
03
AI &
AUTOMATION
KEEPING IT CLEAN
DATA HABITS
Greg Cartwright

A note before you start

If you've read Start Here and Choosing the Right Tool, you've already done most of the thinking that this playbook builds on. Start Here got you noticing your own work. Choosing the Right Tool gave you a vocabulary for the tools you've been handed. This one is about something quieter and slightly more uncomfortable: what can go wrong when those tools start moving data around, and how to develop the habits that keep you and your colleagues out of trouble.

That word “uncomfortable” matters. Most content about data protection, governance, and security falls into one of two camps. Either it's dry, technical, and aimed at people who already know the acronyms, or it's scaremongering: long lists of catastrophes that make you anxious without telling you what to actually do. Neither helps you on a Wednesday afternoon when you've got an email to answer, and the official tool is being slow.

This playbook tries to do something different. It treats governance the way most of us already treat food safety. You don't think of yourself as a food safety expert. You just know not to use the same chopping board for raw chicken and salad. You learned that habit because the consequences are real and recognisable. The aim here is to build the same kind of instinct for data.

About twenty minutes of reading. A cup of tea. By the end you'll have a small set of habits that cover most of what matters, and a clearer sense of when to stop and ask.

> The one idea in this playbook: most data accidents are unforced errors. Good habits prevent the ones you can prevent. The team handles the rest.

Section 1: 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. That image is what gets the headlines. It's seldom what actually causes the daily failures inside real organisations.

The actual incidents are smaller, quieter, and easier to recognise once you know their shape. They happen because someone took a sensible-looking shortcut on a busy day. They happen because the official tool was slow and the unofficial one was right there. They happen because no one had a spare moment to stop and ask.

In the regulatory data published every year by national data protection authorities, the vast majority of reported incidents fall into a small handful of categories. Information was sent to the wrong recipient. Documents shared with the wrong access settings. Devices left where they shouldn't be. Mistakes made by the people who handle the data every day, doing their jobs, intending no harm.

This is the same shape as food safety. The dramatic cases are food poisoning outbreaks that make the news. The vast bulk of cases are domestic, individual, and entirely preventable. You don't prevent them by reading a regulation. You prevent them by developing instincts that cover the moments where things tend to go wrong.

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

Section 2: Food safety is the closest comparison we have

Almost everyone is surprisingly good at food safety. You wash your hands before cooking. You don't use the same chopping board for raw chicken and then prepare a salad on it. You check the date on the milk before pouring it into your coffee. You throw out the leftovers that have been in the fridge for a fortnight, even when they smell fine.

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 of getting it wrong are real and recognisable. Food poisoning is unpleasant. Giving someone food poisoning at a dinner party is worse. You learned to notice the moments when something could go wrong, and you developed habits that cover most of them without thinking about it.

Data is the same. Almost every governance failure in an organisation looks more like a kitchen mistake than a break-in. Someone uses the wrong board for the wrong job. Someone forgot to check the date. Someone leaves the fridge open. The information ends up somewhere it shouldn't be, and by the time anyone notices, it's already happened.

Separation

In a kitchen, you keep raw and ready-to-eat foods apart. Different boards, different sides of the fridge, different utensils. The food itself isn't dangerous in isolation. The danger is in the contact.

Data works the same way. Most of the data you handle every day is fine in the right context. A customer's 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 when it ends up somewhere it doesn't belong.

The mistake to watch for is the one that feels harmless. Pasting a customer's complaint into a free chatbot to help draft a reply. Forwarding an HR document to your personal email because the work account is slow. Sharing a clinical note in a team chat because the right system is too clunky. None of

these feels like a breach in the moment. All of them are the data equivalent of putting cooked chicken back on the board that just had raw chicken on it.

Freshness

Some food is fine today and dangerous next week. The thing in the fridge hasn't changed; the context around it has.

Data has the same property. Information that was fine to hold last year may not be fine to hold now. Information you collected for one purpose may not be appropriate to use for another. A customer's email address given to you for delivery updates is not the same email address given to you for marketing, even though it's the same string of characters in the same database.

Most organisations have rules about how long they hold different types of information and what they're allowed to use it for. You don't need to know those rules in detail. You need to know they exist, and that “we already have it” is not the same as “we can do whatever we like with it.” When in doubt, ask. The same way you'd ask a colleague whether the leftovers in the office fridge are anyone's before throwing them out.

> Data habits work the same way as food safety. You build instincts, not rules.

Section 3: Five places things tend to go wrong

Read these with your own work in mind. You will recognise at least two of them. The point isn't to make you feel bad about doing them in the past. The point is to start noticing them so you can pause before doing them again.

The paste into the chatbot

You have a long email from a customer. You are trying to draft a careful reply. You open a free AI chatbot, paste the email in, and ask it to help you with the response. The customer's name, contact details, complaint, and possibly their account information have now left your organisation's systems and gone somewhere you do not control.

The reply you eventually sent was helpful. The information you sent to get it is now sitting on a server you cannot find, governed by terms you did not 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.

The personal account

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

From your organisation's perspective, that document has now left the building. It is no longer covered by the backups, the access controls, the audit logs, or the legal protections that apply inside the work systems. If your personal account is ever compromised, that document is compromised with it. If you leave the organisation, the document goes with you, whether anyone intends it to or not.

> **Kitchen equivalent:** Taking fresh ingredients home from a restaurant kitchen to finish preparing them in your own kitchen. You may be a perfectly clean cook. The restaurant has no way of knowing. The tools and the staff in the kitchen are fit for purpose; yours are not professional grade.

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. The useful thing is genuinely useful. The data flow underneath it is invisible to everyone who would normally check.

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

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 when they left. The folder also has “anyone with the link can view” turned on because that was simpler when you needed to share something quickly, and you forgot to change it back.

Today the folder contains things it didn't contain last year. The access settings were set for the document that was in there before. The setting hasn't changed; the contents have.

> **Kitchen equivalent:** A cupboard you set up for one purpose two years ago and never reviewed. The shelf labels are still there. The contents have changed.

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 that is being kept somewhere you didn't choose.

The model was doing its job. The output contains information that should not have travelled with it.

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

> If you recognised at least two of these, you're already paying the right kind of attention.

Section 4: The kitchen matters, not just the cook

All five of those moments are about individual choices. Pause before you paste. Don't forward to your personal account. Review what you've shared. They put the responsibility on you, the person at the keyboard, to spot the moment and act on it.

That's only half the picture. The other half is the environment around the action. Where you're working changes the risk profile, often quite dramatically, before you've made any choice at all.

Sanctioned tools and what they actually offer

Sanctioned enterprise tools, the kind your organisation has licensed and approved, are typically operating under contracts your organisation has actually signed. They typically don't train on your data by default. They typically log access in ways your organisation can audit. They typically sit within geographic and legal boundaries that your organisation has chosen.

Free versions of the same tools, or unsanctioned tools you signed up for yourself, don't have any of that protection. The user interface might look identical. The data posture is completely different.

This is the difference between cooking in a kitchen that's been inspected and certified, and cooking on a borrowed camping stove. Both can produce a meal. The level of background protection is nowhere near.

A familiar logo isn't the same thing as a sanctioned tool

This is the part that catches people out most. A well-known brand on the door is not a guarantee that the kitchen behind it is set up for your data. The same major vendor often runs both a consumer service and an enterprise service, and the difference between them is not visible from the user's side.

Free ChatGPT, free Gemini, free Claude, free anything else, all operate under consumer terms. Your data may be used for training. Your conversations may be reviewed. There is no contract between that vendor and your organisation. The enterprise versions of the same products are governed by terms your organisation has negotiated, with data protection commitments your legal team has reviewed.

Knowing which one you're using is part of the habit. Recognising a brand isn't the same thing as having read the small print.

The infrastructure can move under your feet

Even sanctioned tools sit on infrastructure that can be affected by forces well outside your organisation's control. As this playbook was being written, the US government issued an export control directive that forced a major AI provider to disable two of its newest models for every customer worldwide, within hours. The provider was serious. The models were well-built. None of that was the question. A government decision in one country pulled a tool that customers across the world were depending on.

The lesson is not that AI is dangerous or that providers are untrustworthy. It is possible that any single tool can become unavailable for reasons that have nothing to do with how good it is. Knowing which workflows would survive a vendor going dark is not a hypothetical question. Sometimes vendors do go dark.

> The logo on the door isn't the answer to “is this safe for what I'm about to do?”

Section 5: The pause and the three questions

The food safety equivalent of everything in the last two sections is the pause. The split second where a good cook stops, looks at what they're about to do, and asks themselves whether anything in the next action is going to cause a problem. Most experienced cooks don't experience this as a conscious thought. It is the habit underneath the habit.

You are aiming for the same thing with data. A small pause before information moves. Three questions, in order, ten seconds most of the time.

1. What kind of information is this?

You don't need legal definitions. You need to recognise broad categories so the pause has something to land on. Most organisations group information into something like the following:

- Personal data: anything that can identify a real person. Names, contact details, account numbers, and photos.
- Special category data: a stricter subset that gets extra protection in most regimes. Health, ethnicity, sexuality, political views, biometrics, religion.
- Commercial and contractual: anything an NDA or contract covers. Pricing, strategy documents, client lists, supplier terms.
- Confidential internal: business-sensitive material that isn't strictly commercial. Personnel matters, internal plans, drafts.
- Routine internal: the everyday operational stuff. Not secret, but not for public consumption either.
- Public: marketing material, published reports, anything already out there.

Most of what you handle on any given day is routine, internal or public. The serious failures cluster around the first three. The point of the first question is to notice when one of those is involved, not to memorise the definitions.

2. Where is it going to end up?

Pasting into a sanctioned enterprise tool is one answer. Pasting into a personal account is another. Sending to a colleague inside the organisation is one. Sending to a contractor outside it is another. The same piece of information can be perfectly fine going to one destination and a serious problem going to another.

The destination is at least as important as the content. A customer's contact details going into the CRM is part of doing your job. The same details going into a free chatbot is a different action with the same data.

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

This is the question that catches what the first two miss. If you'd be uncomfortable explaining the action to your line manager, to your data protection officer, or to the person whose data you're handling, that is information. It doesn't necessarily mean what you're about to do is wrong. It means it's worth pausing for one more second before you do it.

Most of the time, the answer is fine, and you carry on. Some of the time, the answer makes you reconsider. That moment of reconsidering is the whole point.

> Three questions, ten seconds, most of the time. Same as checking the date on the milk.

Section 6: Frameworks exist. You don't have to memorise them.

Underneath everything in this playbook is a layer of actual regulation. The big frameworks you'll hear about, with brief notes on what each covers, mostly so the names ring a bell when they come up:

- GDPR (General Data Protection Regulation): the EU's data protection law, which the UK kept a version of after Brexit (UK GDPR plus the Data Protection Act 2018). Covers how organisations handle personal data.
- EU AI Act: the EU's regulation specifically about AI systems, with tiered obligations depending on the risk category of the system.
- HIPAA: a US framework specifically covering health information.
- CCPA / CPRA: California's consumer privacy laws.
- Sectoral rules: financial services have their own (FCA in the UK, equivalent regulators elsewhere). Healthcare has its own (NHS Information Governance in the UK). Defence, education, and others have their own.

Your organisation will have a position on which of these apply to you, and what they require you to do. Your job is to know where to find that position, not to learn the regulations themselves. The same way you don't memorise the food hygiene rating system, you trust that the kitchen has been inspected, and you know who to ask if you need to check.

Find your organisation's policy

Almost every organisation has documented policies on data handling, AI use, and acceptable tools. Almost nobody reads them. That's fine for most days. What matters is knowing they exist and where to find them, so that on the rare day you need to check something, you can.

If you don't know where your organisation's policies live, finding out is one of the highest-value five minutes you can spend this month. The person who owns those policies is almost certainly someone you can ask directly.

> 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. A model that was publicly available last week may not be available today. A platform that was approved last year may not be approved this year.

The specifics will keep changing. The habits in this playbook won't. Information that is sensitive will stay sensitive. Destinations that aren't safe will stay unsafe. The pause and the three questions will still work in five years' time, regardless of which tools are in the news.

Anchor in the habits. Let the specifics float.

Section 7: Ask the people who know

Your information governance team. Your data protection officer. Your security team. Most organisations have at least one of these, sometimes all three. 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.

Many people avoid asking because they're worried about looking like they don't know, or because they assume the answer will be no. Both of those concerns are usually misplaced. The data protection community as a profession is mostly populated by people who want to help, and they're far more frustrated by people who don't ask than by people who do.

What to bring to that conversation

- The task. What you're trying to do, in plain English, without specifying a tool.
- The kind of data involved. Personal data? Commercial information? Routine internal stuff?
- Where the data would need to go. Inside the organisation? To an external service? To a personal device?
- Any constraints you already know about. “This involves NHS data,” “This touches contracts under NDA,” “This is personal data about staff.”

That's enough for someone who knows the landscape to either give you a clear answer, point you at the right tool, or honestly say “I don't know, but I'll find out.” Any of those answers is worth more than guessing.

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

Section 8: When something does go wrong

This is the section most governance content skips, or buries at the end in language nobody can follow. It's the most important section in the playbook.

Things will go wrong. Sometimes through your own slip. Sometimes through a colleague's. Sometimes a system does something unexpected. The question is not whether incidents happen. The question is what happens next, and that part is much more in your control than it feels in the moment.

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 handling incidents for a reason: the early notification gives them options. The late notification often leaves them with none.

The instinct in the moment is usually to wait, to see if anyone notices, to hope it sorts itself out. That instinct is wrong. Even when nothing comes of an incident, the act of reporting it promptly is the right thing professionally and almost always the right thing personally.

WHAT TO DO

Tell your line manager. Tell your information governance team or data protection officer. If sensitive data has been involved, tell them today. Most regulations require organisations to act within fixed time windows after a breach is discovered; the clock starts ticking when someone reports it, not when it happened.

The point is not to blame

Reporting an incident isn't a confession. It's a notification. Mature organisations treat it as a normal part of how they handle data, not as an admission of failure. The point of the process is to limit the consequences before they grow, not to apportion blame for the original mistake.

If your organisation has a culture where reporting feels dangerous, that is its problem, not yours. The professional standard, across every sector that handles sensitive information, is that early reporting is the right behaviour.

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

Section 9: You don't have to become a privacy officer

Here's the permission slip, as with each of the playbooks before it. 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.

The habits in this playbook are smaller than the subject sounds. The pause. The three questions. The instinct to ask when you're not sure. The willingness to tell someone quickly when something goes wrong. That's most of the job.

The hard cases, the regulatory specifics, the edge cases with sensitive data classes, are what the team is for. You're not the front line of compliance. You're the part of the system that notices the moment something is about to move, and stops for ten seconds.

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

Glossary

A short reference for terms used in this playbook and likely to come up in conversations with your information governance, data protection, or security teams.

Data Protection Officer (DPO)

The person in an organisation responsible for overseeing data protection. Required by law in some sectors and organisations of certain sizes. The right person to ask when you have a question that touches personal data.

GDPR / UK GDPR

The EU's General Data Protection Regulation, and the UK's version of it post-Brexit. Sets the rules for how organisations handle personal data. Wide-ranging, with material penalties for serious breaches.

Incident

Any event where data may have been accessed, lost, shared, or used in a way it shouldn't have been. Includes near-misses. The act of reporting an incident is not the same as confirming one occurred.

Information Governance (IG)

The broader discipline of managing data well, particularly in regulated sectors. Often the team you'd talk to in the NHS, local government, or financial services about whether a particular use of data is allowed.

Lawful basis

Under GDPR, you can only process personal data if you have one of six legal grounds for doing so. The most common are consent, contract, legal obligation, and legitimate interests. Your organisation has decided which lawful basis applies to which kinds of processing; you don't need to choose.

Personal data

Any information that relates to an identified or identifiable person. Includes obvious things like names and contact details, less obvious things like IP addresses and device identifiers, and combinations that together identify someone even if no single piece does.

Special category data

A stricter subset of personal data. Health, ethnicity, sexuality, political views, religious beliefs, biometric data, trade union membership. Treated more carefully under most regimes.

Sanctioned tool

A tool that your organisation has approved for use, typically under a contract that includes data protection provisions. Distinguished from a tool you signed up for yourself, even if the underlying product is the same.

Subject Access Request (SAR)

A request from an individual asking your organisation to show them what data it holds about them. Time-limited, legally enforceable. If you receive one, pass it to your DPO or IG team immediately.

Where to go next

This playbook gave you the five places things tend to go wrong, the pause that prevents most of them, and the answer to the question of what to do when something gets past you. The next step isn't to memorise it. It's to notice the moments in your own work where the pause would apply.

For one week, just notice. You don't need to change your behaviour yet. Just spot the moments where information is about to move, and ask yourself the three questions. You'll catch one or two things by the end of the week. That's enough to start with.

More playbooks are on the way. Build & Develop will go deeper into the practical side: actually building workflows and prompts that hold up over time. Project Delivery will look at what changes when adoption stops being one-person-trying-something and becomes a deliberate piece of work across a team.

None of them is a prerequisite for the others. Read whichever one fits the question you're sitting with.

If you have questions, suggestions, or examples from your own work that might help someone else, get in touch. The playbooks are unbranded by design. Drop your logo on them, share them inside your organisation, send them to a colleague who'd find them useful. The more they're used, the more useful they get.

Look after your data. Look after each other.