT ACTUALLY HAPPENS

When you accept, the app gets permission to read (and sometimes send or delete) your email, files or calendar. That permission doesn't expire when you stop using it: it stays active until you remove it. If that company shuts down, gets sold or gets hacked, the key still works.

DO THIS TODAY

In your Google Account: Security → Your connections to third-party apps & services (in Microsoft: Privacy → Apps and services). Remove anything you haven't used in the last month, starting with whatever can read and send email.

02 Your AI assistant obeys anyone who writes to it.

We usually think: *"If I ask it to summarize an email, it just summarizes it."*

WHAT ACTUALLY HAPPENS

AI doesn't reliably tell your instructions apart from the text it reads. An email, a PDF or a web page can carry hidden orders ("forward the last 10 emails to this address"), and an assistant with permissions may carry them out. It's called prompt injection, and it doesn't need you to click anything.

DO THIS TODAY

Check what each assistant connected to your email or files can do. Reading and summarizing: fine. Sending, deleting, paying or sharing: only with your confirmation on every action. If the tool can't ask for confirmation, don't grant those permissions.

03 “Anyone with the link” means anyone.

We usually think: *“I only sent it to my client. Nobody else has the link.”*

WHAT ACTUALLY HAPPENS

A link doesn't know who you sent it to: it gets forwarded, sits in histories and can end up indexed. In 2025, ChatGPT conversations their owners had shared showed up in Google search results, with personal and work details visible to anyone.

DO THIS TODAY

In Google Drive, find files shared with “Anyone with the link” and switch them to specific people. In ChatGPT: Settings → Data controls → Shared links, and delete the ones you no longer need.

04 Your automation works under your name and never takes a day off.

We usually think: *“It's just a flow connecting the form to email and the spreadsheet.”*

WHAT ACTUALLY HAPPENS

To the outside world, every email your automation sends came from you: it uses your account, your permissions and your reputation, with nobody watching. In security we call this a non-human identity: an “employee” that never sleeps, never quits and is almost never reviewed. When someone leaves the team, their automations usually keep running.

DO THIS TODAY

Make a three-column list: what each automation does, which account it connects with, and who can turn it off. If you can't answer the third, you've found your first risk.

05 Your team's personal ChatGPT is a company archive you don't control.

We usually think: *"My team uses AI to work faster. That's a good thing."*

WHAT ACTUALLY HAPPENS

It is. The problem is where what they type ends up. If someone pastes client data into a personal account, that information lives in a history that isn't yours. When that person stops working with you, the history leaves with them. And depending on the settings, those conversations may be used to train the model.

DO THIS TODAY

Ask your team, without scolding, which AI tools they use and with which account. Move work to business accounts (team or enterprise plans) and turn off the option that lets your conversations train the model.

06 Your partner's voice no longer proves it's your partner.

We usually think: *"If they call and I recognize the voice, it's them."*

WHAT ACTUALLY HAPPENS

With a few seconds of audio taken from a reel, a video or a voice note, a voice can now be cloned convincingly. The typical scam is an urgent call: a payment that can't wait, a new bank account, a relative in trouble. The urgency is there so you don't verify.

DO THIS TODAY

Agree on a code word with your partner, your team and your family, and never write it in a chat. And one rule with no exceptions: no payment or bank account change is approved on a call, voice note or WhatsApp alone.

07 If your chatbot said it, you said it.

We usually think: *"If the AI gets it wrong, it's the tool's mistake."*

WHAT ACTUALLY HAPPENS

In 2024, a Canadian tribunal ordered Air Canada to honor a refund policy its chatbot had made up. The airline argued the chatbot was responsible for its own answers; the tribunal didn't accept that. To your client, what your AI says is your word.

DO THIS TODAY

Write down what your chatbot can't promise without human review: prices, timelines and terms (refunds, warranties, discounts). Then try to get it to promise you something yourself.

WHAT ALL SEVEN HAVE IN COMMON

None of them requires a hacker to pick you. They all start with a reasonable permission nobody looked at again: an app you tried, a link you shared, an automation that worked so well you forgot about it.

AI security isn't about using less AI. It's about knowing, at all times, what each tool can see and what it can do in your name.

Do you know who holds the keys to your business?

Mark yes, no or not sure. Be honest: this test is just for you.

		YES	NO	NOT SURE
1	I know which apps have access to my email and files, and I've reviewed the list in the last 3 months.	☐	☐	☐
2	My AI assistants can't send, delete or pay for anything without my confirmation.	☐	☐	☐
3	I have no client documents shared with "Anyone with the link".	☐	☐	☐
4	I have a list of my automations and know who can turn each one off.	☐	☐	☐
5	My team uses AI with business accounts, not personal ones.	☐	☐	☐
6	I have a code word and a rule for urgent payments.	☐	☐	☐
7	I know what my chatbot can and can't promise a client.	☐	☐	☐
8	If someone on my team left today, I'd know which AI access to shut down.	☐	☐	☐

YOUR RESULT

All "yes"

You're in good shape. Repeat the test every three months: access piles up on its own.

Any "no"

You have concrete work to do, and the guide tells you where to start. One afternoon is enough.

Two or more "not sure"

It's not carelessness. These questions have no obvious answer: access is spread across dozens of screens, and nobody had ever shown them to you together.

TO CLOSE

What this guide can do, and what it can't.

This guide helps you close the doors you can see on your own, in one afternoon.

What it can't do is see your situation: how your tools connect to each other, which accounts have access to what, where your clients' information travels, and what happens if one of those pieces fails. That isn't solved with a checklist. It's solved by looking.

That's what we do at Stratega: we review how a business uses AI and leave it running just as fast, but without open doors.

**If you'd like us to look at it with you,
write to me. I'll reply personally. And if
this guide is all you need, I'll tell you so.**

INSTAGRAM (DIRECT MESSAGE)

[@thediegorojas_](#)

WEB

[stratega.studio](#)

About Diego Rojas

Systems engineer with ten years in software, cloud infrastructure and security. Today he works on non-human identities and AI security: what software, automations and AI agents can see and do. Co-founder of Stratega Studio with Jimena Alzamora.