

DECISIONFLOW

Release Readiness and Launch Runbook

Production release, monitoring, incident response, rollback, and ownership

Document code	LAUNCH-02
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 Release Readiness and Launch Runbook.
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

1. Release principles
2. Readiness gates
3. Environment and configuration model
4. Pre-launch timeline
5. Deployment procedure
6. Smoke tests
7. Monitoring and alerting
8. Incident command
9. Rollback and recovery
10. Launch-day operating model
11. Post-launch stabilization
12. Release sign-off

Release Principles

- Protect customer data and decision integrity before optimizing launch speed.
- A reversible, observable release is preferred to a large irreversible launch.
- Database migrations, AI prompts, model versions, and configuration are release-controlled artifacts.
- No silent AI fallback, automatic finalization, or ambiguous scoring is permitted.
- The launch owner may stop or roll back the release when a gate is breached.

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.

Readiness Gates

Area	Required evidence	Approver
Product	All P0 scope accepted; finalization, resume, review persistence, personal history, and scope boundaries validated.	Product + Engineering
AI	Evaluation thresholds met; fallback and provenance visible; prompt/model versions recorded.	AI owner + Product
Quality	CI green; migration, integration, E2E, accessibility, browser, and smoke tests passed.	QA + Engineering
Security and privacy	Secrets, sessions, rate limits, authorization, retention, deletion, and provider disclosure approved.	Security / Technical owner
Operations	Staging parity, backups, restore drill, monitoring, alerts, runbooks, and incident owners ready.	Platform owner
Customer	Onboarding, support channel, known limitations, privacy notice, and escalation expectations ready.	Customer success

Environment Model

Environment	Purpose	Data policy	Deployment rule
Local	Developer iteration	Synthetic or approved local data	No production secrets.
CI/Test	Automated quality gates	Synthetic fixtures; mocked AI by default	Every merge runs required checks.
Staging	Production-like release candidate	Synthetic or consented test data; capped AI budget	Same migration and build path as production.
Production	Controlled launch	Customer data under published retention and access policy	Approved change, monitored rollout, rollback available.

Pre-Launch Timeline

Time	Required activities	Exit / evidence
T-14 days	Scope freeze, dependency review, capacity test, security review, support rehearsal, customer list finalization.	Launch blocker register opened.
T-7 days	Production-like rehearsal, migration test, backup/restore drill, E2E run, status-page check, comms approval.	Release candidate approved.
T-3 days	Finalize feature flags, rate limits, AI budgets, dashboards, alert thresholds, owners, contact tree.	War-room checklist signed.
T-1 day	Database backup, production health review, customer notice, deployment package checksum, rollback snapshot.	Go/no-go meeting complete.
T0	Deploy in controlled window, run smoke tests, enable cohorts, observe first workflows.	All critical checks green.
T+1 day	Review incidents, funnel, latency, AI quality, customer tickets, and costs.	Continue / pause / rollback decision.
T+7 days	Stabilization review and first adoption analysis.	Prioritized fixes and customer actions.
T+30 days	Commercial and retention review.	Launch plan updated.

Deployment Procedure

1. Confirm approved release commit, changelog, migration revision, prompt version, model configuration, and feature flags.
2. Verify environment variables and reject default or missing production secrets.
3. Create and verify a restorable production database backup.
4. Deploy backend release with migrations using the approved pre-deploy process.
5. Verify health endpoint, database connectivity, queues, AI provider connectivity, and authorization checks.
6. Deploy frontend release and verify the production API endpoint, content security policy, analytics, and error reporting.
7. Run automated and manual smoke tests using non-sensitive controlled data.
8. Enable the first launch cohort through feature flag or account allowlist.
9. Monitor errors, latency, cost, workflow completion, and support channels for the defined observation period.
10. Record release outcome and either continue rollout, pause, or execute rollback.

Production Smoke Test

ID	Test	Expected result	Stop condition
SM-01	Register and authenticate a test account.	Session created; protected routes and API calls work.	Authentication or isolation failure.
SM-02	Create a second account and verify resource isolation.	Each account accesses only its own records.	Any unauthorized cross-account access.
SM-03	Upload/paste approved test source.	Validation, processing, provenance, and results	Data loss, unsupported silent

ID	Test	Expected result	Stop condition
		complete.	failure, or unlabelled fallback.
SM-04	Edit extraction and reload.	All edits persist and drive later recommendation.	Reviewed data not durable.
SM-05	Generate, regenerate, compare, and finalize.	Versions are stored; selected outcome and immutable final snapshot exist.	Incorrect or duplicate finalization.
SM-06	Search and reopen final decision.	History, evidence, actions, and final outcome are retrievable.	Record inconsistency or missing audit information.
SM-07	Attempt cross-account access.	Access denied for every protected resource.	Any data isolation breach.
SM-08	Archive/delete under policy.	Customer-visible result matches retention and audit rules.	Unrecoverable unexpected deletion.

Monitoring and Alerts

Signal	Warning	Critical	Owner / response
API error rate	> 1% for 10 min	> 3% for 5 min	Engineering investigates; pause rollout if customer workflow affected.
p95 non-AI latency	> 700 ms	> 1,500 ms	Platform scales or rolls back recent change.
p95 AI latency	> 35 sec	> 60 sec	AI owner checks provider, queue, budget, and fallback behavior.
AI failure rate	> 2%	> 5%	Stop new cohort activation; investigate provider and prompts.
Authorization failures	Unexpected pattern	Any confirmed cross-tenant access	Security incident; disable affected path immediately.
Finalization integrity	Any duplicate warning	Any missing/incorrect final snapshot	Pause finalization and preserve evidence.
Cost per finalized decision	> \$0.60	> \$1.00	Review routing, prompts, retries, and abuse.
Support volume	> 8 tickets / 100 active users/day	> 15 / 100	Increase support and identify systemic issue.

Incident Command

Severity	Definition	Initial response	Communication
SEV-1	Security breach, cross-customer data exposure, irreversible data loss, widespread inability to finalize.	Immediate incident commander; stop rollout; preserve evidence; rollback/disable.	Executive, affected customers, legal/security as required.
SEV-2	Major workflow unavailable or high error rate with workaround/limited scope.	Owner within 15 minutes; mitigation target within 60 minutes.	Status page and affected customers.
SEV-3	Degraded experience or isolated defect without material data risk.	Triage same business day; planned fix.	In-product/support communication as needed.
SEV-4	Cosmetic, documentation, or low-impact issue.	Backlog and scheduled correction.	No broad communication required.

Rollback Decision Rules

- Rollback immediately for confirmed authorization breach, data corruption, irreversible finalization error, or secret exposure.
- Pause cohort expansion for sustained critical AI failure, processing latency, duplicate workflows, or support overload.
- Prefer feature-flag disablement when the defect is isolated and data integrity remains intact.
- Database rollback must follow tested migration compatibility; do not reverse a destructive migration without a verified recovery plan.
- Every rollback creates a decision-log entry with trigger, evidence, owner, customer impact, and follow-up actions.

Launch-Day Operating Model

Role	Responsibility during launch
Launch owner	Go/no-go authority, cohort control, stakeholder updates, and final decision.
Engineering lead	Deployment, defects, rollback, and technical coordination.
AI owner	Provider health, AI quality, prompt/model behavior, cost, and fallback.
Platform owner	Infrastructure, monitoring, capacity, backups, and alerts.
Customer success lead	Customer onboarding, support triage, known limitations, and feedback.
Communications owner	Customer notices, status updates, release notes, and internal reporting.
Scribe	Timeline, decisions, incidents, metrics, and unresolved actions.

Release Sign-Off

Gate	Status in illustrative scenario	Evidence owner
Product acceptance	Passed / assumed	Product
AI evaluation	Passed / simulated benchmark	AI + Product
Quality and regression	Passed / assumed	QA + Engineering
Security and privacy	Passed / assumed	Security owner
Operations and rollback	Passed / assumed	Platform owner
Customer readiness	Passed / assumed	Customer success

Portfolio rule - A real release runbook must be completed with named owners, real environments, actual test evidence, approved privacy terms, and verified rollback steps. The scenario statuses above are illustrative.