DECISIONFLOW

Launch Risk, Incident and Feedback Register

Governance for launch risk, production incidents, feedback, and product decisions

Document code	LAUNCH-05
Version	2.0
Status	Final Public Portfolio Edition — Simulated Launch Scenario
Baseline date	14 July 2026
Author	Min Thu Kyaw
Role	AI Product Manager
Evidence rule	Actual artifacts are identified separately from assumed capabilities and simulated outcomes.

EVIDENCE & PUBLISHING INTEGRITY

This public portfolio edition distinguishes actual evidence, assumed continuity steps, simulated case-study outcomes, targets, and recommendations. Simulated customers, commercial results, benchmarks, incidents, quotations, and later-lifecycle outcomes are illustrative and must not be interpreted as results from a live commercial product. Evidence labels are retained throughout the package to support transparent and responsible interpretation.

Document Control

Control	Definition
Purpose	Document the Launch phase for Launch Risk, Incident and Feedback Register.
Primary audience	Hiring managers, product leaders, launch stakeholders, engineering, customer success, and operations.
Evidence classes	Actual artifact / assumed post-development capability / simulated launch outcome / recommendation.
Scenario assumption	All development P0 release gates are assumed complete before the simulated controlled launch begins.
Ethical use	Do not present simulated customer evidence, revenue, testimonials, or performance figures as real-world results.
Change control	Publication rule: retain evidence labels unless content is replaced by documented real-world observations.

Actual Project Artifacts

- Low-Fidelity Prototype
- High-Fidelity Prototype
- MVP Frontend Repository
- MVP Backend Repository

Contents

Risk Governance

The register connects known launch risks, observed incidents, customer feedback, and product decisions. A risk is a possible future event; an incident is an event that occurred; feedback is evidence about a customer problem or response; and a decision records what the team chose and why.

Evidence classification - Actual: public prototypes, repository references, uploaded MVP development baseline, and prior Product/Plan/Develop artifacts. Assumed: the production-ready v1.0 core capabilities documented in DEV-06 as completed for the launch scenario. Simulated: user behavior, commercial results, AI benchmarks, incidents, quotations, and launch outcomes.

Pre-Launch Risk Register

ID	Risk	Likelihood	Impact	Mitigation	Trigger / owner
LR-01	AI extracts or recommends unsupported information.	Medium	High	Human approval, evidence links, schema validation, evaluation, visible fallback.	Hallucination >5%; AI owner.
LR-02	Customers upload content they are not authorized to process.	Medium	High	Disclosure, consent, admin policy, deletion, least privilege.	Privacy complaint; Security/CS.
LR-03	Launch traffic or retries increase AI cost unexpectedly.	Medium	Medium	Budgets, rate limits, idempotency, routing, cost alerts.	Cost >$0.60/finalized decision; Product/AI.
LR-04	New users fail before first finalization.	Medium	High	Guided onboarding, templates, clear errors, office hours.	Activation <70%; Product/CS.
LR-05	The product remains a one-time decision tool rather than a recurring habit.	Medium	High	Second-decision prompts, retrieval use cases, templates, and success plan.	Seven-day second-decision rate below target; Product.
LR-06	Provider outage blocks core workflow.	Medium	High	Retry policy, queue, status, optional labeled fallback, provider abstraction.	Failure >5%; AI/Platform.
LR-07	Authorization or tenant isolation defect exposes data.	Low	Critical	Ownership tests, security review, monitoring, incident plan.	Any confirmed access breach; Security.
LR-08	Finalized record is incomplete or inconsistent.	Low	Critical	Immutable snapshot, validation, audit, E2E tests.	Any missing selected outcome/version; Engineering.
LR-09	Customer expectations exceed supported integration scope.	High	Medium	Clear scope, known limitations, roadmap governance.	Repeated integration blocker; Product/GTM.
LR-10	Simulated case-study results are mistaken for real traction.	Medium	High	Disclosures on every document, portfolio page, and presentation.	External reuse without label; Author.

Incident Process

1. Detect and create an incident record with time, affected capability, customer scope, and evidence.

2. Assign severity, incident commander, technical owner, communications owner, and scribe.

3. Protect customers first: stop rollout, disable feature, isolate provider, or rollback as appropriate.

4. Communicate known facts, impact, workaround, and next update time without speculation.

5. Resolve the immediate incident and verify customer data, workflow, and monitoring recovery.

6. Complete root-cause analysis, contributing factors, corrective actions, owners, and due dates.

7. Link the incident to risks, feedback, release notes, and product decisions.

Illustrative Incident Log

Simulated incidents - The incident records below are invented for the launch case study. They demonstrate governance and learning, not actual production events.

ID	Week	Severity	Incident	Impact	Immediate action	Outcome
INC-001	W1	SEV-2	AI provider timeout spike	17 requests delayed; no data loss.	Paused cohort, increased timeout visibility, provider retry adjustment.	Resolved in 48 min; add provider health dashboard.
INC-002	W1	SEV-3	Password-reset email delay	11 reset emails delayed.	Manual support and queue inspection.	Resolved; email queue alert added.
INC-003	W2	SEV-2	Duplicate finalization after browser retry	4 duplicate audit events; final record remained consistent.	Disabled duplicate action with idempotency key.	Resolved in 2h; E2E regression added.
INC-004	W3	SEV-3	Missing translation labels in Myanmar locale	Settings page showed English fallback.	Used fallback and patched translation keys.	Resolved next release; localization CI added.
INC-005	W5	SEV-3	Unsupported PDF expectation created onboarding friction	Two customers converted files to TXT.	Clarified scope and added conversion guidance.	Resolved in copy; parser remains roadmap.
INC-006	W7	SEV-2	Decision-library query slowdown	Recent decisions loaded slowly for 9 minutes.	Scaled database resources and added latency metric.	Resolved in 76 min; capacity threshold revised.

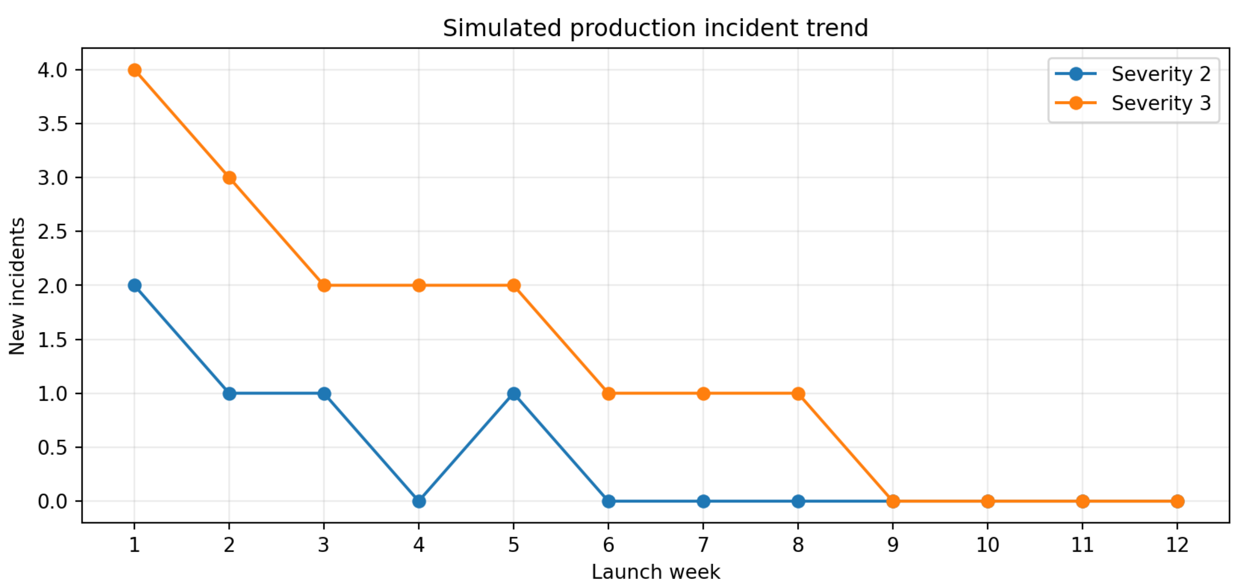

Simulated incidents decline as launch controls and fixes mature.

Customer Feedback Register

ID	Source	Evidence / request	Impact	Product response	Disposition
FB-001	Product Manager	Need clearer evidence before recommendation	High	Increase evidence-first layout and citation visibility.	Prioritized W4 experiment.

ID	Source	Evidence / request	Impact	Product response	Disposition
FB-002	Engineering Manager	Jira/Linear integration required for team workflow	High	Validate integration demand and design action export.	Roadmap next quarter.
FB-003	Founder	PDF/DOCX should work without conversion	High	Release secure parsing after fixture and privacy validation.	Controlled rollout.
FB-004	Product Manager	Owner suggestions are useful but occasionally wrong	High	Improve speaker/context model and require confirmation.	AI evaluation workstream.
FB-005	Operations Lead	Need reusable decision templates	Medium	Test templates for vendor, hiring, and process decisions.	Experiment EXP-05.
FB-006	Product lead	Want shared review, comments, and approval status	Medium	Improve collaboration and review history.	P1 roadmap.
FB-007	Executive	Need cross-project dashboard and export	Medium	Define executive reporting without vanity metrics.	Discovery required.
FB-008	Security reviewer	Need organization-level retention and provider controls	High	Add admin policy and contract documentation.	Business/enterprise package.
FB-009	PM	Mobile is not important for review work	Low	Keep desktop-first; support responsive retrieval only.	Deprioritized.
FB-010	Engineering Manager	Decision search is more valuable after several weeks	High	Measure retrieval and onboarding value over time.	Retention analysis.

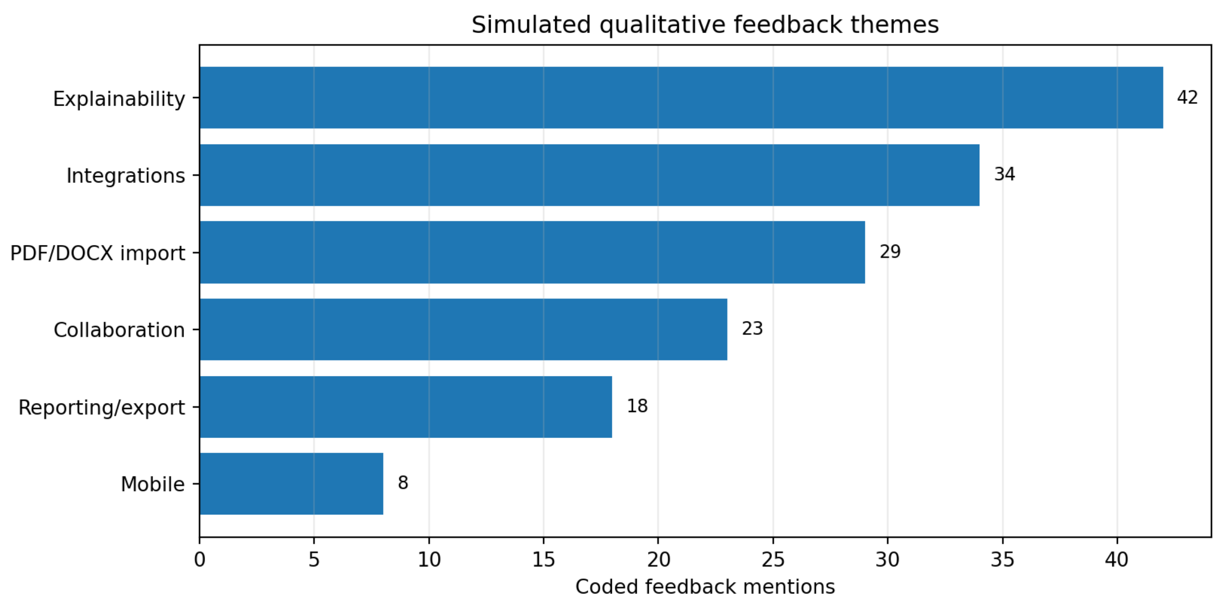

Simulated coded feedback themes from interviews, support, and workflow reviews.

Escalation Matrix

Issue type	Primary owner	Escalate when	Executive / customer action
Security/privacy	Security owner	Potential unauthorized access, disclosure, retention breach, or provider violation.	Stop affected processing; notify according to policy.
Data integrity	Engineering lead	Finalized outcome, audit history, or ownership is inconsistent.	Pause affected workflow and verify records.

Issue type	Primary owner	Escalate when	Executive / customer action
AI quality	AI owner	Hallucination, rejection, or owner/date error exceeds threshold.	Pause model/prompt promotion and sample production outputs.
Availability/performance	Platform owner	Critical threshold is sustained or core launch accounts blocked.	Pause cohort and communicate status.
Customer adoption	Product + CS	Activation or retention misses gate for two reviews.	Run intervention; reconsider ICP/workflow.
Commercial	GTM owner	Qualified conversion or procurement cycle misses gate.	Review positioning, package, pricing, and proof.

Launch Decision Register

ID	Decision	Rationale	Status / review trigger
LD-001	Use controlled cohorts rather than unrestricted public signup.	Protects operations and creates higher-quality learning.	Approved; review after Week 12.
LD-002	Keep explicit user finalization mandatory.	Decision authority and record quality outweigh automation speed.	Approved; permanent principle.
LD-003	Publish AI provider and fallback provenance.	Trust requires users to know how output was produced.	Approved; release gate.
LD-004	Use Starter and Professional individual packages as the v1.0 commercial motion.	The available product is a personal decision workflow; team packaging waits for collaboration evidence.	Approved; review after post-launch discovery.
LD-005	Prioritize explainability, repeat use, and validated collaboration/integration discovery over mobile.	Feedback and usage indicate desktop decision work is more valuable.	Illustrative Week-12 decision.
LD-006	Scale with safeguards after Week 12.	Core value and quality gates pass, but retention and selected AI fields need improvement.	Illustrative decision.

Risk Review Cadence

- Daily during the first two launch weeks for active incidents and critical leading indicators.
- Weekly for risk score changes, corrective actions, feedback themes, and customer account health.
- After every SEV-1 or SEV-2 incident for root cause and systemic risk updates.
- At each release for regression, migration, security, AI, and support readiness.
- At the Week-12 phase gate for residual risk acceptance and scale conditions.