DOCUMENT 1 OF 2 · 14 PAGES

AI Pilot-to-Production Gating Workbook

Stage-gate every AI pilot before it goes live — explicit exit criteria, evidence required at each gate, named decision-makers, common failure modes and a sign-off block.

Audience · AI programme owner, sponsor, model risk, security, privacy, finance, business unit
Companion documents · Stage-Gate Evidence Index
Use as · Audit-grade artefact requiring named human sign-off
Version · v1.0 · 2026-05

§ 1

Stages overview & how this workbook is used

An AI initiative moves through four observable stages. Each stage has an explicit exit gate; the gate is a meeting, not a status field. The workbook is the canonical pack that travels into the meeting.

Stage	Purpose	Typical duration	Exit gate
S0 · Discovery	Problem worth solving; success criteria; constraints	2–4 weeks	G0 → fund a pilot
S1 · Pilot	Evidence the AI works on representative data; safety controls drafted	6–12 weeks	G1 → fund limited production
S2 · Limited production	Live use behind a canary / shadow; full operational controls; measured outcomes	8–16 weeks	G2 → general availability
S3 · General availability	Full population in scope; routine operating cadence	Steady state	G3 → annual review / retire

> **Reading rule.** Every gate page lists exit criteria, evidence required and decision rights. A "yes" without all three is not a gate decision — it is a future audit finding.

§ 2

G0 · Exit criteria from Discovery

Mandatory exit criteria

1. Problem statement with quantified scope (volume, $ value, populations affected).
2. Strategic-fit articulation tied to a named business outcome.
3. Hypothesised AI solution + at least one credible non-AI alternative considered.
4. Data sources identified with lawful basis confirmed by privacy.
5. Initial risk classification (EU AI Act tier; sector-specific equivalents) recorded.
6. Named sponsor, product owner and architecture reviewer.
7. Pilot success criteria (primary metric, MDE, decision rule) pre-registered.
8. Pilot budget & timeline approved.

Evidence required (drop into G0 pack)

- Discovery brief (1–2 pages)
- Use-case risk classification questionnaire (companion EU AI Act doc)
- Privacy preliminary screen + lawful basis statement
- Pilot plan with success criteria and budget

Decision rights. Sponsor approves; programme owner records the decision; absentees do not block but cannot reverse later.

§ 3

G1 · Exit criteria from Pilot

1. Primary metric met or exceeded with pre-registered MDE; CI excludes the null.
2. Subgroup performance reported; no group materially worse without mitigation.
3. Validation report drafted (companion sector pack, e.g. SR 11-7, clinical, industrial).
4. Risk register populated; top-12 risks have named owners and mitigations.
5. Security review of inference path completed; no critical findings open.
6. DPIA (if PII / PHI) drafted; lawful basis confirmed; consent path designed.
7. Human-in-the-loop contract written; automation boundary explicit.
8. Cost of run modelled; ROI working hypothesis defensible to finance.
9. Rollback path defined; canary & kill-switch designed.
10. Limited-production launch plan with target population, channel and exit definition.

Evidence required (drop into G1 pack)

- Pilot results report (with subgroup tables)
- Draft validation report (sector pack)
- Risk register
- Security review attestation
- Draft DPIA
- HITL contract
- Cost model and ROI worksheet
- Rollback design

§ 4

G2 · Exit criteria from Limited Production

1. Live performance on the live population matches or exceeds pilot (with documented variance).
2. Operational SLOs met for the gate-defined window (latency, error, freshness).
3. Permanent holdout in place and reconciles with reported lift.
4. Monitoring contract live; alerts firing into a named on-call; runbooks tested.
5. Drift / fairness guardrails clean for the window; any breaches investigated and closed.
6. Validation report signed by 2nd line; findings rated; remediation on track.
7. DPIA finalised and signed; consent / opt-out path verified end-to-end.
8. Incident-response runbook tested in a tabletop or drill.
9. Cost-per-decision tracked; unit economics defensible.
10. Communications & training complete for all human operators in the loop.

Evidence required (drop into G2 pack)

- Limited-production outcomes report (live data, with holdout reconciliation)
- Monitoring & incident report for the gate window
- Signed validation report + findings tracker
- Signed DPIA
- Incident-runbook drill record
- Unit-economics report
- Training & change-management record

§ 5

G3 · Annual review (GA running)

1. Annual performance, fairness and drift report against original gates.
2. Annual privacy review: lawful basis, retention, sub-processors, transfers.
3. Annual security review: certifications, pen-test, supply chain.
4. Model card refreshed; limitations register reviewed and re-signed.
5. Risk register reviewed; top-12 still valid; controls effective.
6. Cost & benefit reconciliation against the pre-launch business case.
7. Decision: continue · expand · constrain · retire.

Evidence required (drop into G3 pack)

- Annual performance & fairness report
- Annual privacy review
- Annual security review
- Refreshed model card
- Risk-register review
- Realised-benefit report

Anti-pattern. A "GA running" surface that has never received a G3 review is a regulator's favourite finding. Schedule G3 by date on the day GA is granted.

§ 6

Gate decision template (one page per gate)

Field	Detail
Initiative	
Gate	G0 · G1 · G2 · G3
Date / chair	
Attendees	
Evidence pack ref	
Exit criteria status	List each criterion · met / partial / unmet
Open risks & findings	Critical · High · Medium counts
Conditions (if any)	Action · owner · due
Decision	Approve · Approve with conditions · Defer · Reject
Rationale	
Dissent recorded	

§ 7

Gating committee roles & RACI

Role	G0	G1	G2	G3
Sponsor	A	A	A	A
Programme owner	R	R	R	R
Product owner	C	R	R	C
Model risk / 2nd line	I	A (validation)	A	A
Privacy / DPO	A (lawful basis)	A (DPIA)	A	A
Security	C	A	A	A
Finance	C	C	C	A (realisation)
Internal audit	I	I	C	C
Operating sponsor (business)	C	R	R	R

A = Approves · R = Responsible · C = Consulted · I = Informed. Decisions require the named A; conditions can be raised by any R or C.

§ 8

Common failure modes & how to avoid them

Failure mode	Why it happens	Counter
Pilot success without holdout	Convenience sampling; comparison against optimistic baseline	Holdout pre-declared at G0; reconciliation required at G2
"It works in offline AUC but not live"	Offline / online skew; population shift	Shadow + canary mandatory; parity tests in CI
Quiet promotion past a gate	Deadline pressure; informal decisions	Gate is a meeting; decision template signed
Validation arrives after launch	Validator independence late; "we'll get to it"	2nd-line validation as G1 exit; no exceptions
HITL on paper, not in interface	UX bypasses the control	HITL contract verified in the live UX before G2
Cost surprises at scale	Unit-economics not measured in pilot	Cost-per-decision tracked from G1
Retirement that never happens	Inertia; no one owns sunset	G3 owner explicitly empowered to retire

§ 9

Example portfolio view (one-line per initiative)

Initiative	Stage	Risk tier	Last gate	Next gate	Open Crit/ High	Owner
Retail personalisation	S2	Limited	G1 · 2026-02	G2 · 2026-05	0 / 2	
Credit scoring v3	S2	High (SR 11-7 Tier 1)	G1 · 2026-01	G2 · 2026-04	0 / 1	
Triage copilot	S1	High (FDA SaMD)	G0 · 2025-12	G1 · 2026-03	0 / 0	
Predictive maintenance	S3	Limited	G2 · 2025-09	G3 · 2026-09	0 / 0	
Marketing copy assistant	S1	Limited (GPAI)	G0 · 2026-02	G1 · 2026-04	0 / 0	

A weekly programme stand-up reviews this one-pager. Items where Critical > 0 page the programme owner regardless of stage.

§ 10

GenAI / agentic addendum

GenAI and agentic systems introduce additional checks at each gate.

Check	G0	G1	G2
Prompt & system-message registry	Plan	Live	Audited
Retrieval grounding policy (no hallucinated citations)	—	Tested	Monitored
Refusal / unsafe-content tests	—	Sampled	Continuous
PII / secret leakage filters	—	Tested	Monitored
Action permissioning (tools the agent can call)	—	Documented	Audited per call
Cost cap per session / day / agent	—	Live	Live
Human-approval scope (what cannot be autonomous)	Drafted	Live	Audited

Default off. Autonomous customer-affecting decisions (credit, clinical, safety, employment) are off until explicit written sign-off from sponsor + risk + privacy. Re-enabling requires a re-G2 review.

§ 11

Vendor / third-party AI gating addendum

- Vendor scorecard complete at G0 (see Foundation-Model Vendor Scorecard in companion sector pack).
- BAA / DPA / SCCs / right-to-audit executed by G1.
- Sub-processor inventory complete by G1; change-notification clause active.
- Independent validation on bank / hospital / plant data at G1; vendor benchmarks alone do not satisfy.
- Exit plan documented at G2 (data return, fine-tuned model export, transition assistance).
- Annual re-score and contract review at G3.

§ 12

Tabletop drill scenarios (use to test the gates)

Scenario	Stage tested	Goal
"Pilot looks great, sponsor wants to skip G1"	G1	Test that the gate holds under pressure
"Canary monitor breaches at 03:00"	G2 operations	Rollback in < 10 min
"Subgroup performance regression detected post-launch"	G2 / G3	Containment + re-validation
"Vendor announces model deprecation in 30 days"	G3	Exit / re-validation path
"Regulator requests pre-launch evidence pack"	Any	5-business-day SLA met

§ 13

Programme-level cadence

Cadence	Forum	Inputs
Weekly	Programme stand-up	Portfolio one-pager; open Critical / High
Monthly	Steering committee	Gate decisions due this month; incidents
Quarterly	Sponsor + risk review	Realisation vs business case; portfolio health
Annual	Board / audit-committee	Companion briefing pack

§ 14

Sign-off & attestation

Sign-off block

Role	Name	Date	Signature
Programme sponsor			
Programme owner			
Model risk lead			
Privacy / DPO			
Security			
Internal audit			

This workbook is reaffirmed annually or on any material change to the gating policy. Initiative-level gate decisions are recorded against this workbook's version.