AI ENABLEMENT · KEEPING IT CLEAN

Keeping It Clean

Data habits for AI and automation

Most data accidents look more like kitchen mistakes than break-ins.

A 15-minute team intro

Most data failures aren't break-ins

If you asked most people to picture a data breach, they'd describe a hacker. Someone in a hoodie at a screen, breaking in from outside.

Almost none of the daily failures in real organisations look like that. They're quieter, smaller, and almost always preventable. Someone took a sensible-looking shortcut on a busy day. The official tool was slow. Nobody had a moment to stop and ask.

The thing to watch isn't the front door. It's the small habits.

Food safety is the closest comparison we have

You wash your hands before cooking. You don't use the same chopping board for raw chicken and then salad. You check the date on the milk before pouring it into your coffee.

None of this was on a poster you had to memorise. No one made you sit through a training video. You picked it up because the consequences are real and recognisable. You learned to notice the moments when something could go wrong, and you developed habits that cover most of them without thinking.

Data habits work the same way. You build instincts, not rules.

The information isn't dangerous. The context is.

A customer name is fine on the CRM. A diagnosis is fine in the clinical system. A salary figure is fine in HR's spreadsheet. Each one becomes a problem the moment it ends up somewhere it doesn't belong.

> *Like cooked food on a board that just had raw chicken on it. The food was fine. The contact was the problem.*

Watch the contact, not the content.

Five places things tend to go wrong

Read these with your own work in mind. You will recognise at least two of them.

1. The paste into the chatbot
2. The personal account
3. The unsanctioned workflow
4. The shared document
5. The AI output that contains things it shouldn't

1 The paste into the chatbot

You have a long customer email. You're trying to draft a careful reply. You open a free AI chatbot, paste the email in, and ask for help with the response.

The customer's name, contact details, complaint, and possibly account information have now left your organisation's systems. The reply you eventually send is helpful. The information that got you the reply is sitting on a server you can't find, governed by terms you didn't read, owned by a company that has no contract with yours.

> ***Kitchen equivalent:*** *Tasting raw chicken marinade off your fingers. The intent was fine. The action was the problem.*

Intent fine. Action a problem.

2 The personal account

The work system is slow. The file is too big to email. You want to work on it from your phone at the weekend. So you forward the document to your personal email, your personal cloud, or your personal AI account.

From your organisation's perspective, that document has left the building. The backups, access controls, audit logs, and legal protections that cover the work systems don't follow it. If your personal account is compromised, that document is compromised with it.

> ***Kitchen equivalent:*** *Taking ingredients home from a restaurant kitchen to cook in your own. You may be a clean cook. The restaurant has no way of knowing. The food won't turn out the same way it would if you'd ordered the takeaway instead.*

Once it's outside the work systems, the protections don't follow.

3 The unsanctioned workflow

You found a tool that does something useful. You signed up. You connected it to your work calendar, your inbox, your CRM, or your shared drive so it could do the useful thing.

Nobody told you not to. Nobody told you to. The tool is now ingesting data from your organisation and sending it to a third party that nobody has reviewed and nobody is monitoring.

> ***Kitchen equivalent:*** *Bringing your own gadget from home into a commercial kitchen, plugging it in, and using it on food going out to customers. The gadget may be fine. The introduction of it has not been checked.*

The useful thing is real. The data flow underneath it is invisible.

4 The shared document

You shared a folder with a colleague last year. Three colleagues have left since then. Two of them still have access because nobody removed them. The folder also has "anyone with the link can view" turned on because that was simpler when you needed to share something quickly.

Today the folder contains things it didn't contain last year. The access settings were set for the documents that used to be there.

> ***Kitchen equivalent:*** *A cupboard you set up for one purpose two years ago. The shelf labels are still there. The contents have changed.*

Set once. Never reviewed.

5 The AI output that contains things it shouldn't

You asked an AI to summarise something. The summary it produced included a sentence pulled directly from a confidential source.

Or you asked it to draft a reply that referenced a customer, and it correctly used the customer's name, which now appears in a chat history kept somewhere you didn't choose.

The model was doing its job. The output contained information that shouldn't have travelled with it.

> ***Kitchen equivalent:*** *A sauce that turns out to contain something one of your guests is allergic to. You knew what went in. You didn't think about what was coming out.*

Watch what comes out, not just what goes in.

The kitchen matters, not just the cook

Where you do the work changes the risk profile. Sanctioned enterprise tools are operating under contracts your organisation has negotiated. They don't train on your data by default. They log access. They sit within boundaries your organisation has agreed to.

Free versions of the same tools, or unsanctioned tools you signed up for yourself, don't have any of that protection. The model itself may be identical. The kitchen around it is completely different.

Sanctioned tools are a kitchen built for safe handling. Consumer tools are a borrowed camping stove.

The pause

Good cooks pause for a split second before doing something irreversible. The pause isn't a procedure. It's the habit underneath the habit.

You're aiming for the same thing with data. A small pause before information moves. The next slide is what fills it.

The pause is the habit. The questions are how you fill it.

Three questions, ten seconds

1

What kind of information is this?

Personal data about identifiable people, commercial information under an NDA, internal-only content, public marketing copy. You don't need to memorise the categories. You need to notice when one of the more sensitive ones is involved.

2

Where is it going to end up?

A sanctioned enterprise tool, a personal account, an internal colleague, an external contractor. The same information can be fine in one destination and a problem in another.

3

Would I be comfortable if my organisation could see this happening?

The question that catches what the first two miss. If the answer is no, that's information. Pause for one more second before you act.

Ten seconds, most of the time. Same as checking the date on the milk.

Frameworks exist. You don't have to memorise them.

You'll hear about GDPR, the EU AI Act, the UK Data Protection Act, HIPAA, sectoral rules like FCA or NHS Information Governance, and a dozen others.

Your organisation has a position on which of these apply to you. Your job is to know where to find that position, not to learn the regulations yourself. The same way you don't memorise the food hygiene rating system; you just trust that the kitchen has been inspected.

Find your organisation's policy. Find the person who owns it. That's the homework.

The landscape changes. The habits don't.

The rules around AI and data are moving in real time. Models come and go. Regulations update. Vendors disappear overnight. As this playbook was being written, two of the most capable AI models on the market were pulled by a government directive within a week of launch (Anthropic’s Fable).

The specifics will keep changing. The habits in this playbook won't. Anchor in the habits. Let the specifics float.

Anchor in the habits. Let the specifics float.

Ask the people who know

Your information governance team. Your data protection officer. Your security team. They are not there to gatekeep. They are there to tell you what's allowed, what isn't, and what to do when you're not sure.

If you don't know who they are in your organisation, finding that out is one of the highest-value five minutes you'll spend this month.

A five-minute conversation can save you a six-month investigation.

When something does go wrong

Tell someone quickly. The damage from a small incident reported early is almost always smaller than the damage from a small incident hidden. Organisations have processes for this for a reason.

WHAT TO DO

Tell your line manager. Tell your information governance team. If sensitive data has been involved, tell them today. The point is not to apportion blame. The point is to limit the consequences before they grow.

The worst version of any incident is the one nobody knew about until later.

You don't have to become a privacy officer

The job isn't to learn everything about data protection law. The job is to notice three things: what kind of information you're handling, where it's going, and whether to ask before it gets there.

> *Your permission slip: small habits cover most of the risk. The hard cases are what the team is for. You don't need to know everything. You need to notice the moment something is about to move, and stop for ten seconds.*

Small habits, applied often. That's the whole job.

Recap

1. Most data failures are unforced errors, not break-ins.
2. The information isn't dangerous. The context is.
3. Five moments to watch: paste, personal account, unsanctioned workflow, shared document, AI output.
4. Where you work matters. Sanctioned tools have protection built in. Free tools don't.
5. The pause. Three questions. Ten seconds.
6. Find your organisation's policies. Ask the team when in doubt.
7. If something goes wrong, tell someone quickly.

Where to go next

The full playbook

Long-form reference with examples, the five moments expanded, and a glossary.

The quick reference

One-page summary for desks, induction packs, and team rooms.

Companion playbooks

Start Here. Choosing the Right Tool. Build & Develop next.

Who to ask in your org

Information governance, data protection, security, IT.

Questions?

Discussion, examples from your team, things you're unsure about.