

From pilot to enterprise

Why AI succeeds in one department and fails across the enterprise, and how to change that



From Pilot to Enterprise

**Why AI succeeds in one department
and fails across the enterprise, and how to change that**

A practical white paper for chief executives, executive committees and leadership teams

From Pilot to Enterprise

Why AI succeeds in one department and fails across the enterprise, and how to change that

A U-CIE white paper · First edition, October 2026

Legal position and sources checked on 25 October 2026

© 2026 U-CIE.COM. All rights reserved.

This white paper is a work protected by copyright under the French Intellectual Property Code and international copyright conventions. Its text, structure, diagrams, tables and tools (in particular the three-flow map, the four levels of autonomy and the autonomy licence) are the property of U-CIE.

What you may do. You may read this document, share it free of charge and in full, and print it for internal use within your organisation. You may reproduce and complete the tools in the toolkit for your own internal work. You may quote short extracts, provided you clearly cite the source: 'U-CIE, From Pilot to Enterprise, 2026'.

What requires our written permission. Any other reproduction, performance, adaptation or translation, in whole or in part, by any means, requires the prior written permission of U-CIE.COM. This applies in particular to reselling the document, using it in paid training or consulting services, reusing the method or tools under another name, and incorporating them into software. Any unauthorised reproduction constitutes copyright infringement under Articles L. 335-2 et seq. of the French Intellectual Property Code.

Trade marks. U-CIE is a trade mark of UNNAMED COMPANY SAS.

Third-party sources. The works, reports and texts cited belong to their respective authors and publishers. They are referred to with their source, and citing them does not imply that these organisations endorse this document.

Fictional cases. The companies, people and situations described in this white paper are fictional. Any resemblance to real organisations or people is coincidental. They do not describe U-CIE clients or engagements.

Disclaimer. This document offers a methodological and governance framework. It does not constitute legal advice or a recommendation tailored to any particular situation. Regulatory information reflects the legal position on 25 October 2026 and may change. Before making any decision, have your uses assessed by your advisers.

Credits. Cover illustration: U-CIE.COM. Design and writing: U-CIE.

Publisher. UNNAMED COMPANY SAS, 12 rue Vivienne - 75002 Paris. Publication director: Philippe Buschini

Contents

How to use this guide.....	5
The essentials on one page	6
Prologue: six months in the life of a pilot	7
Chapter 1. The misunderstanding: a pilot is not a scale model of the organisation	8
What protects a pilot, and disappears at scale	8
Two mirror-image mistakes	9
Chapter 2. A new lens: seeing the organisation as flows and links.....	10
The flow, the executive's unit of measure	10
The six links	10
Why the word 'integration' is misleading	11
Chapter 3. Cascade effects: how a small error becomes a big problem	12
The right question is not 'can we trust AI?'	13
Chapter 4. Agents: when AI moves from adviser to doer	14
What is an agent, in practice?	14
Four points to watch	14
The autonomy dial: four levels.....	15
Chapter 5. What scaling up really changes: seven workstreams	17
Workstream 1: design the flow before choosing the tool	17
Workstream 2: treat data as a shared product.....	17
Workstream 3: separate the ability to understand from the ability to act.....	17
Workstream 4: name who decides what	17
Workstream 5: evaluate the system, not just the model.....	18
Workstream 6: train for the flow, not for the prompt.....	18
Workstream 7: manage a portfolio, not a collection of prototypes	19
Chapter 6. The European framework, explained simply	20
The principle: the use is judged, not the technology	20
The timetable, checked on 25 October 2026.....	20
AI literacy: an obligation of means, already in force.....	21
The GDPR does not go away	21
When the deployer becomes a provider.....	22
A pragmatic method	22

Chapter 7. Five fictional stories to see the mechanism at work	23
Hélianthe Exploitation: the maintenance that seemed obvious.....	23
Signal Urbain: the report that lost its source	24
Club Partenaires Europe: the sales agent with too many connections.....	24
NordVale Finance: the line between assistance and decision	25
Mécaforge Europe: the recommendation that became a production order	25
Chapter 8. Reading a scale-up budget: what you are really paying for.....	26
Epilogue: back to the original question.....	28
Toolkit	29
Tool 1. Ten questions to ask before approving a scale-up.....	29
Tool 2. The three-flow map.....	29
Tool 3. The autonomy licence	31
Tool 4. Three families of value indicators	32
Tool 5. The leadership team workshop (2 hours).....	32
Tool 6. Five signs that a project is not ready	32
Tool 7. Glossary for the leadership team	34
About U-CIE	36
References.....	37

How to use this guide

This document reads like a story and works like a toolkit. It was written for a leadership team that wants to understand why its AI projects stall just as they start to grow, and then decide what to do about it. No technical background is needed: every concept is explained the first time it appears, and a glossary sits in the toolkit.

The common thread follows an AI pilot in the customer service department of a fictional company, from its success in March to its difficulties in October. Each chapter uses this story to shed light on one mechanism, then ends with two boxes: **Key takeaway** captures the core idea, and **For your leadership team** suggests a question to ask on Monday morning and the tool that helps answer it.

Three ways to use it, depending on the time you have:

Time available	What to read	What you will take away
5 minutes	'The essentials on one page'	The six ideas that change how you look at an AI project
20 minutes	The prologue and chapters 1 to 5	A framework for challenging any project presented to your leadership team
A 2-hour workshop	The fictional cases (chapter 7), then the toolkit at the end of the document	A first map of three critical flows in your organisation

This white paper is one of a pair published by U-CIE, alongside *The Autonomy Dial*. This one explains why an AI project holds up or breaks down when it is scaled. The other sets out how to decide, action by action, what AI is allowed to do on its own.

A word on the status of the information. The legal rules cited were checked on 25 October 2026. The four levels of autonomy and the autonomy licence are a governance method proposed by U-CIE: they do not appear in any legislation. The document always distinguishes what is mandatory from what is recommended. The prologue story, the five sector cases and every company described are fictional; they do not describe U-CIE clients or projects.

The essentials on one page

If you remember only one sentence: **AI almost never stumbles on its intelligence; it stumbles on the seams of the organisation.** Those seams are data that does not mean the same thing from one department to the next, poorly defined access rights, and decisions that nobody knows who owns.

1. A successful pilot proves usefulness, not capability. It shows that a protected team, with hand-sorted data and an expert within earshot, gets value from a tool. It says nothing yet about what will happen with ten teams, five software systems and data that changes every day.

2. Your organisation is not a sum of departments; it is a network of flows. An order becomes a delivery, then an invoice, then a payment. That is the level, the complete journey, at which you need to look at what AI will change.

3. Between two steps in a flow, there is always a link. A piece of data that passes along, a right that applies, a decision that gets made, a responsibility that is exercised. Plugging in an AI system technically resolves none of these links.

4. AI speeds everything up, including errors. Where an experienced employee would instinctively stop a dubious piece of information, a fast system can carry it within seconds all the way to an invoice or a customer.

5. Agents change the nature of the risk. An AI that suggests and an AI that acts do not call for the same level of control. Autonomy is set action by action, on four levels (suggest, prepare, execute, orchestrate), at the pace at which you can monitor it and stop it.

6. Scaling up means funding links, not just tools. Reliable data, tightly bounded rights, named owners, indicators set before launch. It is less spectacular than a demonstration, and it is what makes the difference.

The question that sums up the whole guide: which flows do we want AI to make more capable, and which links must we strengthen so that this capability remains trustworthy at scale?

Prologue: six months in the life of a pilot

An imaginary story, but one that many leadership teams will recognise. We will come back to it throughout the guide.

March. In the customer service department of a services group, a small team is testing an AI assistant. An adviser types a question; the tool finds the contract, summarises the customer's history and proposes a reply. At the demonstration meeting, the effect is immediate: what used to take twenty minutes now takes five. The head of department wants to go further.

June. The leadership team approves the roll-out to three other business units. Licences are bought, training sessions are organised. Enthusiasm is everywhere.

September. The first weak signals appear. In one subsidiary, 'active customer' does not mean the same thing as at head office, and the assistant mixes the two up. Elsewhere, an adviser discovers he can see information he never had access to before. A beautifully written reply quotes a clause that has not existed since the last renegotiation. Finally, the management software is upgraded and a connection stops working, without anyone noticing for ten days.

October. At the leadership meeting, the question lands: 'Is the project working, yes or no?' Nobody can really answer. The users of the original pilot are still delighted. The new business units, much less so.

What happened? The AI model did not get worse between March and October. **The environment changed.** At the start, a handful of people quietly compensated for every imperfection: they sorted the documents, knew the exceptions, caught the errors. At scale, those shock absorbers disappeared, and the dependencies the pilot had hidden became visible.

This is the paradox at the heart of this guide. AI often succeeds **in one department at a time**, because a small scope can be protected. It fails **across the enterprise** when that protection disappears.

To understand why, we first need to clear up a misunderstanding about what a pilot really proves.

Chapter 1. The misunderstanding: a pilot is not a scale model of the organisation

A chef who produces a superb dish for six guests has not yet proved that he can serve three hundred covers every evening. For that, he needs a kitchen brigade, regular suppliers, written recipes, hygiene standards and someone who decides when a plate goes back to the kitchen. The talent is the same. The system around it is entirely different.

An AI pilot is the dinner for six. It answers a real question, and a good one: **under chosen conditions, does AI help a team do a task better?** But the leadership team has to settle a different question: **can the organisation run this use repeatably, safely and measurably, with several teams, several systems, live data and identified owners?** These are not two steps along the same path; they are two different questions, which call for two different kinds of evidence.

What protects a pilot, and disappears at scale

Let us go back to the assistant in the prologue. In March, everything was working in its favour, often without anyone realising:

In the pilot	Across the enterprise
A use case chosen because it is easy to understand	Varied cases, some of them highly unusual
Documents sorted and cleaned by hand	Data that has to be collected, documented and maintained over time
A subject-matter expert available at every step	Users who do not always know whom to call
Volunteers who are already convinced	Diverse groups, some of them sceptical
Rare exceptions, caught by the project team	Daily exceptions that nobody has ever written down
Errors with little consequence	Errors that affect a payment, a customer, an employee

These protections are not a mistake: they make it possible to learn quickly and at low risk. The problem arises when people forget they were there. The OECD notes that organisations adopting AI often have to set up systematic collection of data from their processes, and that the quality, structure and integration of heterogeneous data remain recurring difficulties [1].

Two mirror-image mistakes

The first mistake is to read a successful pilot as implicit permission to roll it out everywhere. The second is to conclude from difficulties in scaling up that the pilot was worthless. In both cases, two levels of decision are being confused.

A pilot can be excellent and its scaling-up laborious, because scaling up requires investment in processes, data, skills and governance. Economists who study so-called 'general purpose technologies', those that transform the whole economy as electricity once did, describe this gap: the gains often appear only once the organisation has been reinvented around the technology [4].

That is why the risk management framework of NIST, the US standards and technology institute, asks organisations to document from the outset the system's purpose, its users, its contexts of use, its limits and the skills needed to operate it [2] [3]. This is not paperwork: it is what turns a demonstration into a system that the organisation understands and can account for.

Key takeaway. A pilot proves that a use is valuable under protected conditions. It does not prove that the organisation can run it without those protections.

For your leadership team. *Question to ask:* what exactly has this pilot proved, and what has it not proved yet? *Tool:* the ten questions to ask before approving a scale-up (toolkit, tool 1). Every empty box is a workstream to fund.

If the pilot is not enough to judge, we need another way of looking at the organisation.

Chapter 2. A new lens: seeing the organisation as flows and links

In a 4 × 100 metres relay, a team rarely loses because one runner is slow. It loses at the baton changeover. An organisation works the same way: its performance depends as much on the handovers between steps as on each step itself. And it is precisely there, in the handovers, that AI slots in.

The flow, the executive's unit of measure

In this guide, a **flow** is the complete sequence that turns a trigger into a business outcome. Not the written procedure: what actually happens, with files sent by email, informal checks, re-keying, phone calls and the shortcuts teams have invented to keep work moving.

Take a customer complaint, like the ones the prologue's assistant handled. The flow begins when the message arrives and ends when the customer is satisfied. In between, someone identifies the customer, retrieves the contract, assesses the urgency, chooses who should deal with it, proposes a solution, records the action, then closes or reopens the case. AI can step in at almost every stage: sorting the message, finding a clause, summarising the history, drafting a reply, preparing an update in the CRM (the software that centralises the customer relationship).

At each of these points, it depends on what we call **links**: whatever allows one step to receive the information, the right or the decision that the next step needs. They fall into six families.

The six links

Link	In plain terms	The question that reveals it
Data	Is the information accurate, up to date, and does it mean the same thing everywhere?	'What exactly do we mean by an active customer, and which system is the reference?'
Applications	How AI connects to your software (CRM, accounting, email, scheduling)	'If this connection fails, who notices, and how quickly?'
Identity and rights	On whose behalf AI acts, and what it can access	'Can it see or do something the user is not allowed to?'
Decision	What AI does with its output: suggest, prepare, execute or orchestrate (chapter 4)	'Who decides, and by what rule, when a case is out of the ordinary?'
Accountability	Who detects, corrects,	'Who has the right to press the

Link	In plain terms	The question that reveals it
	suspends, informs	stop button?’
Measurement	How we know the flow is genuinely improving	‘What was our starting figure, and what do we look at each month?’

Reread the September episode of the prologue with this table in mind: each incident corresponds to a link. The definition of ‘active customer’ is a **data** issue. The adviser who sees too much information, a **rights** issue. The connection broken for ten days, an **applications** and **accountability** issue: nobody was responsible for monitoring it. And the inability to say in October whether the project works is a **measurement** issue.

A few pointers from the reference literature take this further. On data, connecting two systems is not enough if they do not speak the same business language: a study by the European Commission’s Joint Research Centre shows that this layer of ‘meaning’, known as the semantic layer, is decisive in most of the cases studied [5]. On rights, security specialists recommend that AI always act in the context of the user concerned, with the minimum access necessary, and that the software systems themselves, not the AI model, check permissions [6] [7]. On accountability, NIST organises risk management around four functions (govern, map, measure, manage) and expects clear roles, an inventory of systems and explicit commitment from senior leadership [2]. On measurement, it recommends continuously comparing results observed in real conditions with those obtained during testing [8].

Why the word ‘integration’ is misleading

People readily talk about ‘integrating’ AI into the information system. The word suggests plugging in a cable. But scaling up also requires integrating meaning, authority, control and accountability. An AI system can be technically connected to ten applications and still be unusable in a critical flow if nobody knows which data prevails, who approves, how exceptions are handled, or when to stop everything.

Key takeaway. Look at AI projects at the level of the complete flow, and examine each of the six links. A project that is technically well connected can remain fragile on all the others.

For your leadership team. *Question to ask:* for this project, can you show me the six links and tell me which ones are still marked ‘unknown’? *Tool:* the three-flow map (toolkit, tool 2).

A fragile link is not an isolated problem: it becomes the starting point of a chain reaction.

Chapter 3. Cascade effects: how a small error becomes a big problem

In a manual flow, your employees play a role nobody sees: they act as circuit breakers. They spot the odd figure in a spreadsheet, recognise the sensitive customer, know that a status shown in the system does not reflect reality on the ground. This expertise is valuable, but it is rarely written down anywhere.

AI puts these circuit breakers to the test, because it moves information from one step to the next in seconds: it summarises a note and feeds a ticket, reads a document and prepares an order in another system. Well framed, this speed creates a great deal of value. Without a stopping point, it carries errors further and faster. This is called a **cascade effect**: the output of one step becomes the input, decision or action of the next, and an initial error spreads, growing as it goes.

These cascades take five forms, matching the links in the previous chapter.

The data cascade. A customer reference is wrong in a source file. The AI picks it up, rephrases it neatly without flagging any doubt, and another system then attaches the request to the wrong account. The AI ‘invented’ nothing: it faithfully passed on false data. But because the tool is seen as reliable, the error gains credibility at every step.

The decision cascade. The AI classifies a request as low priority because certain expected words are missing. The manager, trusting it, assigns the resource elsewhere. The delay grows; the customer escalates. No component has ‘failed’: it is the chain of a recommendation, excessive trust and an allocation decision that does the damage.

The rights cascade. An assistant has been given the right to search a document base, write to the CRM and send emails, all under a single identity with very broad access. One ambiguous or malicious instruction is then enough to trigger an action that goes far beyond the need. The best-known case is called **prompt injection**. Imagine an executive assistant who carried out every instruction found in the post, including a sticky note slipped into an envelope by a stranger. That is what can happen when an AI system reads a document or email containing hidden instructions: it may follow them as if they came from you. NIST explains that in these systems, the boundary between ‘data to read’ and ‘orders to carry out’ becomes unreliable [9]. OWASP, a global reference in application security, lists among the possible consequences the leaking of information and the triggering of unauthorised actions in connected systems [6].

The confidence cascade. A reply is well written, well structured and convincing. The user lowers their guard. The text is copied into a note, the note into a report, and the report feeds a decision. NIST uses the term **confabulation** for the confident production of false content, and points out that even the sources and explanations cited by AI may be wrong [10]. The real danger is not the wrong answer: it is the wrong answer that is credible enough to pass several checks. The out-of-date clause quoted by the prologue’s assistant is a typical example.

The accountability cascade. The business chose the tool, IT configured the connection, the vendor hosts the model, security manages access, legal reviewed the contract, operations use the output. Each sees one piece. Nobody owns the flow end to end, or the right to stop it. The project looks well governed because it has many stakeholders; in reality it is fragile because it has no owner.

The right question is not ‘can we trust AI?’

Put that way, the question only invites a vague answer, and artificially pits humans against machines. The useful question is more precise: **in this flow, which AI outputs do we trust, to what degree, to trigger which action, under what control and with what recourse?** It forces us to separate what has no consequence from what affects a right, a sum of money, a customer relationship, a safety instruction or a decision about a person.

Key takeaway. AI does not just create new errors; it accelerates and amplifies existing ones. You need to know where the stopping points are before the cascade starts.

For your leadership team. *Question to ask:* if the AI makes a mistake at the first step, how far can the error travel before someone or something stops it? *Tool:* the row ‘How far an error or misuse could spread’ in the three-flow map (tool 2), played out as a simulation during the workshop (tool 5, step 4).

This risk takes on a new dimension when AI no longer just produces text but acts directly in your systems.

Chapter 4. Agents: when AI moves from adviser to doer

On their first day, you do not hand a new employee the company credit card and the authority to sign contracts. You grant them progressive delegations, based on what they have demonstrated and what you know how to control. That is exactly how to think about AI agents.

What is an agent, in practice?

A conventional AI assistant answers: it drafts, summarises, searches. An **agent** goes further: it combines an AI model with tools and access rights that allow it to **act** in your software. It can look up information, compare, create a ticket, prepare an email, book a slot, change a priority. So it is not its intelligence that makes it autonomous; it is the doors the organisation opens for it.

Imagine the prologue's assistant becoming an agent. In its cautious version, it reads the request, finds the contract, summarises past incidents and proposes a reply. In a more autonomous version, it creates the ticket itself, changes its priority, books an intervention, informs the customer and closes certain requests. Each addition seems harmless. Taken together, they turn a drafting aid into an operational cog in your organisation.

A useful legal point: the term 'agent' is not a separate category in EU law. Obligations depend on the organisation's role, the system's purpose and the risk level of the use, as we will see in chapter 6 [11]. But operationally, autonomy and access to tools make behaviour harder to predict, test and monitor. That is the finding of the UK's National Cyber Security Centre, the NCSC [12].

Four points to watch

Understanding is not the same as having the right. An agent may fully grasp that a customer is asking for a refund. That does not give it the authority to grant one. It is the finance system that must check who can approve, within what limit, for which customer and with what supporting documents.

Everything the agent reads must be treated as untrusted. An email, an attachment or a web page may contain booby-trapped instructions, like the sticky note in the previous chapter. The NCSC stated that, at the time of its analysis, there was no foolproof defence against prompt injection [19]. So the answer is not to ban all reading of external documents, but to make sure that nothing the agent reads can extend its rights.

Speed reduces reaction time. What used to pass through several hands in a day now runs in seconds. If human approval has not been placed in the right spot, that is, before the irreversible action, it comes too late [12].

You must be able to pull the plug. An agent in production must be able to be restricted, paused, revoked or returned to a known state. That means knowing which version, instructions, sources

and access rights it was using, and keeping **logs**: a record of what it received, produced and triggered, errors included [7].

The autonomy dial: four levels

The simplest way to manage these issues is to think in levels, as with a delegation of authority, and to set them **action by action**, not tool by tool: the same assistant may suggest for one task and execute for another. Two questions are enough to place each task: who triggers the action, and who chooses what happens next? These four levels are part of the U-CIE method and do not correspond to any category in the EU AI Act.

Level	What AI does	Example	Condition for reaching it
1. Suggest	It informs a decision: summary, options, alerts	Summarising a file before a meeting	A framed use, verifiable sources, trained users
2. Prepare	It produces a ready-to-go act, which a person approves before execution	Pre-filling a reminder in the CRM, sent after the salesperson reviews it	A named approver with the time, information and authority to review properly
3. Execute	It carries out on its own an operation whose scenario is set in advance	Moving an appointment within authorised slots	A bounded, reversible action; detectable errors; logs; alerts; ability to roll back
4. Orchestrate	It chooses for itself the sequence of several authorised operations to reach a goal	Handling an incident end to end within a defined scope	Explicitly limited options for the agent, real monitoring, tested fallback, ability to stop

Two nuances prevent misunderstandings. First, human approval is only a safeguard if it is real: an approver clicking through eighty files a day in a few seconds each is no longer controlling anything. In some cases, a well-framed level 3 is safer than a level 2 that has become mechanical. Second, level 4 is not banned in principle; it simply requires much stronger evidence. The NCSC recommends starting with tightly bounded tasks and widening access only after satisfactory evaluation and monitoring [12].

The NCSC also offers a common-sense test that any executive can adopt: **if the organisation cannot understand, monitor or contain an agent's actions, it is not ready to be deployed** [12]. The point

is not to predict every sentence the AI will produce, but to know what it can do, where to observe it, what limits it and how to stop it.

This setting is formalised in an **autonomy licence**: a written, dated and revocable authorisation that sets out, for a specific action, what AI may do on its own, what it must have approved and what it will never do. The white paper *The Autonomy Dial* sets out the full method.

Key takeaway. AI autonomy is delegated like signing authority: action by action, with ceilings, and according to the evidence gathered. It can be withdrawn as easily as it is granted.

For your leadership team. *Question to ask:* what level is each AI task at today, what level do we want to take it to, and what evidence do we expect before moving up a notch? *Tool:* the autonomy licence (toolkit, tool 3).

We now know what makes an AI project fragile at scale. The practical question remains: what needs to be built to make it solid?

Chapter 5. What scaling up really changes: seven workstreams

Scaling up AI is not just a project, a platform or an ethics charter. It is a capability built in layers, rather like a house: fine furniture is pointless if the foundations, plumbing and wiring are not up to the job. Each of the seven layers below must match what is at stake in the flow concerned. None replaces the others.

Workstream 1: design the flow before choosing the tool

Before deciding what AI will do on its own, you need to map the expected outcome, the trigger, the steps, the data, the decisions, the exceptions and the end point [3]. This exercise forces questions that are often missing from demonstrations. Which step causes the most delay? Which decision is genuinely difficult? Which exception does the expert handle without ever writing it down? And which act should not be entrusted to the machine, even if it is technically possible?

Workstream 2: treat data as a shared product

Data that matters for a flow must have a definition, an owner, a known origin, a quality rule and an update frequency, like a product delivered to internal customers. An AI system can produce a brilliant answer from contradictory documents. It cannot decide on its own which reference data set is authoritative: that decision belongs to the organisation [5]. The UK government points out that preparing data for AI is never just a question of format: quality, context, security and compliance are all part of it [13].

Workstream 3: separate the ability to understand from the ability to act

This is probably the most important architectural principle in this guide. AI can interpret a request and prepare an intention. But it is the systems that hold the data or carry out the action that must check identity, rights and business rules. Otherwise, the assistant effectively becomes **a super-user disguised as a conversation**: an account that can see and do everything, which people talk to like a colleague.

The principle of ‘least privilege’ means giving each use only the access it strictly needs. An agent that proposes meeting slots has no reason to access payroll files. Temporary access, limited to one action and logged, is often better than permanent permissions granted to speed things up [7].

Workstream 4: name who decides what

There is no need to create a heavy organisation. But for each AI-assisted flow, you should be able to fill in this table without hesitation. In a mid-sized organisation, one person may hold several roles; what matters is that none is left empty.

Role	What it covers
Business owner	The expected outcome and the value created; holds the autonomy licence
Process owner	How the flow runs, its steps and its exceptions
Data owner	The quality, definition and updating of reference data
Technical owner	Operations, connections, monitoring
Risk and compliance	Analysis proportionate to the level of exposure
Holder of the right to stop	The decision to suspend, restart or change the system

NIST does not prescribe a standard organisation chart, but it asks for documented roles and places responsibility for risk decisions at senior leadership level [2]. The international standard ISO/IEC 42001, for its part, offers an organisation-wide AI management system based on continuous improvement [16].

Workstream 5: evaluate the system, not just the model

A model can shine on a set of test questions and fail when faced with incomplete documents, an ambiguous request, a broken connection or a rule that has changed. NIST warns against conclusions drawn from narrow or anecdotal evaluations [10]. To judge the **value** of a use, a useful evaluation looks at three things at once:

- **output quality:** accuracy, completeness, traceable sources;
- **flow quality:** lead times, rework, escalations, errors detected, workload shifted onto other teams;
- **consequences:** impact on customers, finances, people, regulation or safety.

A faster answer improves nothing if it creates twice as many checks downstream. That is often what lies behind a productivity gain that is announced but cannot be found in the accounts. Before granting more **autonomy**, you also need to check control (does the system respect its limits and rights?) and resilience (what happens in the event of change, failure or attack?), as *The Autonomy Dial* explains in detail.

Workstream 6: train for the flow, not for the prompt

The OECD stresses that the skills required go well beyond technical expertise: they include basic digital skills, transferable skills and the managerial ability to decide what belongs to people and

what can be left to the machine [17]. In practice, operators must be able to recognise the tool's limits and report an anomaly; managers must understand how performance is measured; control functions must be able to review a usable audit trail. OECD surveys also link training and consultation of employees with better outcomes for them [18].

Workstream 7: manage a portfolio, not a collection of prototypes

Not every flow deserves the same investment or the same autonomy. It is better to start where value, risk, data maturity and oversight capacity are aligned. A highly profitable task fed by unreliable data may first require a data workstream. A regulated task may call for a slower roll-out without being a bad use case.

An organisation's maturity is therefore not measured by its number of prototypes. It is measured by its ability to make these trade-offs explicitly, and then to repeat the method from one flow to the next.

Key takeaway. Scaling up rests on seven complementary layers: the flow, data, access, roles, evaluation, skills and the portfolio. The weakest one sets the level of autonomy you can reach.

For your leadership team. *Question to ask:* of these seven workstreams, which are actually covered by the plan being presented to us, and which are implicitly assumed to be sorted?
Tool: the five signs that a project is not ready (toolkit, tool 6).

Several of these workstreams have a regulatory dimension. Before turning to the practical cases, it is worth knowing what EU law requires, and from when.

Chapter 6. The European framework, explained simply

Compliance is not a stamp applied at the end. Properly understood, it is a design aid: it forces you to ask, at the right time, the questions this guide recommends anyway.

The principle: the use is judged, not the technology

The EU regulation on artificial intelligence (often called the *AI Act*) does not create a legal category of ‘agentic AI’, and nowhere mentions levels of autonomy. It reasons along three criteria: **the role** of your organisation, **the purpose** of the system and **the risk level** of the actual use [11].

On roles, two concepts are enough to start with. The **provider** is whoever develops an AI system and places it on the market or puts it into service under its own name. The **deployer** is the organisation that uses it in the course of its professional activity. Most organisations are deployers of most of their tools.

A very practical consequence: the same tool can be harmless in one use and sensitive in another. An internal drafting assistant is not, by nature, a high-risk system. The same kind of tool used to screen job applications, evaluate employees or assess the creditworthiness of individuals may fall under high-risk uses. These uses are listed in Annex III of the regulation, a list of areas deemed sensitive (employment, credit, education, access to essential services, among others); classification always depends on a precise analysis of the use [11].

The timetable, checked on 25 October 2026

Regulation (EU) 2026/1744, known as the ‘Digital Omnibus on AI’, entered into force on 27 July 2026. It postponed the heaviest obligations, those for high-risk systems, without postponing the text as a whole [22] [25].

Date	What applies	Status on 25 October 2026
2 February 2025	Prohibited practices, definitions, AI literacy obligation (Article 4)	In force
2 August 2025	Governance rules and obligations for providers of general-purpose AI models	In force
27 July 2026	Regulation 2026/1744 enters into force; Article 4 rewritten	In force
2 August 2026	General application of the regulation, including	In force

Date	What applies	Status on 25 October 2026
	transparency obligations (with transitional periods for certain systems already on the market)	
2 December 2027	Obligations for high-risk systems in Annex III (employment, credit, education, access to essential services, etc.)	Upcoming (originally due on 2 August 2026)
2 August 2028	Obligations for high-risk systems built into regulated products (Annex I: machinery, medical devices, etc.)	Upcoming (originally due on 2 August 2027)

Beware the optical effect of the postponement. It concerns only the regime for high-risk systems. The GDPR, the prohibitions, AI literacy, transparency and your general security obligations continue to apply. And an HR system that will need to comply in December 2027 must be prepared now: compliance involves process changes that cannot be improvised in a few weeks.

AI literacy: an obligation of means, already in force

Since the Omnibus, Article 4 no longer requires organisations to ‘ensure’ a sufficient level of AI literacy. Providers and deployers must **take measures to support its development** among their staff and the people who use the systems on their behalf, without being required to guarantee a set level for each individual [22] [26]. In legal terms, this is a shift from an obligation of result to an obligation of means: you must be able to show that you have taken action.

For a leadership team, the practical stakes remain the same: everyone needs to understand the tool’s capabilities and limits in their role, the usage rules, the risks of disclosure and what to do when in doubt. That is exactly workstream 6 in the previous chapter.

The GDPR does not go away

As soon as an AI flow handles personal data, whether in the documents provided, the instructions, the logs or the outputs, a GDPR analysis remains necessary: lawful basis, purpose, data minimisation, accuracy, retention period, security, processors, transfers outside the European Union and, where needed, an impact assessment [20]. Two pitfalls come up often:

- **pseudonymising is not anonymising.** Replacing a name with a code does not make data anonymous if the person can still be identified. The European Data Protection Board points

out that whether a model trained on personal data is anonymous must be assessed case by case [21];

- **a fully automated decision** that produces a legal effect or significantly affects a person is governed by Article 22 of the GDPR [20]. Involving an approver is not, on its own, enough to fall outside this regime: their involvement must be genuine.

When the deployer becomes a provider

Using a model through a major vendor's service does not make you take on all of that vendor's obligations. However, if you put a high-risk system into service under your own name, substantially modify it or change its purpose to the point of making it high-risk, you may inherit the provider's obligations [11]. This needs checking before every extension of use, and it is one reason why a harmless pilot can, once extended, become a compliance issue.

A pragmatic method

The CNIL, France's data protection authority, recommends starting from a concrete need, defining authorised and prohibited uses, checking what the supplier does with the data, training users and involving the business, IT, the security officer and the data protection officer from the outset [23]. That is the very logic of this guide.

Key takeaway. EU law judges the use, not the tool. Postponing the high-risk regime to late 2027 suspends neither the GDPR, nor AI literacy, nor transparency.

Disclaimer. This white paper offers a governance framework, not legal advice. Classification under the AI Act, the GDPR or sector-specific rules depends on the actual use, the organisation's role and the context. Organisations operating outside the European Union should check which rules apply to them.

For your leadership team. *Question to ask:* for this use, are we provider or deployer, what exactly is its purpose, and could a planned extension change its risk category? *Tool:* the tenth question in tool 1, combined with an initial legal classification by your advisers.

The principles are in place. Let us now see how they play out in concrete situations.

Chapter 7. Five fictional stories to see the mechanism at work

The companies below are fictional. These cases do not describe U-CIE clients or results. They show, in five different worlds, how the same links determine whether scaling up succeeds or fails. Each story follows the same structure: the pilot, the breaking point, the lesson.

Case	Sector	Link that breaks	Lesson in one sentence
Hélianthe Exploitation	Renewable energy	Data (definitions)	A good summary does not say which definition is authoritative
Signal Urbain	Field marketing	Data (status of information)	AI formats information; it does not turn an estimate into proof
Club Partenaires Europe	Sports hospitality	Identity and rights	Preparing an offer is not the same as having the right to commit to it
NordVale Finance	Lending	Decision	Summarising a file is not judging a file
Mécaforge Europe	Manufacturing	Decision and accountability	A quality recommendation must not drift into a production order

Hélianthe Exploitation: the maintenance that seemed obvious

The pilot. This operator of solar and wind farms is testing, on one site, an assistant that reads maintenance reports and produces a summary: equipment concerned, fault, part mentioned, possible next action. A manager reviews everything. Nothing goes automatically to scheduling or purchasing. The technicians save time.

The breaking point. Extended to all the farms, the tool reveals what the pilot had hidden. The same piece of equipment has three different codes in the supervision, maintenance and purchasing systems. 'Unavailability' is not calculated the same way by operations, maintenance and finance. Subcontracting agreements do not allocate responsibilities the same way. The AI still produces

excellent summaries, but it cannot decide on its own which definition prevails, or whether a recommendation justifies a work order.

The lesson. The real workstream becomes the flow ‘from alert to return to service’: a single source of truth for each piece of equipment, a criticality rule, a person who approves, a measurement method. Until this flow is consolidated, AI stays at levels 1 and 2. Success is no longer judged by drafting time saved, but by a shorter qualification time without more errors or repeat interventions [8].

Signal Urbain: the report that lost its source

The pilot. This field marketing agency turns its field managers’ voice notes into campaign reports, sorted by location, time slot, number of contacts and incidents. Each author reviews their text; use remains internal: everything works.

The breaking point. The agency connects the tool to its CRM, the client portal and invoicing. Suddenly, the status of every piece of information matters. ‘Qualified contact’ does not mean the same thing to the client, the sales team and the field staff. The notes contain opinions, photos and sometimes contact details. An estimate made on the spot in the street can end up in a client report as a firm figure, simply because the AI worded it well.

The lesson. Raw observations, management indicators and what can be shared with the client must be kept apart. Contractual figures come from identified sources, every publication has an owner, and personal data is reviewed under the GDPR [20]. AI speeds up formatting; it does not turn uncertain data into proof.

Club Partenaires Europe: the sales agent with too many connections

The pilot. This company organises hospitality packages around major sporting events: boxes, meals, guest reception, logistics. An assistant retrieves past exchanges with a client and proposes a reply, which the salesperson reviews, corrects and sends themselves. The assistant touches neither bookings nor contracts. This is level 1, then level 2 once proposals are pre-filled in the CRM.

The breaking point. The plan is to connect the tool to booking and invoicing: checking availability, placing an option on a box, chasing the client. But an offer depends on availability, contractual clauses and payment terms, and an option once placed may be invoiced. Overly broad access also exposes negotiation information, and a booby-trapped incoming email could influence an agent that has tools to act.

The lesson. Four tasks are separated: search, recommend, prepare, execute. The agent may search and prepare. The CRM, the booking system and the finance tool check identity, availability, pricing rules and commitment ceilings before any action [6]. Anything that commits the company (firm offer, option, discount, payment) is not delegated without an explicit decision by the leadership team. The white paper *The Autonomy Dial* follows this company for twelve months, licence by licence.

NordVale Finance: the line between assistance and decision

The pilot. A lender's team is testing an AI tool that summarises the documents in a loan file, flags missing items and prepares a note for the analyst. The analyst keeps access to the sources, their judgement and the decision. The tool saves research time, nothing more.

The breaking point. Extending it to the decision system, external data and the customer onboarding journey completely changes the picture. Assessing the creditworthiness of individuals is among the high-risk uses listed in Annex III of the AI Act, depending on its actual purpose [11]. A fully automated decision that significantly affects a person is also governed by Article 22 of the GDPR [20].

The lesson. The line between summary, recommendation and decision must be held. For the summary, you check faithfulness to the documents, detection of gaps, traceability and confidentiality. For a recommendation or a decision, you must make visible the criteria, the data used, the oversight, the explanation and the routes of appeal. The dangerous shortcut would be to treat time saved on reading as a justification for automating judgement.

Mécaforge Europe: the recommendation that became a production order

The pilot. On one site, an AI tool reads non-conformities, maintenance history and operator comments to prepare quality meetings. It helps with diagnosis; it triggers nothing. Level 1.

The breaking point. Rolling it out across several plants, connected to production control, management and planning systems, brings AI into the flow 'from quality deviation to production decision'. But material references, units, tolerance thresholds and quarantine rules vary by site, product or customer. An operator's comment is a hypothesis, not a measurement. A wrong recommendation to stop or adjust a line can affect quality, safety or deliveries.

The lesson. Three tasks are distinguished: detect, recommend, act. Each requires qualified data, explicit rules, an authorised person and appropriate traceability. Without a shared language for business objects, connecting systems does not remove ambiguity [5]. Automatically opening an analysis ticket on a precisely defined deviation may fall under level 3; any decision affecting production stays at level 2.

Key takeaway. In all five stories, the AI works. What breaks is a link: a definition that is not shared, an unclear information status, an overly broad right, a decision boundary crossed without anyone saying so.

For your leadership team. *Question to ask:* which of these five stories do we recognise ourselves in most, and on which current project? *Tool:* choose that project as one of the three flows in the workshop (toolkit, tool 5).

These stories have one thing in common: each required investment that was not in the pilot's budget. How should this be read in a funding plan?

Chapter 8. Reading a scale-up budget: what you are really paying for

When a scale-up budget for AI is presented to the leadership team, it often looks like a technology line: a platform, connectors, licences. That is an incomplete reading. What you are really paying for is **the organisation's ability to transform its critical flows without losing control of its decisions**.

The seven workstreams in chapter 5 translate, in a budget, into six investment lines. The table below helps check that none has been forgotten.

Investment line	Workstreams covered	What to look for in the budget
Operational reality	1	Time from business experts to describe the real flow, exceptions included
Data and connections	2, 3	Reference data, quality, sharing rules, monitoring of connections
Security and control	3	Bounded rights, logs, suspension mechanism
Operating model	4, 7	Named roles and time allocated to fulfil them
Evidence	5	Indicators measured before launch, decision thresholds
Learning	6, 7	Flow-based training, regular portfolio reviews

Operational reality. Describing the flow as it really runs, with its workarounds, exceptions and trade-offs, takes time from people who have little to spare. It is less spectacular than a demonstration. Yet it is what prevents AI being plugged into a theoretical procedure that nobody follows any more.

Data and connections. Not all data needs to be perfect for a first use. However, any data that triggers a significant action must have a status, an owner and a known level of confidence. In the prologue story, this is the line that would have prevented the confusion over 'active customers' and the broken connection.

Security and control. ENISA, the EU agency for cybersecurity, recommends a layered approach: first the cybersecurity fundamentals, then AI-specific measures, then sector-specific requirements [24]. In other words, sophisticated AI never makes up for a weak security foundation.

Operating model. A central team can provide standards, reusable building blocks, an evaluation method and legal and security expertise. The business keeps ownership of the outcome and of the transformation of its flow. The right organisation is neither fully centralised nor fully decentralised: it is the one that makes end-to-end responsibilities visible and decides quickly.

Evidence. Define before launch the indicators for quality, lead time, rework, risk and value, along with the thresholds that will lead to extending, correcting, suspending or stopping. This discipline guards against two temptations: keeping a tool because it impresses, or abandoning too early a transformation that needed some adjustment. Economists talk about a 'J-curve': measured productivity may first stall, or even fall, while the organisation invests in invisible assets (processes, data, skills), before the benefits appear [4]. This is not a promise of future gains; it is a reminder to judge a programme against indicators set in advance, rather than on the impression of the first few months.

Learning. A mature organisation does not try to make all its flows autonomous. It learns, flow by flow, what degree of assistance is relevant, where data quality limits automation, where human control still adds value and where deeper integration becomes justified. That is what turns a series of isolated bets into a portfolio of deliberate decisions.

Key takeaway. An AI budget that funds only the tool is really funding one more pilot. The links cost time, organisation and discipline, and they are what make the tool pay off.

For your leadership team. *Question to ask:* in this budget, what share funds the tool, and what share funds the links that will make it reliable? *Tool:* the six-line table above, and the three families of indicators (toolkit, tool 4).

Back to the prologue's leadership team, in October.

Epilogue: back to the original question

In October, the prologue's leadership team was asking whether its project 'worked, yes or no'. This guide suggests replacing that binary question with a more useful one: **which flows do we want AI to make more capable, and which links must we strengthen so that this capability remains trustworthy at scale?**

Asked this way, the answer is no longer a verdict on a tool. It is a plan: chosen flows, identified links, levels of autonomy deliberately set action by action, indicators defined in advance. A plan the leadership team can follow, adjust and defend.

Toolkit

Seven tools, to be used in order or separately. They do not replace a detailed process, security or compliance analysis; they help bring to light, quickly, the links that will determine whether scaling up succeeds.

Tool 1. Ten questions to ask before approving a scale-up

To be ticked off in the meeting. A box left empty is not a refusal: it is a workstream to fund.

- ☐ Do we know what the pilot has proved, and what it has not proved yet? (*chapter 1*)
- ☐ Is the complete flow described as it really runs, workarounds included? (*chapter 2*)
- ☐ For each critical piece of data, do we know the authoritative source, its owner and how up to date it is? (*chapter 2*)
- ☐ Does the AI act with the user's rights, and is it our software that checks permissions? (*chapters 3 and 5*)
- ☐ For each AI task, what level are we at (suggest, prepare, execute, orchestrate), and who decided? (*chapter 4*)
- ☐ Are the planned human approvals genuine, and placed before any irreversible action? (*chapter 4*)
- ☐ Has an end-to-end owner been named, with the right to stop the system? (*chapter 5*)
- ☐ Were quality, flow and consequence indicators measured before launch? (*chapters 5 and 8*)
- ☐ Can we pause, roll back and reconstruct what happened? (*chapter 4*)
- ☐ Have our role under the AI Act and the GDPR analysis been established for this specific use? (*chapter 6*)

Tool 2. The three-flow map

Choose three flows with different consequences: one that creates or protects revenue, one that delivers an operation or service, and one that grants or controls an important right (an expense, an access right, a contractual commitment, a decision about a person). Bring together the business owner, the process owner, **someone who does the work day to day**, the representative for the data or software concerned, and security, the data protection officer or compliance if the flow warrants it.

Golden rule: when nobody knows the answer, write 'unknown'. That in itself is valuable information.

What to map	Flow 1	Flow 2	Flow 3
Name of the flow and expected business outcome			
What triggers it, what really ends it			
The real steps, workarounds included			
Decisions and exceptions			
Critical data: source, owner, freshness			
Connected software, files and suppliers			
Who is allowed to do what, up to what amount			
Tasks entrusted to AI and the level of each: suggest, prepare, execute, orchestrate			
Human approval and right to stop			
How far an error or misuse could spread			
Audit trails, evidence and indicators of value and risk			
Decision: keep the current level, move up a level, or first resolve a dependency			

After the workshop, do not look for an overall score. Look for unresolved links. A promising flow may first require work on reference data or rights. Conversely, a flow that seems simple should stay at level 1 or 2 until the organisation can understand, monitor and contain the actions involved.

Tool 3. The autonomy licence

One licence per action, not per tool. It is held by the business owner and granted by the leadership team. It sets out in black and white what AI may do on its own, what it must have approved and what it will never do, along with the conditions under which this right may change or be withdrawn. The full method, with a completed example, is in the white paper *The Autonomy Dial*.

Section	Question to answer	Your answer
Use and purpose	What problem is solved, what value is expected, how is it measured?	
Action concerned	What can AI read, prepare, create, modify, send or coordinate?	
Authorised level	Level 1 to 4, effective date, review date	
Scope	Data, tools, rights, amounts, volumes, hours, areas, exclusions	
Prohibitions	What AI will never do, even to reach the objective	
Decision and oversight	Holder, approver, alert recipient, who can stop and restart	
Evidence	Tests, reference sources, indicators, known incidents, logging	
Resilience	Critical dependencies, degraded mode, fallback procedure, date of last exercise	
Review	Events that trigger a reassessment; conditions for moving up and down	

Tool 4. Three families of value indicators

Family	What it measures	Examples
Output quality	What AI produces	Accuracy checked on samples, traceable sources, completeness
Flow quality	What happens to the work around it	End-to-end lead time, rework rate, escalations, cancelled actions, workload shifted onto control
Consequences	What the organisation experiences	Complaints, incidents, financial discrepancies, reports, customer and employee satisfaction

Measure these three families **before** launch, so you have a baseline. Without a baseline, no gain can be demonstrated and no drift can be spotted. That is precisely what the prologue's leadership team was missing in October.

Tool 5. The leadership team workshop (2 hours)

Objective: leave the session with a first map of three flows and three priority workstreams.
Participants: chief executive, finance director, IT director, the business owners of the chosen flows and, where relevant, legal, compliance or data protection.

1. **Framing (15 min).** The executive sponsor explains why the organisation is looking at AI through flows, and the group chooses the three flows to study.
2. **The real flow (30 min).** The person who does the work describes a recent case, from trigger to outcome. The others listen without correcting.
3. **Spotting the links (30 min).** The three-flow map (tool 2) is filled in, writing 'unknown' without hesitation.
4. **Cascade simulation (20 min).** For each flow, the group asks: 'If the AI gets it wrong here, how far does the error go?'
5. **Levels of autonomy (15 min).** For each task, the group sets the reasonable level today and the evidence needed to move up a notch.
6. **Decisions (10 min).** Three priority workstreams, an owner for each, a review date.

Tool 6. Five signs that a project is not ready

- Nobody can say who has the right to stop the system.

- The assistant uses a technical account with broader rights than its users.
- Success is measured only by the satisfaction of the first users.
- Exceptions are 'handled by the project team' without being written down anywhere.
- Nobody can reconstruct what the AI read, produced and triggered last week.

Tool 7. Glossary for the leadership team

Term	In plain terms
Agent (agentic AI)	An AI system that can use tools and chain actions in software, rather than just producing text
API	A software 'socket' that allows one program to read information from another or ask it to act
Autonomy licence	A written, dated and revocable authorisation that sets out what AI may do on its own for a given action (U-CIE method)
Confabulation	The confident production of false information by an AI system; also called hallucination
CRM, ERP, CMMS, MES	Management software: customer relationship, enterprise resource planning, maintenance, production execution
Deployer	Under the AI Act, the organisation that uses an AI system in its professional activity
Flow	The real sequence of steps, decisions and handovers that turns a trigger into a business outcome
High-risk system	A category in the EU AI Act (employment, credit, education, etc.) subject to stricter obligations from 2 December 2027
Least privilege	The principle of giving a system only the access strictly necessary for its task
Levels of autonomy	Suggest, prepare, execute, orchestrate: the four possible settings for a task entrusted to AI (U-CIE method)
Link	Whatever allows one step in a flow to receive the information, right or decision the next step needs
Logging	Recording what the system does, when, with which data and which version

Term	In plain terms
Model	The AI 'engine' that understands and generates text; it is only one part of the system
Pilot	A trial on a restricted, protected scope, designed to check whether a use is valuable
Prompt	The instruction or question addressed to the AI
Prompt injection	Malicious text, often hidden in a document or email, that tries to hijack an AI system's behaviour
Provider	Under the AI Act, whoever develops an AI system and places it on the market or puts it into service under its own name
Reference data	The data source recognised as authoritative for a business object (a customer, a piece of equipment, a contract)
Semantics	The meaning given to words and data; two systems can be connected without sharing the same meaning

About U-CIE

U-CIE helps organisations think through, build and steer their AI from A to Z. Rather than adding yet another tool, we start by mapping with your teams how your organisation really works, to identify where AI can make a difference.

Think. Map the flows that matter and the links to strengthen before any scale-up.

Build. Design with you generative or agentic AI uses, setting out for each what it may do on its own, what it must have approved and what it will never do.

Steer. Put in place the indicators, reviews and rules that make it possible to measure value, monitor risks and adjust autonomy at the pace the organisation can control.

Our conviction fits in one sentence: AI should transform your organisation without ever slipping out of your control.

To map your first three critical flows with your leadership team, contact U-CIE at u-cie.com.

References

Legal position and sources checked on 25 October 2026.

1. OECD, [The Adoption of Artificial Intelligence in Firms: New evidence for policy making](#)
2. NIST, [Artificial Intelligence Risk Management Framework \(AI RMF 1.0\)](#)
3. NIST, [AI RMF Playbook: Map](#)
4. Brynjolfsson, Rock, Syverson, [The Productivity J-Curve: How Intangibles Complement General Purpose Technologies](#), NBER
5. Joint Research Centre (JRC), [Artificial Intelligence for Interoperability in the European Public Sector](#)
6. OWASP, [Top 10 for LLM Applications v2025](#)
7. NSA, CISA and partners, [Deploying AI Systems Securely](#)
8. NIST, [AI RMF Playbook: Measure](#)
9. NIST, [Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations](#)
10. NIST, [AI RMF: Generative Artificial Intelligence Profile](#)
11. European Union, [Artificial Intelligence Act, consolidated text](#)
12. NCSC, [Thinking carefully before adopting agentic AI](#)
13. UK Government, [Guidelines and best practices for making government datasets ready for AI](#)
14. European Commission, [Ethics guidelines for trustworthy AI](#)
15. DSIT, [Portfolio of AI assurance techniques](#)
16. ISO, [ISO/IEC 42001:2023, Artificial intelligence: Management system](#)
17. OECD, [Skill needs and policies in the age of artificial intelligence](#)
18. OECD, [Using AI in the workplace: Opportunities, risks and policy responses](#)
19. NCSC, [Thinking about the security of AI systems](#)
20. European Union, [General Data Protection Regulation](#)
21. European Data Protection Board, [Opinion 28/2024 on certain data protection aspects related to AI models](#)
22. European Union, [Regulation \(EU\) 2026/1744, Digital Omnibus on AI](#)

23. CNIL, [Comment déployer une IA générative ? La CNIL apporte de premières précisions](#) (in French)
24. ENISA, [Multilayer Framework for Good Cybersecurity Practices for AI](#)
25. European Commission, [AI Omnibus enters into force](#), 27 July 2026
26. AI Act Explorer, [Digital Omnibus: replacement of Article 4](#) (in French)