

Make AI Adoption Safe

The complete public Knowledge Layer of J01: six questions between an AI idea and a governed, evidenced use case

Asset ID	EYW-JRN-J01-PUB-GUIDE-001 v3.0
Journey	J01 Make AI Adoption Safe
Role	Public Journey Guide: canonical substantive public narrative
Classification	PUBLIC
Release profile	RP-PUBLIC
Publisher	Simon Owen / EmPOWER Your World
Date	26 September 2026
Controlling sources	Narrative v3.5.2 (SRC-015); PJP-01 (SRC-016); legacy Public Journey Guide v1.0 (SRC-017, migration source); Modules M15 v1.6, M10 v2.8, M11 v2.3, M12 v1.6, M13 v1.4, M14 v1.4 (SRC-006 to SRC-011)
Contract	Specification v3.3 §§9.1-9.2; PJP-01; Narrative v3.5.2

How to read this guide

This guide is the Knowledge Layer of Journey J01, published in full rather than trailed. It explains why AI adoption becomes unsafe even when activity looks healthy, sets out the EYW point of view and the six-stage Risk Spine, and works through each stage with its failure patterns, concepts, a controlled synthetic example, a reflection prompt and a picture of stronger practice.

What is not here is the operating system: instruments, calibration, working canvases, facilitation methods and assessment logic. Section 15 sets that boundary out plainly, and explains why it exists.

Sections 1 to 4 set up the problem and the model. Sections 5 to 10 take one stage each. Sections 11 to 16 help you interpret what you have read and decide, without pressure, what to do next. Reading time is about an hour end to end.

1. The trigger

AI is being used before anyone has agreed how it should be used. Ideas arrive from several directions at once, each plausible and each carrying a different amount of risk, and there is rarely a consistent line between the ones that can proceed quietly and the ones that need a conversation first.

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.

The person who feels it first is usually in risk, security or governance leadership. Typical triggers are regulation, an incident, an assurance gap or adoption that has already outrun its controls.

You will recognise this Journey if

- AI tools are already in use, and the list of where is longer than anyone could produce from memory.
- Use cases enter through inconsistent routes: some through architecture, some through procurement, some as a licence someone already bought, some never announced at all.
- Different teams make different risk decisions about work that carries similar risk.
- Controls are discussed before risk is understood.
- Evidence is assembled when someone asks for it rather than captured when the thing happened.
- If an AI system did something surprising on a Friday afternoon, nobody could say quickly who owns the first decision.

If none of that describes your organisation, because a defined front door, repeatable sizing, owned controls, live oversight and designed evidence are already in place, the honest advice is to sustain what you have. That is a legitimate outcome, and section 16 treats it as one.

2. The hidden problem

An organisation can look as if it is adopting AI well while being structurally unsafe. Activity is visible; governance is not. It helps to separate four states that are easy to confuse.

State	What it means	How it looks from outside
Visible adoption	People are using AI and talking about it; usage is growing.	Healthy, often celebrated.
Controlled adoption	Each use case has been sized, carries proportionate controls, and those controls have named owners.	Much the same as visible adoption, from outside.
Evidenced adoption	The controls leave a trace as work happens, so the organisation can show what operated and when.	Indistinguishable from the above until someone asks.
Sustainable capability	The organisation notices when something changes, knows who owns the call, and improves from what it learns.	Only distinguishable over time.

Visible adoption and controlled adoption look almost identical, from outside and from inside, until someone asks a question only controlled adoption can answer: how was this use case sized, which controls does it carry, who owns them, what would wake governance up, and what would we show if asked for proof?

Why documentation alone is not enough

Most organisations respond by writing something down: a policy, a set of principles, a responsible-AI statement. Documentation is necessary and it is not sufficient. It tends to fail in four recognisable ways.

Failure	What it looks like
Abstract	A policy that says “human oversight” cannot be operated; two reviewers reading it reach two different answers about the same use case.
Uncalibrated	One control set applied everywhere puts the same burden on a read-only summariser and on an agent that can act.

Failure	What it looks like
Moment-in-time	Heavy review at the gate, then nothing once the use case is live, while the model, data and integrations keep changing.
Unowned	A control named in policy with nobody holding its evidence, cadence or escalation.

A control nobody owns is a wish. A size label without controls has no consequence. A control checked once is already out of date.

3. The EYW point of view

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

Capability is present only when the declared subject can perform the work repeatedly in context, with authority, evidence, correction and learning visible.

The promise. 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 point of view, stated plainly so that you can disagree with it:

- **Size risk before choosing controls.** A control chosen before the size of the use case is known is either excessive or irrelevant. Proportionality is what makes governance survivable.
- **Controls must be owned and operable.** A material control has an owner, a trigger, a requirement, an exception route, an escalation path and proof that it operated. If it cannot be named, owned and evidenced, it is not yet a control.
- **Governance must stay alive after launch.** Models, data, providers, policies and people change continuously. Governance has to respond to the change, not only to the calendar.
- **Evidence is designed, not reconstructed.** Capturing decisions, controls, signals and outcomes at source is cheaper and more credible than rebuilding them when an audit is announced.
- **Surprise needs a declared owner.** Anticipated incidents have runbooks. Novel behaviour needs a named person who can declare a surprise and own the first decision.
- **The six parts form one dependent system.** Each stage hands the next something it needs, so sequence matters and a missing element weakens everything after it.

4. The Risk Spine

The Risk Spine is the six-stage sequence between an AI idea and a governed, evidenced, reviewable use case. Each stage answers one question and hands the next stage something it needs.

The Risk Spine: six stages between an AI idea and a governed, evidenced use case.

Step	Stage	Frozen Module source	The question it answers	Capabilities
1	Front Door	M15 Use-case Intake and Gates v1.6	Where do AI use cases enter, and what bypasses that route?	EYW-CAP-052 to 055
2	Risk Size	M10 AI Risk Sizing v2.8	How big is this, and would two reviewers agree?	EYW-CAP-033 to 036

Step	Stage	Frozen Module source	The question it answers	Capabilities
3	Control Floor	M11 AI Risk Controls v2.3	Which controls are required, and who owns them?	EYW-CAP-037 to 040
4	Living Loop	M12 Living Governance v1.6	What changes after launch, and what wakes governance up?	EYW-CAP-036 and 041 to 043
5	Evidence Chain	M13 AI Evidence & Audit-Readiness v1.4	If proof were requested tomorrow, where would it come from?	EYW-CAP-044 to 047 (linked to 040)
6	Response Net	M14 Incident & Surprise Response v1.4	Who declares an AI surprise, and who owns the first decision?	EYW-CAP-048 to 051

Why the sequence matters

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 a missing stage costs

Missing stage	What it looks like
Front Door	Everything downstream is optional. The organisation governs the use cases it happens to know about and discovers the others in production.
Risk Size	Controls are uniform or arbitrary: over-governed where risk is low and under-governed where it matters.
Control Floor	The size is a label with no consequence.
Living Loop	Governance describes the system as it was at launch; drift is invisible until it becomes an incident.
Evidence Chain	Good practice cannot be shown. Under assurance pressure, unprovable and untrue look the same.
Response Net	The first genuinely novel failure is handled by whoever happens to notice.

The spine runs both ways

The spine also runs backwards as a learning system. A response teaches the loop what to watch; the loop's signals become evidence; evidence shows which controls actually operate; control experience refines what a size requires; patterns in sizing show which use cases should not have passed the front door.

Each of the next six chapters follows the same shape: the question, why it matters, the common failure, the EYW point of view, the concepts, a controlled synthetic example, a question to think about, what stronger practice looks like, and the hand-off to the next stage. Worked examples are controlled synthetic examples. They illustrate the idea and claim no observed outcome.

5. Stage 1: Front Door

The question. Where do AI use cases enter, and what bypasses that route?

Why this question matters

Every later stage assumes a use case arrived through a gate, and most did not. The cost of a skipped gate climbs at each step: a short conversation at intake, a redesign at design, rework and retrofitted controls at build, and an incident or audit finding in production.

The common failure

Use cases are discovered after build, near-duplicates enter through different channels, sizing is used retrospectively to justify work, and approval cannot be reconstructed.

The EYW point of view

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.

The concepts

These ideas are owned by Module M15 Use-case Intake and Gates v1.6; J01 sequences and composes them.

- **One front door, no side doors** What is not declared cannot be sized, controlled or owned. Undeclared use is surfaced and placed at the right gate retrospectively rather than punished or ignored.
- **Four lenses at intake** Risk and regulatory; architecture and feasibility; value and business case; people and organisational impact. A clear pass needs all four, and any one lens can redirect or refuse.
- **Five gates, three outcomes** Intake, size, design, build and deploy, each with one question and the same three outcomes: Pass, Redirect or Refuse. “Park it and revisit” is a deferral, not a decision.
- **The gates own the decision** The rest of the spine owns the depth: sizing, controls, evidence and live oversight each belong to their own stage.

Controlled synthetic example

An undeclared call-summary assistant is registered retrospectively; ownership and data boundary are clarified and the authorised gate records Pass, Redirect or Refuse.

Illustrative only: no observed outcome is claimed.

A question to think about

Where do ideas enter, and what bypasses the route?

What stronger practice looks like

Recent requests trace to entry, acknowledgement and owner; use cases carry intended user, action, data, benefit, affected party, boundary and non-goal; gate decisions carry criteria, conditions and rationale.

How it connects onwards

A defined, triaged use case passes to Risk Size.

Capabilities	EYW-CAP-052 to 055
Module Card	C15 (supporting reference in the Module Cards pack; not a substitute for this chapter)
Public here	One-front-door concept, example, reflection and stronger practice.
Protected	Intake fields, gate criteria, records, instrument and facilitation.

6. Stage 2: Risk Size

The question. How big is this, and would two reviewers agree?

Why this question matters

Generic risk models struggle with AI because the system can act, its outputs can be hard to explain, one model can serve many uses so the risk sits in the use case, and a legal floor now exists.

The common failure

Teams size by instinct, average dimensions so the decisive risk disappears, or score a no-go into acceptance.

The EYW point of view

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.

The concepts

These ideas are owned by Module M10 AI Risk Sizing v2.8; J01 sequences and composes them.

- **Gate, then size, then respond** A no-go is stopped at the gate and never sized into acceptance. Only what passes is sized.
- **Pattern, not sum** The use case is read across four dimensions: Agency, Reach, Data and Governability. The dimension that sits highest sets the floor; averaging it away is the classic failure.
- **One effective size** Each use case resolves to one effective size, S, M, L or XL, with a flag where a general-purpose model is involved. A use the law already treats as high-risk is not sized down.
- **FIRM and FLEX** Some judgements may be debated on evidence; some may not. Knowing which is which keeps sizing honest.

Controlled synthetic example

A low-volume assistant can send messages. Reach is limited but Agency is high, so the high dimension is retained and human approval required.

Illustrative only: no observed outcome is claimed.

A question to think about

Would two reviewers agree, and where might they differ?

What stronger practice looks like

Context precedes sizing; judgements carry evidence and dissent; class is explainable and revisited by named authority when a trigger invalidates the earlier judgement.

How it connects onwards

Risk context, dimensions, class and trigger pass to Control Floor.

Capabilities	EYW-CAP-033 to 036
Module Card	C10 (supporting reference in the Module Cards pack; not a substitute for this chapter)
Public here	Dimensions, pattern-not-sum, dissent, example and reflection.
Protected	Instrument, thresholds, calibration, decision records and legal mapping.

7. Stage 3: Control Floor

The question. Which controls are required, and who owns them?

Why this question matters

A size is not a plan; controls are. Controls still fail when risk is sized correctly if nobody can operate them, they are the same everywhere, they stop at launch or nobody owns them.

The common failure

Controls become abstract, uncalibrated, moment-in-time or unowned, so a tier label is mistaken for an operating plan.

The EYW point of view

Build the smallest sufficient floor, define material controls operationally, name owner, exception, escalation and proof.

The concepts

These ideas are owned by Module M11 AI Risk Controls v2.3; J01 sequences and composes them.

- **A floor in four layers** A Baseline under every use case, Axis Modifiers for the dimension that sits high, Tier Deltas that stack as size climbs, and a GPAI Overlay where a general-purpose model is used. The floor is the union sized to this use case, not the sum of everything.
- **FIRM and FLEX decide who may decide** The baseline, a named human owner on every control and logged override authority are firm. Cadence, sampling and borderline placement flex on evidence, never below the firm floor.
- **Design once, apply to every use case** The standard is designed once across technology, process and people, then applied consistently.

Controlled synthetic example

The assistant gets data boundaries, human review, logging and a service owner; irrelevant model-development controls are excluded with rationale.

Illustrative only: no observed outcome is claimed.

A question to think about

Which controls operate and which are only policy?

What stronger practice looks like

The floor shows included and excluded controls; each control has trigger, requirement, operation, exception and escalation; owners accept proof duties.

How it connects onwards

Owned controls pass to Living Loop and Evidence Chain.

Capabilities	EYW-CAP-037 to 040
Module Card	C11 (supporting reference in the Module Cards pack; not a substitute for this chapter)
Public here	Control-floor concept, proportionality and ownership.
Protected	Catalogue, mapping method, cards, exception workflow and canvas.

8. Stage 4: Living Loop

The question. What changes after launch, and what wakes governance up?

Why this question matters

Governance was built for systems that hold still, and AI does not. Capability ships often, providers change the engine under the same name, data sensitivity changes, regulation moves and people change role.

The common failure

Governance follows a calendar while models, data, providers, policy and people change continuously, and silence is mistaken for health.

The EYW point of view

Sense, decide, act and leave a trace. Define expected signals, suspicious silence, triggers and authority.

The concepts

These ideas are owned by Module M12 Living Governance v1.6; J01 sequences and composes them.

- **Sense, decide, act, loop back** Catch each change across platform, model, data, policy and people; route the decision to the right hands; act with evidence; feed the trace into the next cycle.

- **Event, not calendar** A scheduled review confirms; the event is what catches. Silence is a signal: a sensor that has gone quiet is suspect, not healthy.
- **A governable signal** A noticed change becomes governable only when it has a named source, a defined shape, an expected cadence and a routed line.
- **Hold the human floor** Non-delegable accountability, a novel pattern or an irreversible impact always require a named human.

Controlled synthetic example

A provider model change routes to the owner before the calendar review; sizing and controls are revisited and the decision recorded.

Illustrative only: no observed outcome is claimed.

A question to think about

What change or missing signal wakes governance up?

What stronger practice looks like

Signals trace to authorised decisions, actions and updated records; review rationale, participants, inputs, thresholds and intervention routes are explicit.

How it connects onwards

Actions leave evidence; reassessment returns to Risk Size.

Capabilities	EYW-CAP-036 and 041 to 043
Module Card	C12 (supporting reference in the Module Cards pack; not a substitute for this chapter)
Public here	Loop, event-not-calendar, silence-as-signal, example.
Protected	Trigger library, rhythm, thresholds, workflow and templates.

9. Stage 5: Evidence Chain

The question. If proof were requested tomorrow, where would it come from?

Why this question matters

Most organisations treat AI evidence as an audit-time project, and by then it is too late. Evidence captured at the moment of decision is cheaper and more credible than evidence rebuilt later.

The common failure

Evidence is reconstructed at audit time from emails, screenshots and fragmented records with weak provenance and ownership.

The EYW point of view

Capture decisions, controls, signals and outcomes at source. Custody is Create → Capture → Store → Retain → Retrieve → Defend.

The concepts

These ideas are owned by Module M13 AI Evidence & Audit-Readiness v1.4; J01 sequences and composes them.

- **Four evidence streams** Decisions (who chose what, and why), controls (did the safeguards run), signals (what the system showed) and outcomes (what actually happened).
- **Capture at source** Records fall out of how the system already runs: at the gate, the control test, the signal and the incident.
- **A six-stage custody chain** Create, Capture, Store, Retain, Retrieve and Defend. The weakest stage sets how defensible the evidence is, and retention follows the size of the risk.

Controlled synthetic example

Gate, risk, owner, review and change records link by use-case ID; retrieval exposes known gaps honestly.
Illustrative only: no observed outcome is claimed.

A question to think about

Could the organisation retrieve and defend evidence without reconstruction?

What stronger practice looks like

One use case has linked records maintained during work; owners accept quality, correction, availability and expiry duties; assurance questions are answered by retrieval with limitations.

How it connects onwards

Evidence supports every stage and later assurance requests.

Capabilities	EYW-CAP-044 to 047 (linked to 040)
Module Card	C13 (supporting reference in the Module Cards pack; not a substitute for this chapter)
Public here	Chain, capture-at-source, limits, example.
Protected	Schemas, descriptors, custody controls, retention and prove-it method.

10. Stage 6: Response Net

The question. Who declares an AI surprise, and who owns the first decision?

Why this question matters

Playbooks were built for failures recognised in advance. AI can fail in shapes nobody scripted: novel outputs, behaviour that drifts with no code change, and causes that sit in prompts, data or the world.

The common failure

Novel behaviour is forced into the nearest familiar runbook, roles are stale and learning is not converted into owned change.

The EYW point of view

Distinguish anticipated incidents from surprises, name authority beforehand, contain proportionately, preserve evidence and convert learning into change.

The concepts

These ideas are owned by Module M14 Incident & Surprise Response v1.4; J01 sequences and composes them.

- **Two failure modes, one discipline** Incidents are events the organisation anticipated; surprises are events its playbook did not. Classify on the first decision: forcing a surprise into the nearest runbook is the costly mistake.
- **Name the Surprise Owner in advance** A named seat able to declare “this is a surprise” must exist before novelty arrives.
- **Every response feeds the loop** The lesson becomes a new control, signal or playbook entry, so the next occurrence is an incident rather than a surprise.

Controlled synthetic example

After an update, an assistant invents commitments; external use is paused, evidence preserved, scope confirmed and unrelated use stays live.

Illustrative only: no observed outcome is claimed.

A question to think about

Who declares surprise and owns the first decision?

What stronger practice looks like

Classes carry severity, rationale and revision rules; roles, contacts, scenarios, access and authority are rehearsed; learning becomes verified change.

How it connects onwards

Learning returns to Living Loop and may update entry criteria.

Capabilities	EYW-CAP-048 to 051
Module Card	C14 (supporting reference in the Module Cards pack; not a substitute for this chapter)
Public here	Incident/surprise distinction, authority, containment, example.
Protected	Playbooks, scenario library, escalation paths, simulation and contacts.

11. Subjects, capability states and confidence

These definitions let a reader interpret J01 language consistently. Every profile declares one subject, scope, context and time. Capability state is recorded per declared subject, scope, context and time. It is not a composite score.

Subject	Meaning	No silent transfer
Organisation	Operating roles, authority, process, information, technology, governance and relationships in declared scope.	Individual evidence does not prove organisation capability.
Practitioner	A consenting individual in a declared role, context and authority ceiling.	Organisation evidence does not accredit an individual.
Partner or delivery subject	A separately authorised delivery entity or practitioner under an explicit profile.	Experience or access does not grant delivery authority.

State	Controlled meaning
Unobserved	No relevant practice has been observed or described in the declared subject, scope, context and time. This is not proof of absence.
Emerging	Elements are attempted but incomplete, inconsistent, person-dependent or not repeatable.
Applied	Used in at least one relevant context with a bounded result; repeatability or coverage is not yet established.
Repeatable	Performed consistently across comparable contexts with known ownership, method and evidence.
Embedded	Integrated into normal work, survives reasonable change and reduced support, and is renewed through feedback and learning.
Insufficient Evidence	The evidence is inadequate for a state conclusion. It is not a low state.
Unable to Conclude	Conflict, unstable scope, inadequate authority or another limitation prevents a responsible conclusion.
Not Applicable	Outside the declared subject or scope, with rationale and approving authority recorded.

Evidence confidence is the trustworthiness and sufficiency of the evidential basis. Interpretation confidence is the strength with which material supports the reading. Neither converts automatically into capability state.

Profiles are primary. Every material finding has a factual correction and interpretation-challenge route. Role, seniority, purchase, attendance or completion are not evidence of capability.

Authority ceilings. L0 is orientation only and makes no capability-state claim. L1 is bounded subject self-report and makes no external assurance claim. L2 is evidence-reviewed contextual judgement only when separately authorised. L3 comparison is not active.

12. Reading the pattern

Having read the six chapters, you will have formed an impression of your own organisation. This section offers a way of thinking about that impression. It is educational, not an assurance assessment: there is no scoring, no benchmark, no comparison with other organisations and nothing that produces a rating.

- **Shared language.** If conversations keep stalling on what "governance", "control" or "oversight" mean, a shared vocabulary and a shared picture of how a use case travels from idea to production may matter more than any new control.
- **Consistency and ownership.** If the pieces exist but are used inconsistently, or are owned by functions rather than named people, the useful question is which of them would survive someone actually testing them.
- **Proof and learning.** If use cases are sized and controlled but the organisation could not quickly show what operated, or does not learn from what happens, the evidence and response stages deserve attention.
- **Uncertainty.** If you simply do not know how a stage works in practice, that is itself worth recording honestly. Not knowing is different from doing badly.

The six stages form a dependent system, so a gap in one stage affects the others. That observation helps a conversation; it does not select a route. Section 16 explains how a route is chosen.

A caution about self-reflection. One observer inside a system usually reports it as more consistent than it is. Comparing impressions independently with a colleague in a different function is often more informative than either view alone.

13. A reflection you can use

The L0 Journey Check turns the six stage questions, plus two questions about fit and action, into a short private reflection. It asks for no personal or organisational data, produces no score and makes no capability-state claim. You can use it alone in about twenty minutes or as a conversation with colleagues.

Stage	Reflection prompt
Front Door	Where do ideas enter, and what bypasses the route?
Risk Size	Would two reviewers agree, and where might they differ?
Control Floor	Which controls operate and which are only policy?
Living Loop	What change or missing signal wakes governance up?
Evidence Chain	Could the organisation retrieve and defend evidence without reconstruction?
Response Net	Who declares surprise and owns the first decision?

14. What this Journey changes

Stage	From	Towards
Front Door	Use cases arrive by whatever route they find.	One declared route, with three honest outcomes and a recorded reason.
Risk Size	Sizing by individual judgement, or after the build.	Repeatable sizing that two reviewers can explain, before controls are chosen.
Control Floor	Generic controls, or a size label with no consequence.	Controls proportionate to the size, each named, owned and evidenced.
Living Loop	Governance set at launch.	Governance that responds to the change, with silence treated as a signal.

Stage	From	Towards
Evidence Chain	Evidence assembled when someone asks.	Evidence designed and captured as work happens.
Response Net	Improvisation by whoever notices.	Declared decision and response routes, including for failures nobody predicted.

J01 reaches sufficient state when the declared subject can, at the authority level its configuration allows:

- receive AI demand through an owned and visible route;
- define a use case well enough for proportionate treatment;
- size risk across relevant dimensions and revisit judgement when context changes;
- select, define and own proportionate controls;
- operate a living governance loop with triggers and intervention routes;
- create, retain and retrieve trustworthy governance evidence;
- prepare for, classify, contain and learn from incidents and surprises;
- translate findings into a Capability Transition Plan and, where authorised, a Capability Development Plan;
- choose a neutral continuation route without forced commercial escalation.

Sufficiency is contextual: PUBLIC needs understanding and L0 orientation; REGISTERED adds bounded L1 self-report; SELF supports complete internal application; Authorised L2 remains deferred.

15. What this public product does and does not provide

This guide gives you understanding: the problem, the point of view, the six stages, the concepts, controlled examples, the questions worth asking and a picture of stronger practice. You can explain J01 to a colleague after reading it and recognise where it applies. That value is yours.

Public explanation may explain, recognise, reflect and orient. Some material is protected because it does a different job: diagnosing, deciding, implementing, facilitating, assessing, assuring or authorising. Applying a model inside a specific organisation needs instruments, calibration and judgement that a document cannot responsibly supply on its own. Protected material includes:

- Facilitator guides, session scripts and facilitation methods.
- Assessment instruments beyond L0, their interpretation guidance and any assessment logic.
- Working canvases and templates used to build a front door, sizing record, control floor, governance rhythm, evidence chain or response net.
- Calibration methods, thresholds and the control catalogue.
- Evidence descriptors, schemas and assessment protocols.
- Delivery, accreditation, partner or brand rights of any kind.

Launch configurations

Configuration	Receives	Explicitly withheld / ceiling
RP-PUBLIC	Complete Knowledge Layer plus L0 orientation.	Formal diagnosis, implementation methods, L1/L2 records, facilitation and assurance authority. L0 makes no capability-state claim.

Configuration	Receives	Explicitly withheld / ceiling
RP-REGISTERED	PUBLIC plus bounded L1 self-report, interpretation, uncertainty and correction.	External assurance, benchmarking, composite maturity score and delivery authority. L1 is self-report only.
RP-SELF	Knowledge Layer, internal activities, tools, transition/development/change/evidence records.	External commercial delivery, EYW-branded L2, accreditation and brand rights.
Deferred	Authorised L2, Community Workshop and later service/partner profiles.	No production until explicit activation. L2 is separately authorised evidence-reviewed judgement; L3 is inactive.

A release profile selects an audience-appropriate rendering; it does not redefine J01 or expose protected methods merely because this guide mentions them.

Boundaries

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.

- No employee-performance, redundancy, promotion or disciplinary use.
- No hidden surveillance or passive capability collection.
- No public person ranking, and no unsupported benchmark or causal claim.
- Private by default: learning contribution, anonymous use, named case study and testimonial each require separate permission.
- Purchase, registration, attendance, membership or completion do not confer capability, accreditation, assurance or delivery authority.
- Legal interpretations require specialist review. Synthetic examples claim no observed outcome.

16. Where to go next

There are five neutral routes out of J01. None of them is a default, and none is chosen by a score.

Route	When it fits	What to record
Sustain	Sufficient state is present and the main need is to maintain, monitor or renew it.	Record the persistence or review trigger.
Deepen J01	J01 remains the right problem frame and a material gap or uncertainty remains.	Record the gap, evidence and intended outcome.
Another Journey	A different admitted problem is materially more important or better matched.	State why. Do not default merely because J01 evidence is incomplete.
No further action	No proportionate action is justified now.	Record the reason, owner and reconsideration trigger.
Poor fit / redirect	J01 is not the right intervention, authority is missing, the need is outside EYW's boundary or a specialist route is required.	Record the safe 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.
- Routes remain valid even when deeper paid work is available; J01 findings stay owned.
- Safety, legal, privacy, security, regulatory or authority concerns pause product routing and trigger the appropriate specialist or governance route.
- Insufficient Evidence or Unable to Conclude is recorded before routing where applicable.
- Commercial preference cannot override product fit.

The separate Where to Go Next guide sets out these routes in more detail, including the no-sale outcomes.

About this guide

M15 and M10 to M14 retain intellectual ownership of their content. J01 owns problem framing, sequence, composition, hand-offs, audience experience and neutral route. Module onward references, including M27, M33 and M35, are Module content and are not a J01 route.

Guide section	Controlling source
1 to 4	Narrative v3.5.2 §§1-5; legacy Public Journey Guide v1.0 §§1-4 (migrated where consistent)
5 to 10	Narrative v3.5.2 §§6-11; Modules M15, M10, M11, M12, M13, M14 (concepts as curated in the controlled narrative)
11	Narrative v3.5.2 §§2, 12, 15
12 to 13	Narrative v3.5.2 §§6-11 reflection prompts; legacy guide §§11-12 (adapted: no scoring, no weakest-area routing)
14 to 16	Narrative v3.5.2 §§2-4, 14; legacy guide §§13-15 (adapted: neutral routes only)

Migration note: the legacy guide used a recurring fictional company for its worked examples. This edition uses the controlled synthetic examples from the canonical narrative so that one set of examples is used across the Journey.

Worked examples are controlled synthetic examples. They illustrate the idea and claim no observed outcome. PUBLIC fixed renditions may be retained, copied and presented internally with attribution. They confer no external republication right and no authority.