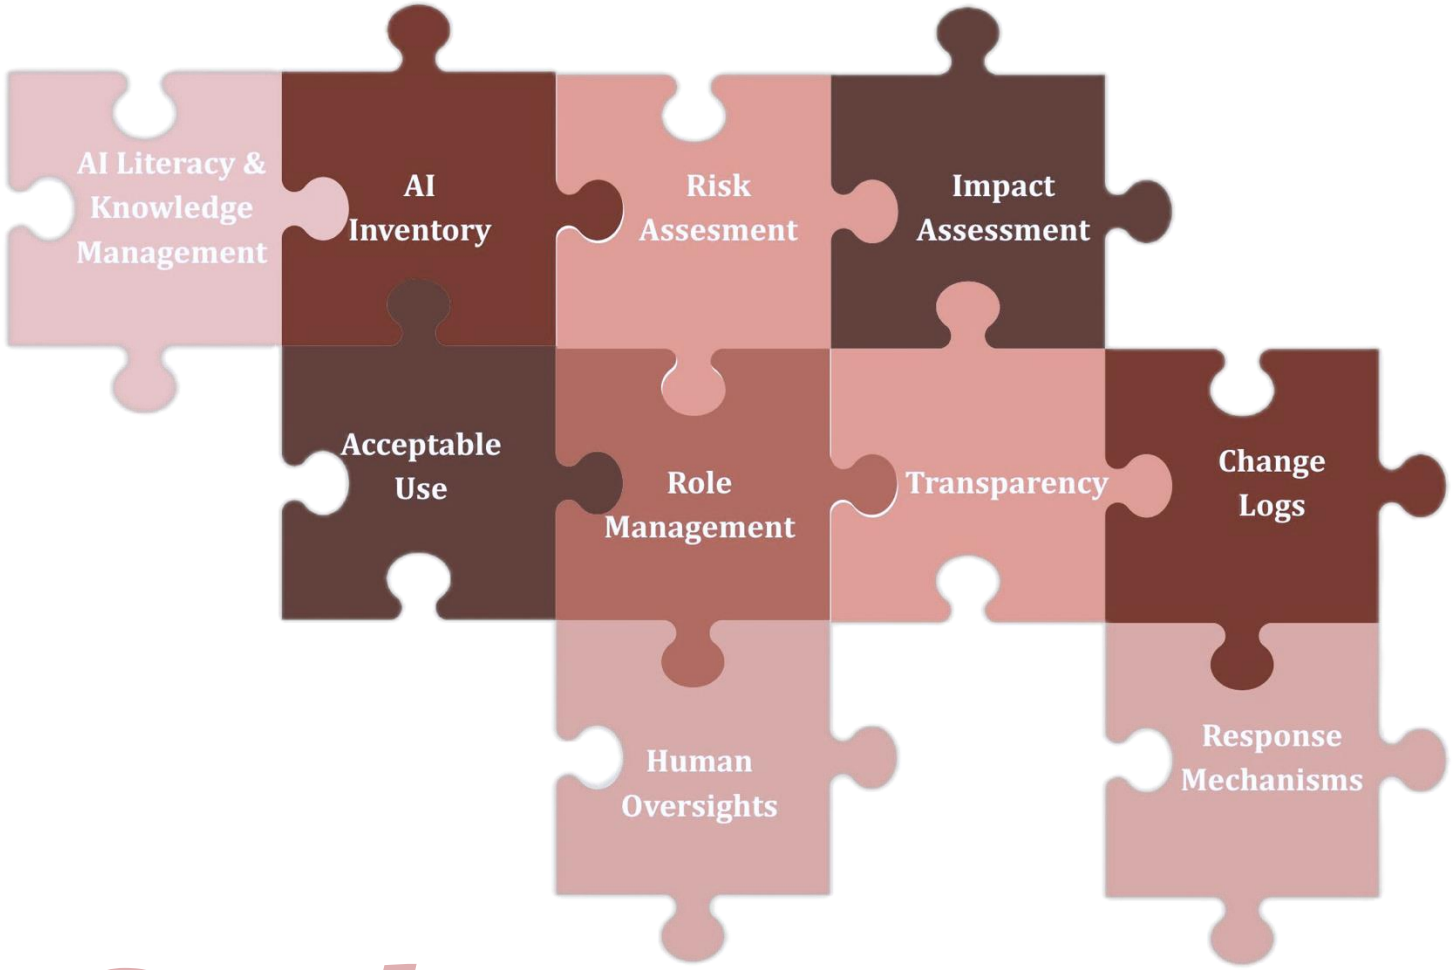

Building Your *AI Governance Stack*

How small teams can govern AI without building a second compliance function

In this *session:*

1. **The audience** — *governance owned by someone who already has a job*
2. **The structure** — *carry it, adapt it, or add it — the general idea in one slide*
3. **The baseline** — *one list, seven columns, and how to go about shadow AI*
4. **The assessments** — *risk and impact assessments for getting the ball rolling*
5. **The human layer** — *transparency, accountability and functioning oversight*

A policy is *not a control*

86%

of respondents work at organizations **they believe** have AI policies in place

72%

believe they know AI better than their own tech teams

88%

have shared work-related information with public AI tools such as ChatGPT, Claude, or Gemini.

81%

believe leadership operates under a different set of rules

66%

used unauthorized AI tools in their work

31%

have shared financial information or confidential company documents and strategies

PagerDuty Survey of 1,250 office professionals · US, UK, AU, JP · companies over $500m revenue — organisations that have a compliance function.

60

seconds

Write down every place AI touches your organisation.

Paper, phone, whatever. Go.

The four *you (most likely) didn't count*



The AI you didn't buy

Your CRM, helpdesk, HR platform, notetaker.

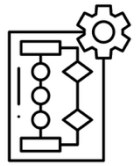

AI in your vendors' vendors

Your processor added an AI subprocessor.



Private tools

Individual accounts, personal logins.



AI you're subject to

The screening tool that filters grant applications.

Sometimes there is no AI governance gap.

But there is a routing problem.

You already have

- Lawful bases and a record of processing
- Vendor onboarding with a security review
- Access controls
- Data classification processes
- A breach procedure
- Contracts with data protection terms/DPAs
- An annual training cycle

~60–70%

of the control surface is already installed. So why does it feel like nothing?

Because no process you own has a sign on it that says AI, so the use case walks straight past them.

Carry · Adapt · *Add*

CARRY

It already works?
Don't touch it.

- Access controls & authentication
- Encryption and retention
- Data classification
- Vendor security baseline
- Contracting machinery
- Incident intake channel
- Training infrastructure

ADAPT

It exists, but doesn't quite work?
It needs new fields and new triggers.

- The software overview
- Risk & impact assessments
- Acceptable uses
- Procurement procedure
- Incident taxonomy and logging

ADD

Genuinely new?
There are a couple you'll have to add…

- AI literacy, role-specific
- AI Inventory
- Human oversight
- Pre-deployment evaluation
- Post-deployment monitoring

The Adapt and Add columns are the work. Today we start it.

Why the third column *has to exist*

Software your controls were built for

Fixed rules, meaning

- Same input, same output
- It only changes when someone files a change request

AI systems

Machine learning as the foundation:

- Same input, different output
- The vendor updates the model on a Tuesday, and your control environment has changed

Nothing in a classic compliance stack watches for systems that quietly change.

No test fails. No notification you'd recognise. No change request. That is what the *Add* column is for.

Building Your
AI Governance Stack
The Baseline

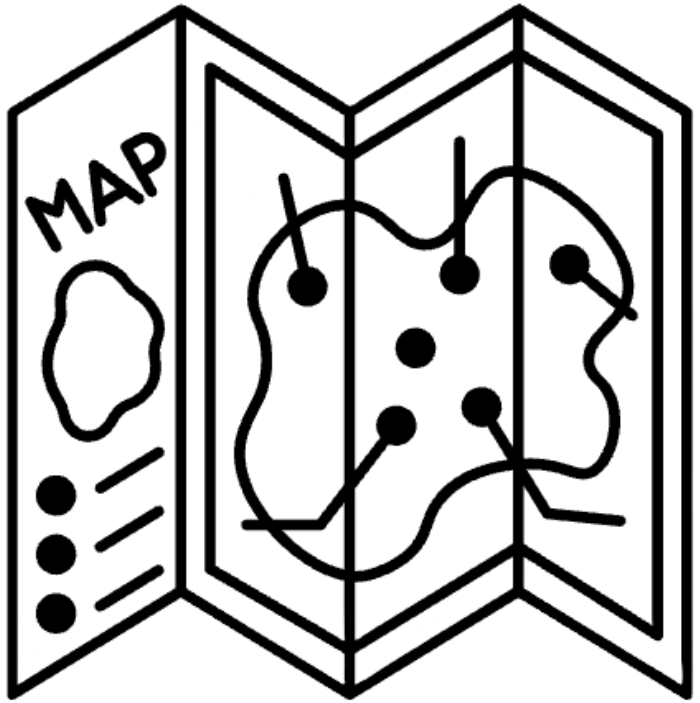

Shadow AI is *a rational response*



Under the waterline: *four exposures*



Personal data

beneficiary and staff data processed with no basis, no DPA, no transfer mechanism



Confidentiality & IP

case files, strategy, legal advice disclosed to a third party



Quality & harm

unchecked outputs reaching real decisions: a mistranslated interview, a wrong entitlement, a hallucinated deadline



Duties you can't meet

no inventory, no oversight, no logs — you can't answer questions about a system you didn't know existed

Transparency, *facing inward*

Staff should never wonder...

- which AI tools are approved (and which died)
- what each tool is allowed to touch
- who owns a system when something looks off
- where the rules and the list actually live

The scrappy version

- one intranet page: tools, owners, rules
- a standing agenda slot in team meetings
- announce every new tool + its one-pager
- the same page links the request fast lane

The policy *people actually follow*

Why we use this AI system	two sentences tying AI use to the mission and to the limits of that use
Our red lines	what we will not do, ever regardless of efficiency (e.g. no AI-only decisions about people)
Our risk appetite	which risk tiers we accept, which need exec sign-off, which we refuse
How we govern	one paragraph naming the components: inventory, assessments, training, oversight
Who answers	the four hats from subpart A — by role title, not by person's name
Review	yearly, by the board — and after any serious incident

Keeping data out: *controls that actually hold*

1. **Provide a safe alternative:** an approved tool that fits the purpose — the single most effective control
2. **Make the safe path the easy path:** logged in, bookmarked, in the workflow
3. **Teach redaction:** strip names and identifiers before pasting; abstraction usually works just as well
4. **Technical guardrails:** block known endpoints, DLP rules on obvious identifiers — a band-aid, not a strategy
5. **Name the red line, once:** no beneficiary or confidential data in public tools

Controls 1 and 2 do 80% of the work. Blocking alone just relocates the behaviour to personal devices.

The AI discovery hunt *(gently)*

1 · Ask, don't blame	a short anonymous survey + a no-blame window, announced by leadership	*surfaces the most, fastest — if trust holds*
2 · Follow the clues	expense claims, card statements, team subscriptions, previous requests	*finds the paid tools, and the enthusiasts*
3 · Check the plumbing	SSO logs, network traffic, browser extensions — proportionately, and announced	*finds what people forgot to mention*
4 · Make it recurring	a light quarterly pulse keeps the picture current	*a one-off hunt ages in weeks*

The AI Inventory: One list *to rule them all*

Seven columns beat an unused governance platform

System	*what it is*
Purpose	*what it's for*
Owner	*one name, accountable end to end*
Data it touches	*and whether it leaves your tenant*
Risk tier	*from your triage with a link to the AI Act*
Vendor & terms	*including whether your data trains their model*
Status	*approved, and the next review date*

If you have one, don't build a second one

add a view to the register you already have — two lists disagree by month four

Owner is the key column

a system without a name next to it is ungoverned

Risk tier is a connector

it's what connects this list to your assessment routine

This is the load-bearing wall

assessments, training records, and contracts all hang off it

Transparency, *facing outward*

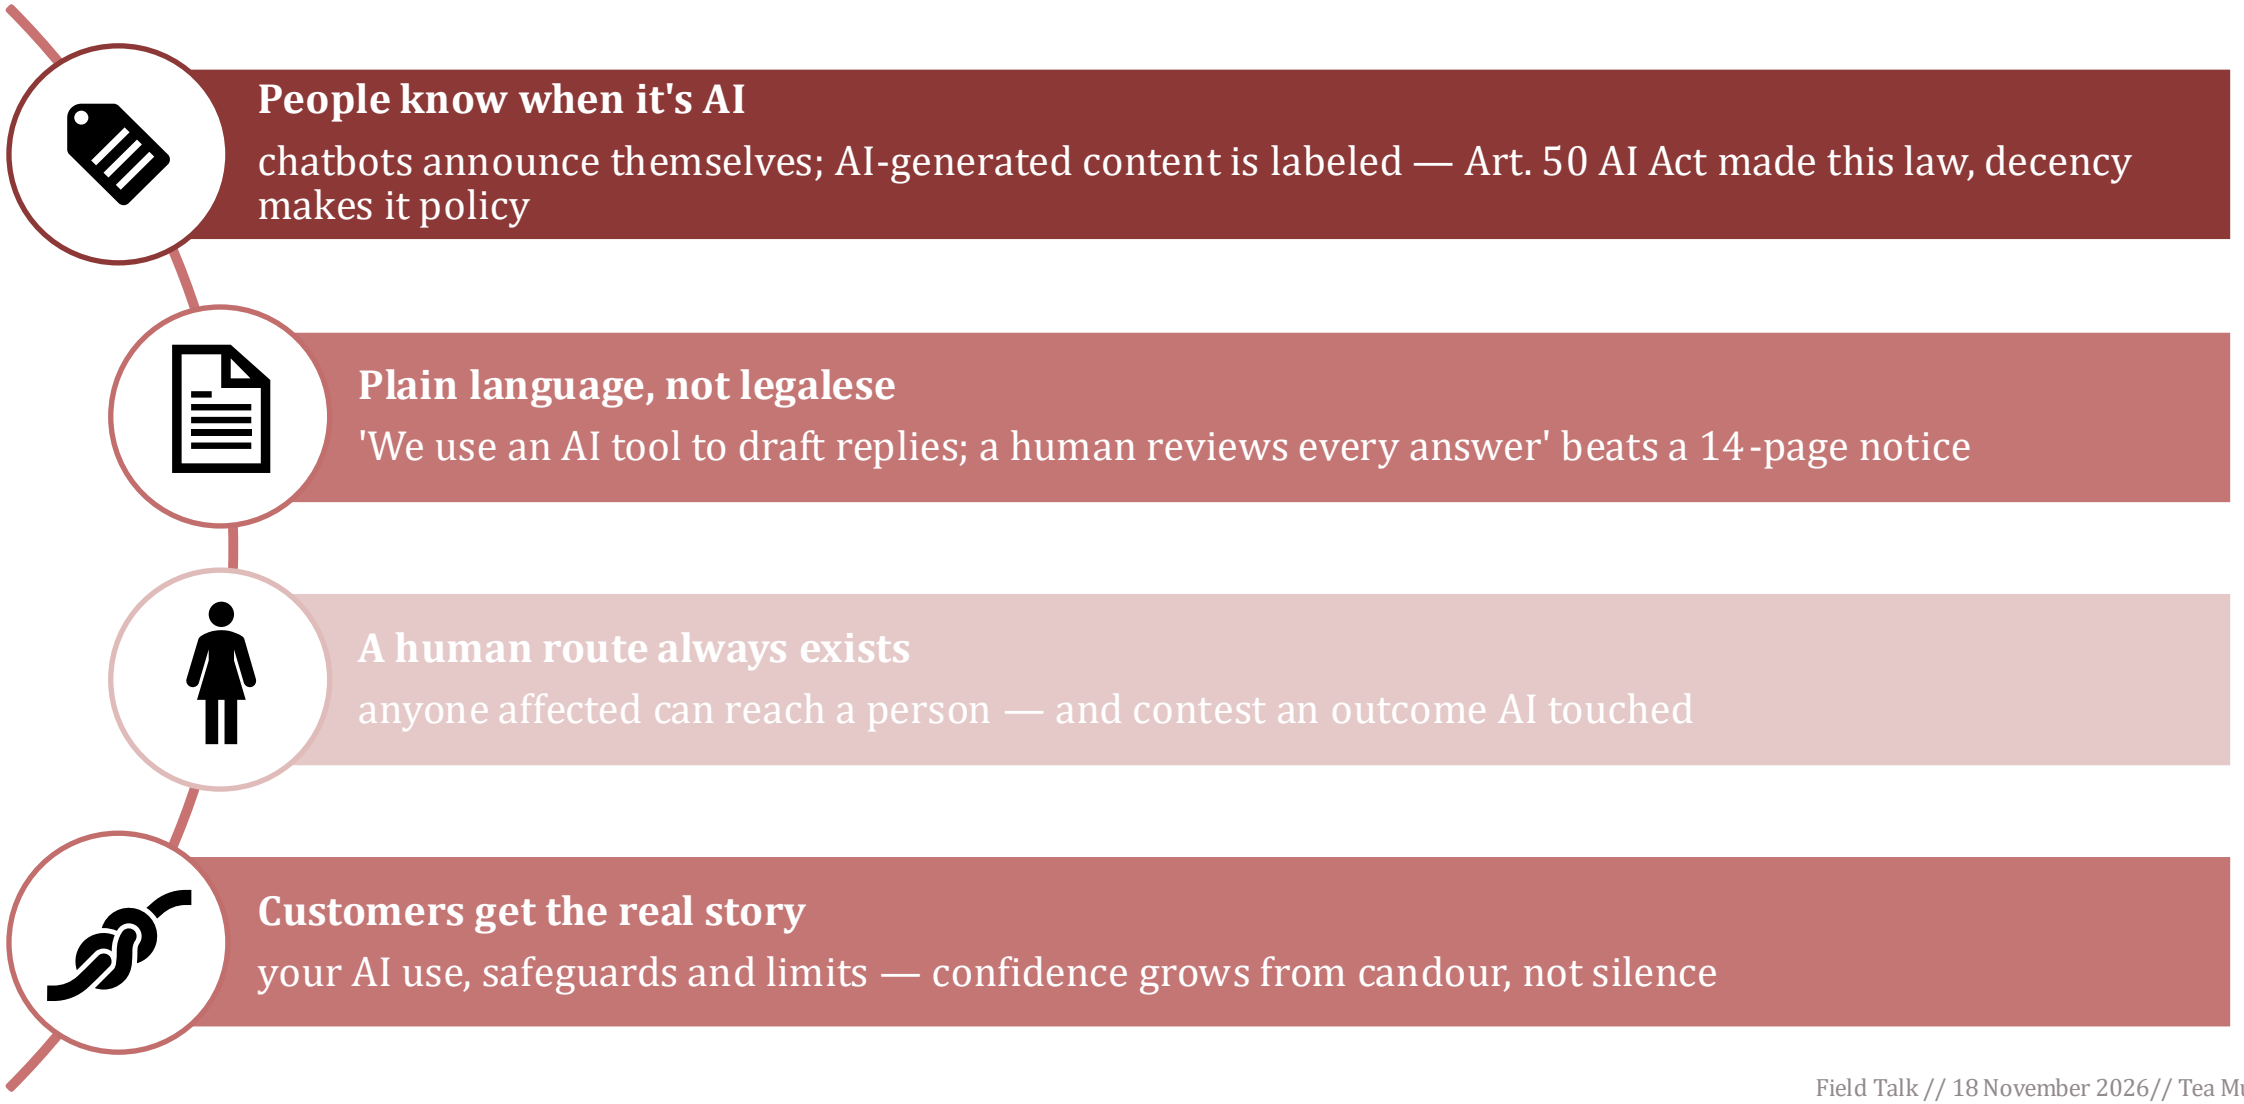

Building Your
AI Governance Stack
The Assessments

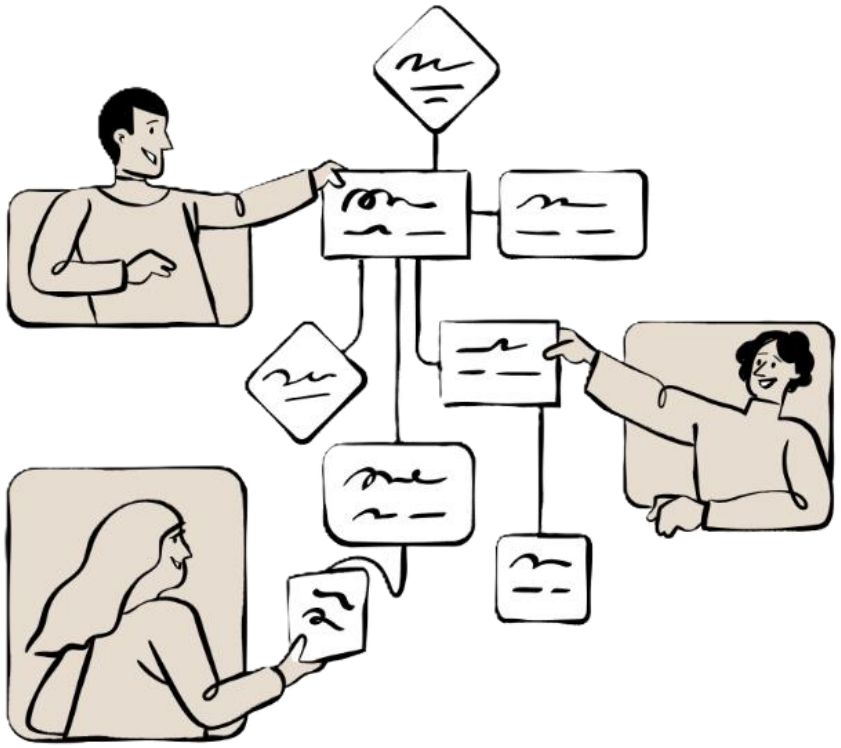

AI Impact Assessment(s) *Project Initialization*

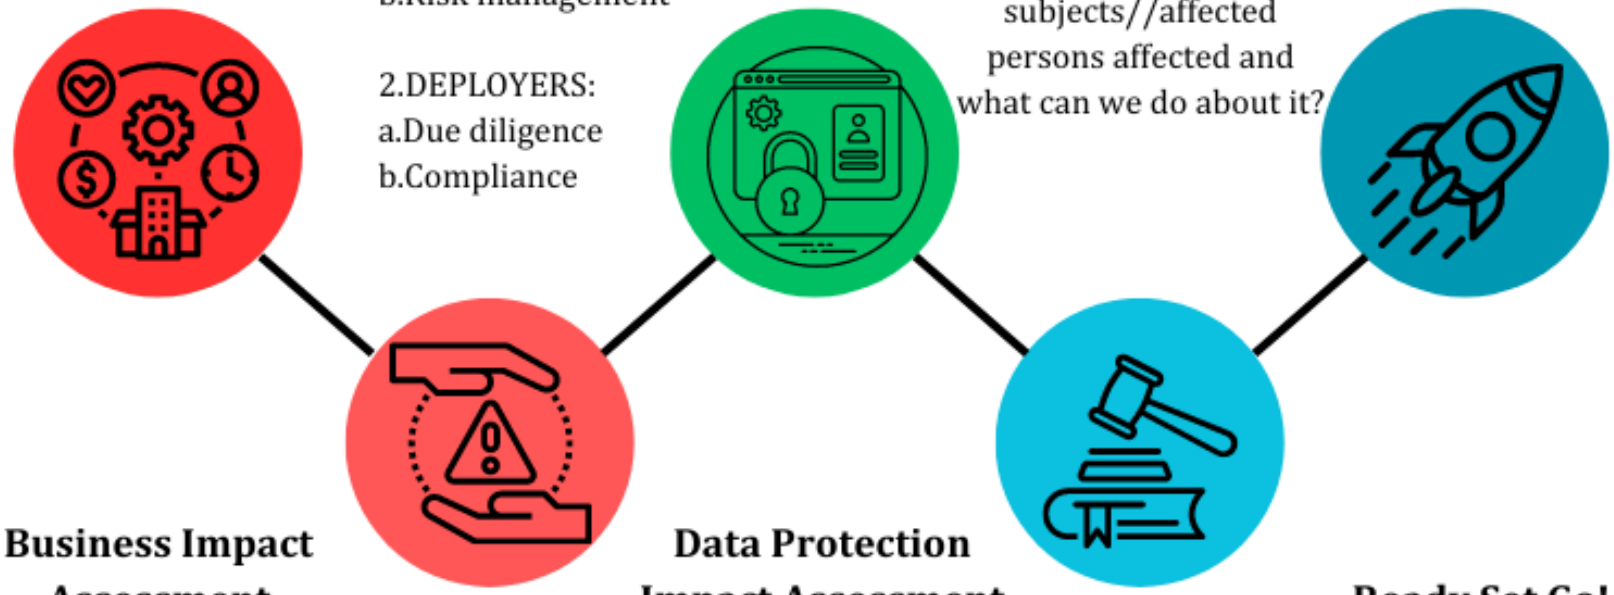

Business Impact Assessment

Use cases, expected benefits, expected costs, etc.

Risk Assessment

1.PROVIDERS:
a.Risk ordering
b.Risk management

2.DEPLOYERS:
a.Due diligence
b.Compliance

Data Protection Impact Assessment

How are the data subjects affected and what can we do about it?

Fundamental Rights Impact Assessment

How are the model subjects//affected persons affected and what can we do about it?

Ready Set Go!

Compliance by design time.

Shadow AI is not a discipline problem.

It's a process problem and you build the queue.

If every AI request gets the same review, you are the bottleneck and people rarely wait.

Your employees will route around you.

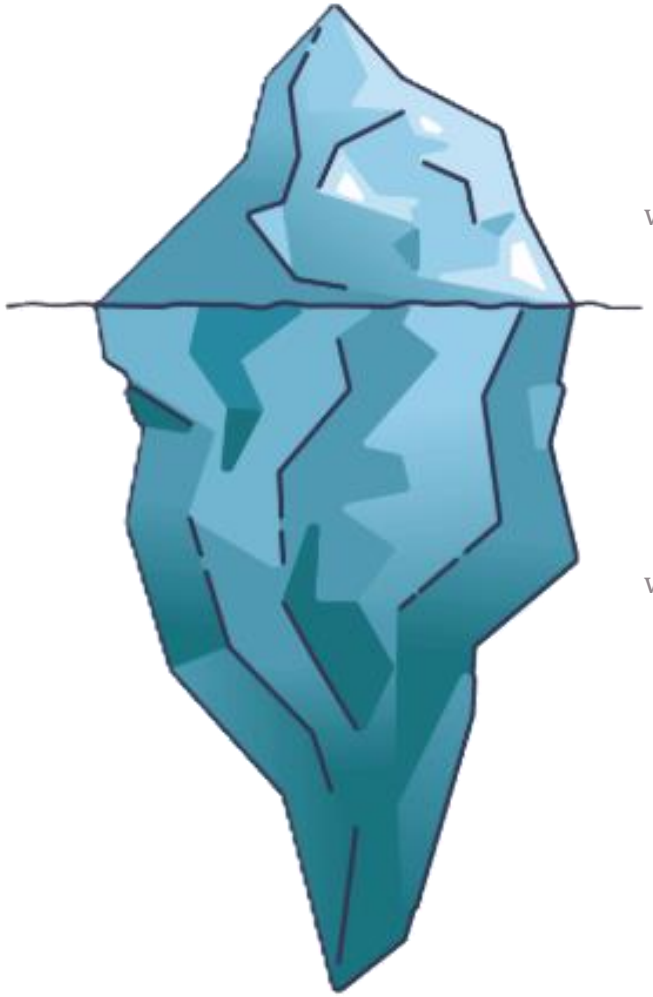

The risk triage: *three questions, two minutes*

1. **Does it touch people?**
   outputs seen by, about, or affecting humans — or just internal plumbing?
2. **Does it make decisions about people?**
   influence on access, money, opportunities, wellbeing — even 'just recommendations'?
3. **Does it touch sensitive data?**
   personal data, beneficiary records, special categories, confidential material?

The risk triage: *three questions, two minutes*

0 × yes → fast lane	*log it in the inventory*
1 × yes → start with a one-pager	*a lightweight assessment to begin with*
2-3 × yes / high-risk → full treatment	*structured and thorough assessment(s)*

The one-pager: *think before you deploy*

Lightweight AI Impact assessment — one page, seven boxes

Purpose	what job does this system do, for whom?
Data	what goes in, what comes out, where is it stored?
People affected	staff, beneficiaries, donors, the public?
What could go wrong	worst realistic failure — write it in one honest sentence
Mitigations	checks, oversight, limits that reduce that failure
Owner & sign-off	who runs it, who approved it
Review date	when we look at this again — 12 months max

30–60 minutes

with the system owner in the room

The hard truth

Q4 is the main point

Filed, not framed

linked from the inventory row it belongs to

The full treatment: *AI Risk Assessment Cycle*

The one-pager

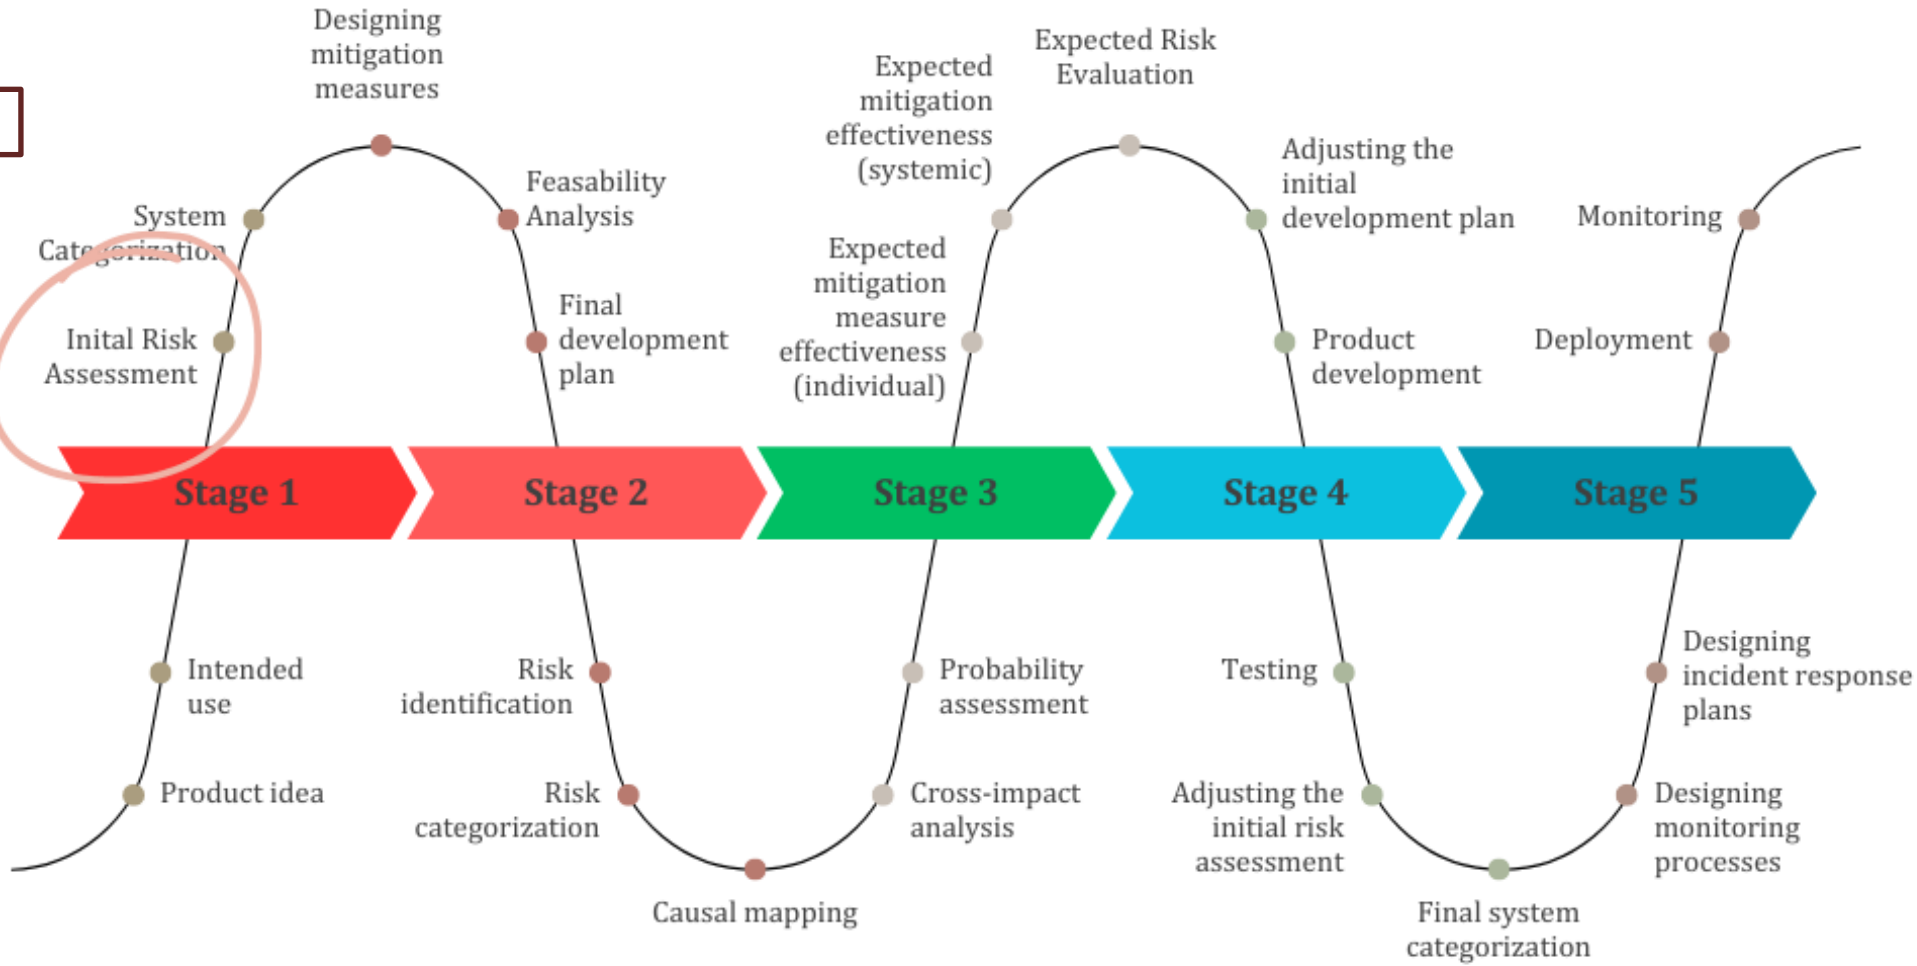

The full treatment: *Risk Management*

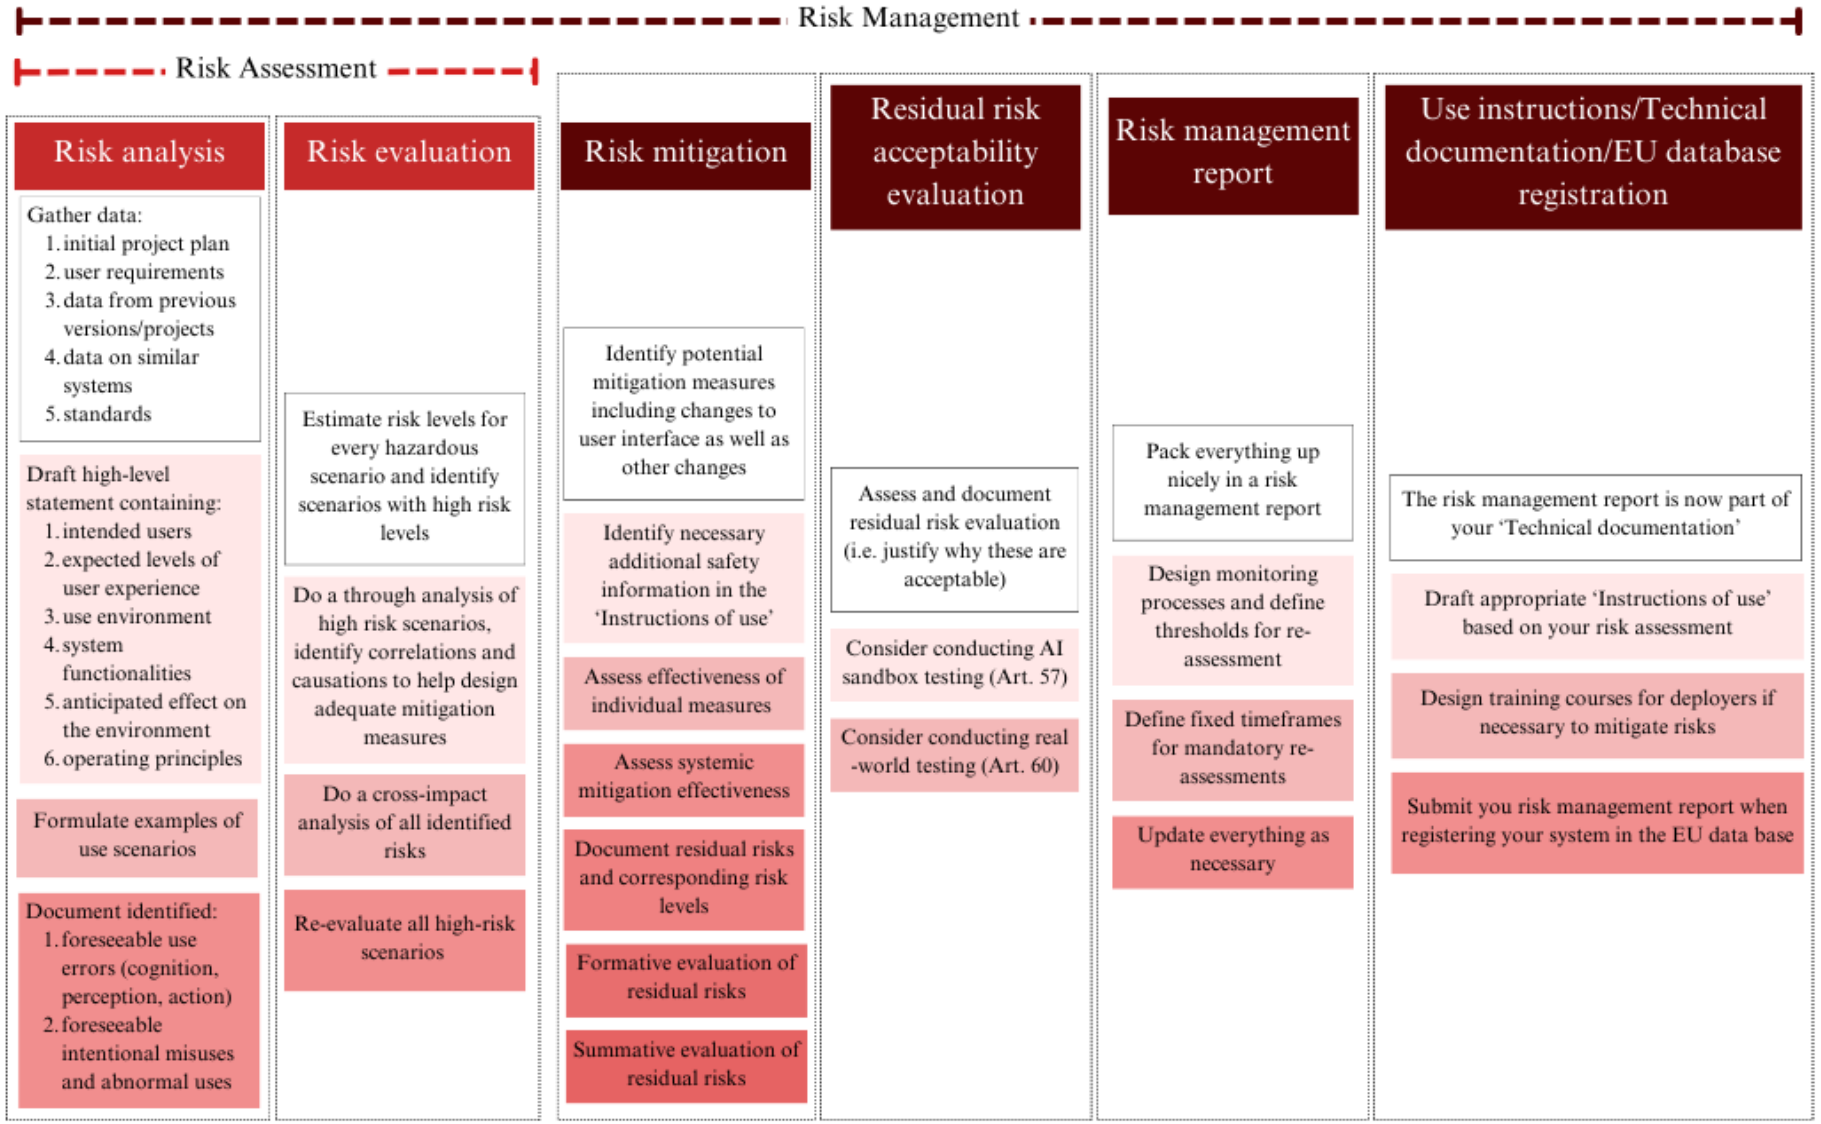

Assessments are *verbs, not documents*

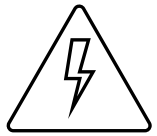

Re-open on triggers

material change, vendor update, new user group, incident, or 12 months — whichever comes first



Findings become tasks

every 'what could go wrong' gets a mitigation with an owner and a date — tracked like any other work

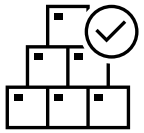

Linked to the inventory

the assessment lives one click from the system's row — not in a drive folder named 'final_v3'

Building Your
AI Governance Stack
The Human Layer



Roles: *names, not departments*



Every system has an owner

the inventory column made flesh: one name per system, accountable end to end



One program coordinator

part of a role, not a hire: keeps the inventory, the calendar and the training log alive and updated



A leadership sponsor

someone at the top who owns the risk appetite and unblocks decisions



HR wires it in

AI duties in job descriptions, onboarding and appraisals — so roles survive staff turnover

Oversight: *the four abilities*

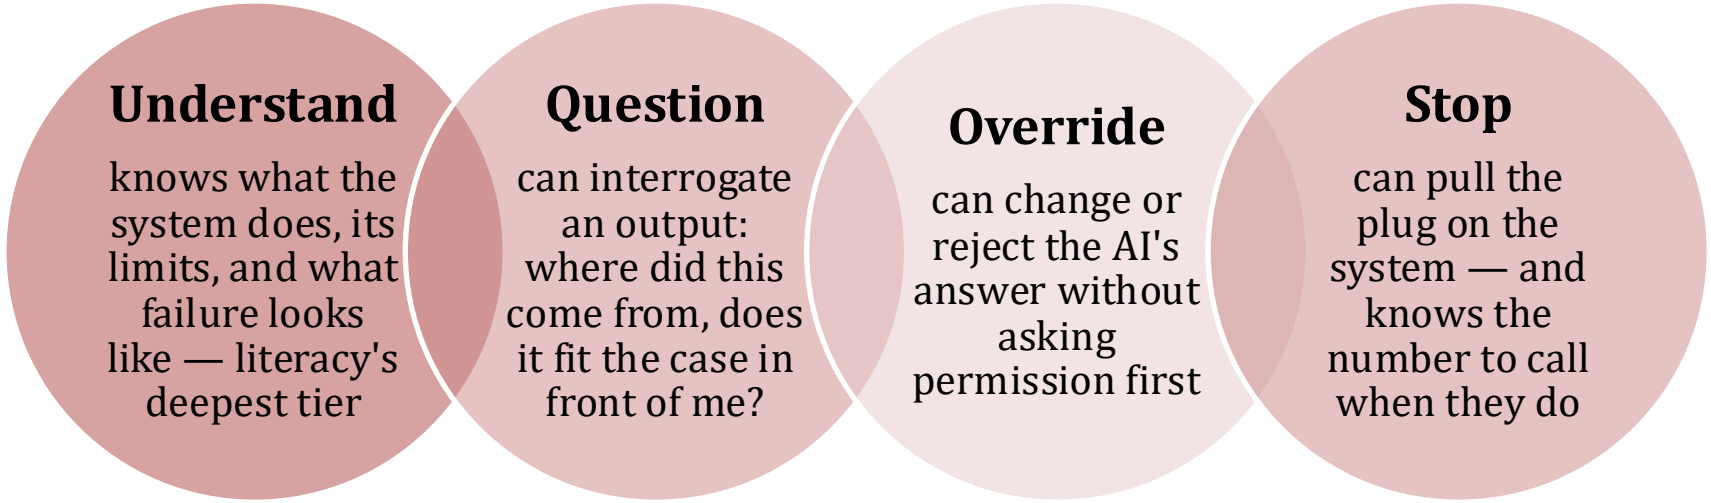

The rubber-stamp trap

An overloaded human clicking 'approve' 200 times a day is not oversight.

Give overseers time, authority, and a workload that allows actual looking.

Seven blocks, *one machine*

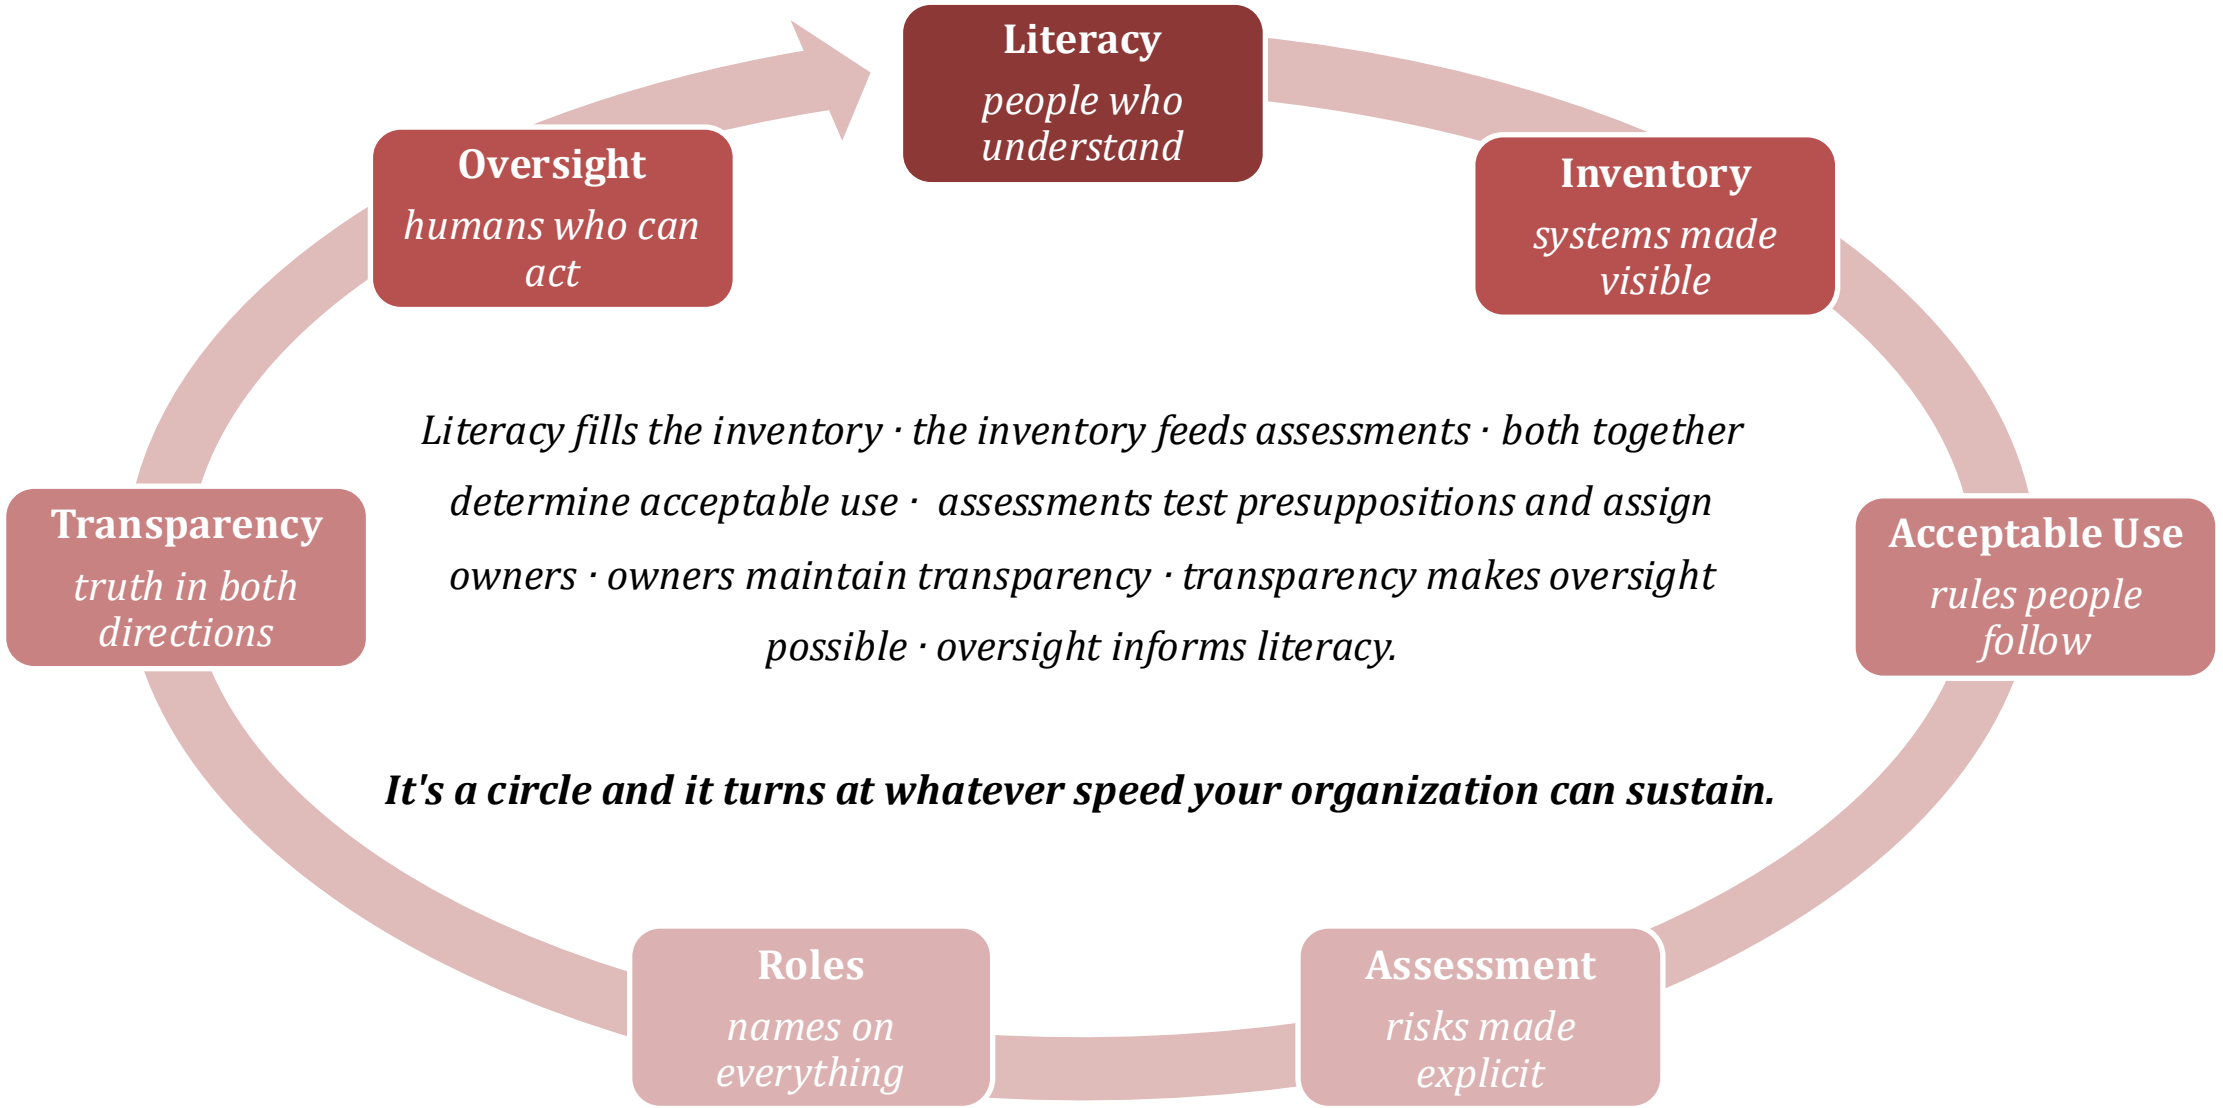

The AI Act calendar moved. *The exposure didn't.*

Applies now

- Article 5 prohibitions
- Article 4 AI literacy
- GPAI obligations
- Article 50 transparency
- **All of GDPR**

Deferred runway

2 December 2027

Annex III high-risk: employment, education, access to essential services, border management

2 August 2028

Annex I: AI embedded in regulated products

The AI Omnibus entered into force on 27 July 2026 and added an SME tier extending simplified documentation and sandbox access.

Three things *to take with you*

1

Add columns, rows and steps, not documents

Open the register you already have and add: AI system — yes or no — and who owns the output.

2

Put a fast lane on the front door

Three questions, two minutes, three lanes. Publish the answers.
This is what keeps people inside the system.

3

Set an oversight and accountability mechanism

For your three highest-exposure use cases. Names, triggers, and actions.

Contact

Tea Mustać

tea.mustac@upstreambytm.com

LinkedIn

Applied AI Governance



Your organisation is using AI. But are you prepared to manage the risks?
What can staff safely share with AI tools? How do you navigate GDPR, protect confidential data and intellectual property, recognise bias a AI mistakes, and ensure accountability?

Move from uncertainty to practical action.
Our Applied AI Governance course helps you and your organisation assess risks, develop internal polic an AI governance framework your staff will actually follow.

4 practical modules | 20 hours | 2 live expert sessions | Professional certificate

Course launch: 16 November 2026

ENROL IN THE PROGRAMME

