

Make AI Adoption Safe

A concise entry point to the J01 Knowledge Layer. It does not replace the Public Journey Guide.

Asset ID	EYW-JRN-J01-PUB-SUMMARY-001 v3.0
Journey	J01 Make AI Adoption Safe
Role	Public Summary (derived from the Public Journey Guide v3.0)
Classification	PUBLIC
Release profile	RP-PUBLIC
Publisher	Simon Owen / EmPOWER Your World
Date	26 September 2026
Controlling sources	Narrative v3.5.2 (SRC-015); legacy Public Journey Guide v1.0 (SRC-017)
Contract	Specification v3.3 §9.2 Summary contract; Narrative v3.5.2

The problem

J01 admits a specific problem. Adoption may appear active while the organisation cannot consistently show how use cases enter, how risk is sized, which controls apply, who owns them, how governance stays live, where evidence is held, or how incidents and surprises are handled.

Visible adoption and controlled adoption look the same until someone asks how a use case was sized, which controls it carries, who owns them, what would wake governance up and what proof exists.

The EYW point of view

Safe AI adoption is an organisational capability problem, not simply a policy, technology or compliance problem.

J01 enables a declared subject to establish and sustain a proportionate, evidence-based approach with uncertainty and authority boundaries visible. It does not promise universal safety, legal compliance or certainty.

The six stages

Stage	Question	Stronger practice in one line
1. Front Door	Where do AI use cases enter, and what bypasses that route?	One front door makes demand visible. The only gate decisions are Pass, Redirect and Refuse. Deferral pending a named prerequisite is not a fourth outcome.

Stage	Question	Stronger practice in one line
2. Risk Size	How big is this, and would two reviewers agree?	Gate no-gos, then size Agency, Reach, Data and Governability by pattern, retain dissent, apply controlled overrides and define reassessment triggers. M01 is not a J01 dependency.
3. Control Floor	Which controls are required, and who owns them?	Build the smallest sufficient floor, define material controls operationally, name owner, exception, escalation and proof.
4. Living Loop	What changes after launch, and what wakes governance up?	Sense, decide, act and leave a trace. Define expected signals, suspicious silence, triggers and authority.
5. Evidence Chain	If proof were requested tomorrow, where would it come from?	Capture decisions, controls, signals and outcomes at source. Custody is Create → Capture → Store → Retain → Retrieve → Defend.
6. Response Net	Who declares an AI surprise, and who owns the first decision?	Distinguish anticipated incidents from surprises, name authority beforehand, contain proportionately, preserve evidence and convert learning into change.

The order is deliberate. A use case is admitted and defined before it is sized; sizing frames proportionate controls; controls need live governance; operation must create defensible evidence; incidents and surprises feed learning back to governance and entry criteria.

What it does not provide

These materials are learning, self-service and working resources. They are not a formal benchmark, assurance review, certification, legal advice, regulatory advice or professional advice.

L0 is orientation and reflection only. It makes no capability-state claim, produces no score, and is not a diagnosis, benchmark or assurance.

Where to go next

Five neutral routes: Sustain, Deepen J01, Another Journey, No further action, Poor fit / redirect. Routing is a human-owned decision based on subject, scope, capability need, evidence, context, urgency, authority and alternatives. No score or weakest area selects the next product.

Read the full argument, examples and boundaries in the Public Journey Guide v3.0. PUBLIC fixed renditions may be retained, copied and presented internally with attribution. They confer no external republication right and no authority.