AI

Safe AI Adoption for **Business Owners**

AI is already inside your business. It shows up in meeting summaries, draft emails, marketing ideas, and quick research.

That is where AI leakage starts. Not with bad intent, but with friction. A task takes too long. A tool is missing. A policy is vague. So people route around the system.

This guide is for owners, COOs, finance leads, and department managers who want progress without chaos.

Use it with your leadership team, your IT partner, and the people doing the work. You will define what tools are approved, what data is off-limits, which workflows are driving workarounds, and what controls make the safe path easier than the risky one.

You do not need to become an AI expert.

You need clarity, ownership, and a plan you can execute in weeks, not quarters.

Shadow AI
is a visibility problem

If you only learn where AI is used after something goes wrong, you do not have an adoption plan. You have a surprise plan.

Shadow AI is what happens when employees use AI outside your approved tools, accounts, and processes. That might be a personal ChatGPT login, a browser extension, or a “free trial” app connected to company data. Once work moves into that shadow, leaders lose basic answers: who is using what, what was shared, where output went, and how it can be audited.

Most businesses reach for the same lever first: ban it. The problem is that bans do not remove the need. They just push the activity into quieter corners.

The better play is to treat AI like email and cloud sharing. You do not stop people from sending messages. You make sure messages are sent through systems you can secure, monitor, and improve.

In the rest of this guide, we build visibility first, then guardrails, then support. The goal is simple: make safe AI use the easiest way to get work done.

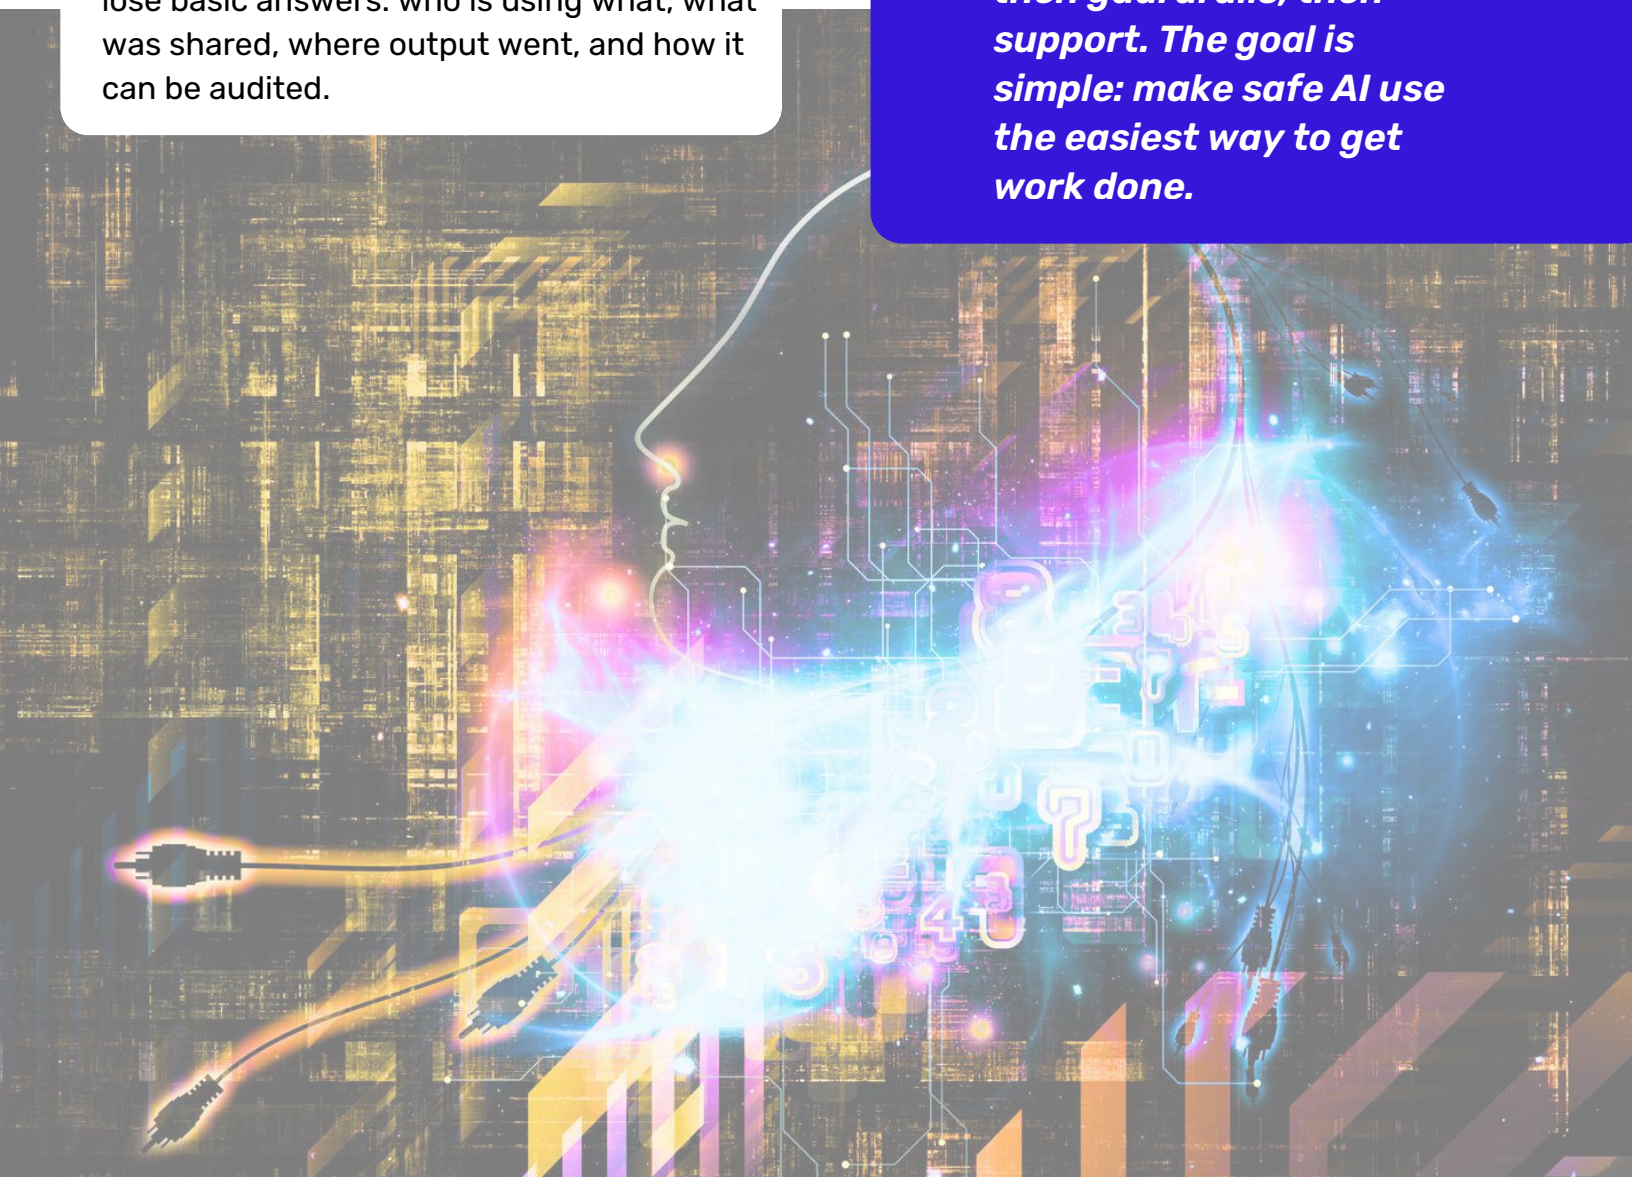

The AI leakage test **(4 questions)**

Set a timer for 15 minutes.
Answer each question with Yes, No, or Unsure.
Treat Unsure as No.

1. **Which AI tools are approved, and which are explicitly not?**
2. **What data is off-limits to paste, upload, or summarize (client lists, contracts, PHI, employee data, financials)?**
3. **Where are employees already using AI today?**
4. **Which workflows are creating the workarounds?**

If you cannot answer these with confidence, start here before you buy anything new.

This is not about catching people. It is about finding friction. Every "No" points to a missing tool, a missing rule, or a missing workflow fix. Leadership owns the fix.

Your goal is not perfection. It is clarity you can act on this month.

Build your AI inventory in **five buckets**

Most companies do not have an “AI program.”

They have a pile of AI touches that grew over time.

Start by listing what is actually in play today. Do it on one page.

Use five buckets:

1. **Tools:** chat assistants, writing tools, meeting note tools, image tools, code helpers, and any browser extensions.
2. **Accounts:** work accounts vs personal accounts, who has admin roles, and how access is removed when someone leaves.
3. **Data paths:** what people paste or upload, where outputs are stored, and whether prompts or files are retained by the tool.
4. **Integrations:** plugins, connectors, and automation tools that can pull data from email, cloud drives, CRM, and finance systems.
5. **Workflows:** the repeatable tasks where AI is used (support replies, proposals, recruiting, claims notes, scheduling, analysis).

This snapshot is not busywork. It is how you spot hidden risk: one unapproved tool in finance, one shared login, one “free” plugin connected to sensitive data.

When you can see the full picture, every decision gets easier: what to approve, what to block, what to train, and what to fix first.

Define what data **never goes into AI**

Most AI risk is not “the model turns evil.” It is an employee copying the wrong thing into the right looking box.

Create a simple, written rule set with examples. Use three bands:

- **Green:** public or already approved to share (public website copy, job postings, generic policies).
- **Yellow:** internal, but low sensitivity (meeting agendas, internal how-to docs, de-identified summaries). Allowed only in approved tools and only when the output stays inside your environment.
- **Red:** sensitive data. Never paste, upload, or summarize it in any AI tool unless your security team has explicitly approved a protected workflow.

Red typically includes:

- Customer and patient information, student records, claims data, donor lists
- Contracts, pricing, M and A details, banking, payroll, tax information
- Credentials, API keys, security configurations, incident response details
- Anything covered by confidentiality, attorney-client privilege, or regulator rules

Give people a fallback move. If they need help with Red data, require one of these:

1. remove identifiers and numbers, then rephrase the question at a process level
2. use an approved internal tool that is authenticated, logged, and covered by your policy
3. open a short request so IT or security can provide a safe workflow

Write the rule in plain language: “If it would be harmful on the front page of the newspaper, it does not belong in a prompt.”

Then make the safe alternative easy: a secure internal knowledge base, a ticketed workflow for high-risk requests, and a fast way to ask, “Is this safe?”

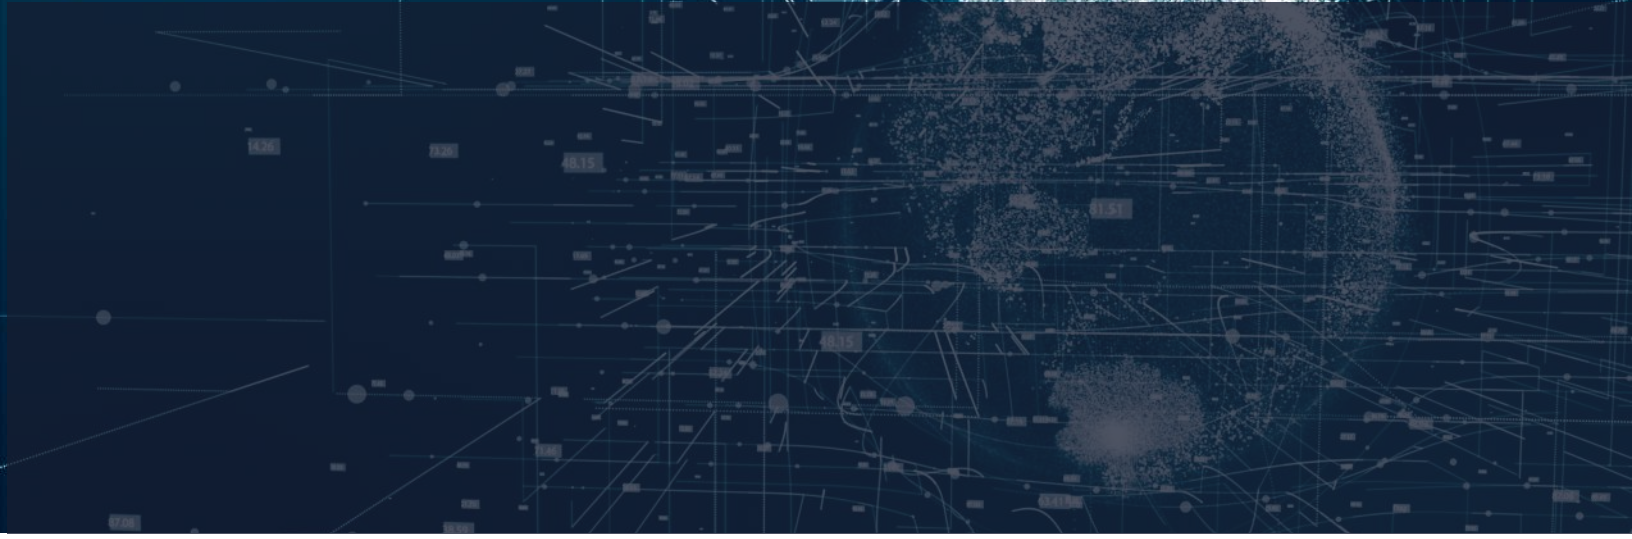

Approve tools, **then make the safe path the default**

Employees use unapproved AI when the approved option is slower, harder, or missing. Your job is not to eliminate curiosity. It is to remove the need for workarounds.

Start with two categories:

1. **Productivity AI inside your core platforms** (for example, AI features in Microsoft 365 or Google Workspace).
2. **Specialized tools for specific teams** (marketing, support, development, analytics), approved through a short review.

For each approved tool, decide:

- Work accounts only, no personal logins for company work
- Single sign-on, MFA, and least-privilege access
- Clear retention and logging settings, aligned to your policy
- No “bring your own plugin” without review
- A documented owner for licenses, renewals, and support

Do basic vendor due diligence. Confirm, in writing, what happens to your data: whether content is used for training, how long it is retained, and how confidentiality commitments are enforced. If the vendor cannot explain this in plain English, keep shopping.

Then close the gaps that drive risky behavior:

- If people paste data because search is terrible, fix search.
- If people copy reports because templates are messy, standardize templates.
- If people connect random apps to move files, build a supported automation path.

One more trap: buying inside an AI chat. If a tool can trigger purchases or subscriptions, extend your procurement rules to cover it. Define who can buy, what can be bought, and where spend is visible.

When the safe option is one click away, adoption stops being a guessing game.

Secure the workflow, **not just the chat window**

If you are building or connecting AI into business systems, treat it like any other software integration, with extra caution around inputs, outputs, and permissions.

Three risk patterns show up again and again:

1. **Prompt injection:** the model is tricked into ignoring instructions or exposing data.
2. **Sensitive data disclosure:** the model sees or produces information it should not.
3. **Excessive agency:** an assistant is allowed to take actions (send email, edit files, place orders) with too much reach.

Practical guardrails:

- Keep AI read-only by default. Add write actions only after testing.
- Limit what the model can access. Use role-based access and least privilege.
- Validate outputs before they trigger downstream actions. Humans review high-impact steps.
- Log prompts, tool calls, and data access in a way your team can audit.
- Treat plugins and connectors as suppliers. Review them, version them, and be able to disable them fast.

A safe AI assistant is not only accurate. It is contained, observable, and recoverable.

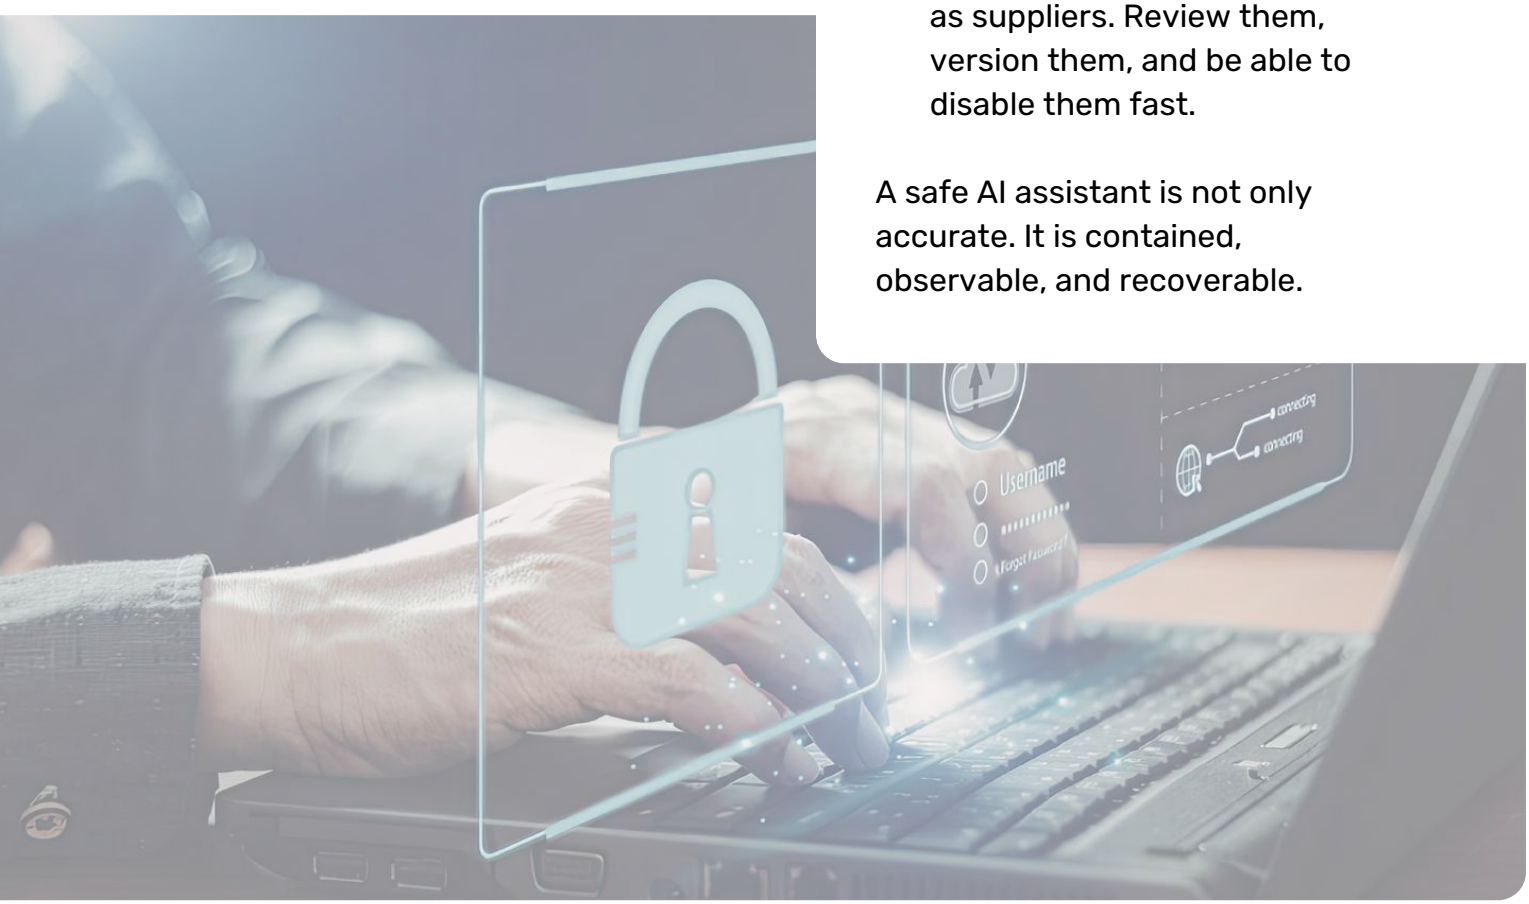

Governance that feels like operations, **not bureaucracy**

Responsible AI governance fails when it is written like a legal memo and ignored like a legal memo. For business owners, governance is simply decision-making with ownership.

Start small with an AI-council (5 to 7 people): an executive sponsor, IT/security, legal or compliance, finance, and two workflow leaders (sales, service, operations). Their job is to set boundaries, approve tools, and remove friction.

Use a simple cadence:

- Monthly 30-minute review: new-tool-requests, incidents and near-misses, top workflows, and what is being retired.
- Quarterly risk check: run the four-question AI-leakage test again, then update your approved list and off-limits examples.

Keep a living approved-tool register. For each tool, capture: business-owner, technical-owner, SSO/MFA, role-based-access-control, data-loss-prevention, retention-window, training-use, logging-level, plugin-policy, procurement-path, support-SLA, and kill-switch-contact.

Borrow the NIST-AI-RMF style of thinking without turning it into homework:

- **Govern:** who owns decisions and exceptions
- **Map:** where AI touches customers, money, credentials, or regulated data
- **Measure:** evals, red-team testing, bias-checks, privacy-checks
- **Manage:** change-control, monitoring, and rollback-plans

Common no-go zones until proven safe: payroll-automation, termination-letters, clinical-recommendations, legal-advice-drafting, wire-instruction-changes, and anything that can edit systems-of-record without review.

Document three things every time:

1. intended-use and non-use (your no-go zones)
2. risk-tolerance (what must be human-reviewed)
3. auditability (logs, retention, and access-controls)

This is how AI becomes a managed capability instead of a side hobby.

Training that matches **real work**

One policy email will not change behavior. Training has to be short, specific, and recurring.

Run a 30-minute kickoff per department, then a 10-minute refresher monthly. Focus on:

- What tools are approved, and where to find them
- What never goes into prompts, with concrete examples
- Where outputs should live (approved folders, tickets, CRM notes), and what should not be emailed
- How to verify output before it leaves the building (facts, tone, numbers, names)
- How to report a “this feels off” moment without getting in trouble

Give people prompt templates that steer them to safer behavior, like:

- “Rewrite this using the same meaning, remove names and identifiers.”
- “Summarize these notes, but do not include any customer data.”

Most mistakes are made by good people moving fast. Training is how you slow the right moments down.

Visibility loops, **not surveillance**

If you cannot see AI usage, you cannot manage AI risk. But visibility does not mean spying on employees. It means logging the right events, in the right systems, for the right reasons.

At minimum, you want answers to:

- ☐ Which AI tools are used, by department
- ☐ Which accounts are used (work vs personal)
- ☐ What data sources are connected (email, drives, CRM)
- ☐ What high-risk actions occurred (file exports, bulk copy, external sharing)

Use the same approach you use for other cloud apps: centralized identity, access reviews, and audit logs. Then add a lightweight discovery step. Once a quarter, ask teams, “What tools are helping, and what is getting in the way?” You will learn more from honest friction reports than from guessing.

When you find an unapproved tool, avoid panic. Ask: what problem was it solving? Can you approve it safely, or replace it quickly?

Visibility turns AI from a rumor into something you can improve.

A 30-day rollout plan **you can actually finish**

Week 1

Run the AI-inventory (tools, accounts, integrations).

Pick an executive sponsor.

Publish a one-page use-AI-safely rule with Green/Yellow/Red examples.

Week 2

Choose the approved-tool set (enterprise-grade accounts only).

Turn on SSO/MFA, and set baseline controls: role-based-access-control.

Data-loss-prevention.

Retention-settings.

eDiscovery-holds.

Audit-logging.

Week 3

Fix the top workflow-friction points that drive copy-and-paste behavior (search, templates, approvals, reporting).

Add a fast exception-path for high-risk requests, plus a procurement-visibility rule for in-chat purchasing.

Week 4

Train each department with real examples.

Ship prompt-templates.

Start monitoring for shadow-tool use via identity and app-discovery.

Deliverables to keep in one folder

approved-tool-register, data-classification-matrix, prompt-library, plugin-approval-checklist, vendor-privacy-and-confidentiality-attestation, model-change-log, systems-of-record-write-restrictions, customer-facing-human-review-requirements, incident-playbook-for-AI, secure-usage-metrics-dashboard, and a quarterly-review-calendar.

Controls to verify by day 30

prompt-injection-testing, insecure-output-handling-sanitization, sensitive-data-exfiltration-prevention, connector-scope-limitations, retrieval-augmented-generation-source-filtering, model-denial-of-service-rate-limits, supply-chain-package-integrity-checks, access-review-schedule, logging-retention-alignment, red-team-tabletop-exercise, model-theft-protection, secret-scanning, token-budget-alerts, kill-switch-contact-list, and emergency-disable-steps.

If you want help, start with the first “No,” assign an owner, and set a due date this wee.

If you're thinking, "We're probably using AI already… but I'm not sure it's safe or visible," you're not alone.

Safe AI adoption is less about picking the perfect tool and more about making sure your team can use AI inside clear guardrails, with the right controls, and without creating workarounds.

We can help you:

- Map where AI is already being used across your teams and tools
- Define approved AI tools and a simple "what data is off-limits" policy
- Lock down accounts, access, plugins, and data paths so AI stays inside your controls
- Identify the workflows creating shortcuts and fix the friction driving shadow AI
- Build a practical training plan that matches real work, not policy jargon
- Set up monitoring and a lightweight governance cadence so you can keep improving

If you want a second set of eyes from a local team serving Chicago and beyond, reach out and we'll start with a short, structured AI visibility and guardrails checkup.

Get in touch.

CALL: (312) 985-6810
EMAIL: info@reintivity.com
WEBSITE: www.reintivity.com

REINTIVITY technology solutions

Serving the Chicagoland