

DECISIONFLOW

Launch Phase Overview

Controlled commercial launch and illustrative outcome summary

Document code	LAUNCH-00
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 Phase Overview.
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. Executive summary
2. Launch assumptions and evidence classification
3. Product at launch
4. Launch objectives and decision gates
5. Target market and positioning
6. Launch model and timeline
7. Launch readiness summary
8. Illustrative outcome dashboard
9. Launch recommendation
10. Document set map

Executive Summary

DecisionFlow is presented in this portfolio scenario as a production-ready v1.0 core decision intelligence product launched through a controlled commercial release. The launch scope is intentionally narrower than the long-term organizational-memory vision. The launch strategy prioritizes evidence, trust, and operational safety over a high-volume public release. The product helps knowledge workers transform meeting notes or transcripts into structured decisions, evaluate options, review explainable AI recommendations, and preserve a durable decision record.

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.

Launch Thesis

Launch thesis - A decision-first AI workflow can earn repeat usage when it shortens the path from ambiguous discussion to a trusted, finalized decision while preserving explicit human authority.

Product at Launch

Capability	Assumed launch state	Customer value
Decision intake	Pasted text and UTF-8 TXT supported with secure validation	Customers can start from approved text-based meeting artifacts.
Structured extraction	Decision topic, options, criteria, stakeholders, risks, actions, and advisory owner/date suggestions when present	Reduces manual structuring while preserving confirmation for uncertain fields.
Human review	Every AI-generated field is editable and requires approval	Maintains control and creates trusted records.
Recommendation and comparison	Explainable recommendation plus transparent weighted comparison	Supports reasoning without replacing judgment.
Finalization and history	Selected outcome, recommendation version, actor, timestamp, and immutable snapshot	Creates an auditable decision record.
Personal decision library	Reliable resume, search, history, archive, and immutable final snapshots	Supports repeat individual use and durable retrieval.
Operations	Production monitoring, backups, rate limits, incident process, analytics	Supports a controlled external launch.

Launch Objectives

1. Prove that target individual decision owners can activate quickly and complete the end-to-end decision workflow without assisted recovery.
2. Validate AI quality, trust, and explainability under realistic customer content and operating conditions.


3. Establish repeat behavior around decision finalization, retrieval, and a second completed decision.
4. Generate credible willingness-to-pay and conversion evidence without over-scaling customer acquisition.
5. Create an operational learning loop connecting analytics, support, incidents, interviews, and roadmap decisions.

Target Market and Positioning

Element	Launch choice
Beachhead	Individual Product Managers, Engineering Managers, Project Managers, founders, and decision owners inside 20-500 employee technology organizations.
Primary buyer	Individual professional, founder, or functional leader sponsoring a named decision owner.
Primary user	Product Manager, Engineering Manager, Project Manager, or decision owner.
Category	AI Decision Intelligence platform.
Positioning	DecisionFlow turns unstructured notes or transcripts into a reviewed, explainable, durable decision record.
Competitive boundary	Not positioned as a transcription tool; integrates with or follows existing meeting capture systems.

Launch Model

Illustrative launch timeline



Illustrative phased launch from readiness through the Week-12 gate review.

Phase	Purpose	Exit condition
Readiness	Verify production, support, security, data, content, and analytics.	No open critical launch blocker; rollback and incident owners assigned.
Recruitment	Select design-partner and early commercial accounts.	Customer profile, consent, onboarding owner, and success plan confirmed.
Controlled launch	Release to small cohorts with active monitoring.	Activation and operational stability meet minimum thresholds.
Adoption	Improve onboarding, usage, and repeat-use habits.	Repeat decision creation and retrieval are visible.
Optimization	Run pricing, messaging, workflow, and trust experiments.	Evidence supports scale, iteration, or stop decision.
Gate review	Evaluate customer, AI, business, and operational evidence.	Formal Scale / Iterate / Stop decision recorded.

Release Gates


Dimension	Definition	Gate
Activation	New registered users who finalize one decision within 24 hours	>= 70%
Workflow completion	Started decisions that reach finalization	>= 70%
Week-4 activated-user retention	Activated users active again by Week 4	>= 60%
AI decision precision	Correct extracted decisions among extracted decisions	>= 90%
AI decision recall	Actual decisions successfully extracted	>= 85%
Hallucination rate	Unsupported material AI outputs	<= 5%
Trust rating	Average user rating of AI-supported workflow	>= 4.0 / 5
Reliability	Monthly platform availability	>= 99.9%
Commercial signal	Onboarded organizations converting to paid	>= 30%

Illustrative Outcome Dashboard

Simulated results - The following values demonstrate how a launch case study could be evaluated. They are not actual DecisionFlow customer or commercial results.

Metric	Illustrative result	Classification
Organizations invited	20	Simulated
Sponsoring customer accounts onboarded	15	Simulated
Organizations converted to paid	8	Simulated
Individual users invited	118	Simulated
Registered users	94	Simulated
Activated users	73	Simulated
Week-12 retained users	43	Simulated
Decision workflows started	386	Simulated
Decisions finalized	291	Simulated
Workflow completion	75.4%	Simulated
Median time to finalization	8m 42s	Simulated
User trust rating	4.2 / 5	Simulated
SUS	82 / 100	Simulated

Metric	Illustrative result	Classification
Monthly recurring revenue at Week 12	\$2,960	Simulated
Platform uptime	99.93%	Simulated
p95 AI processing time	28 seconds	Simulated



Simulated unique-user funnel. Multiple decisions per activated user are reported separately.

Launch Recommendation

Illustrative decision: Scale with safeguards - The scenario meets activation, workflow completion, precision, recall, reliability, trust, and commercial-signal gates. Scale should remain controlled because owner and due-date accuracy are below the original 90% ambition, activated-user retention declines from 68.5% at Week 4 to 58.9% at Week 12, and customers request post-launch collaboration, integrations, richer ingestion, and deeper explainability.

Document Set Map

Code	Document	Primary question
LAUNCH-00	Launch Phase Overview	What launched, why, and what did the scenario indicate?
LAUNCH-01	Go-to-Market and Commercial Launch Plan	Who is the launch for and how will demand be created?
LAUNCH-02	Release Readiness and Launch Runbook	How is the release executed, monitored, and recovered?
LAUNCH-03	Customer Onboarding and Support Guide	How do customers reach value and get help?
LAUNCH-04	Launch Measurement and Experiment Plan	How are value, quality, trust, and growth measured?

Code	Document	Primary question
LAUNCH-05	Launch Risk, Incident and Feedback Register	How are launch risks, incidents, and feedback governed?
LAUNCH-06	Launch Results and Retrospective	What happened in the simulated scenario and what changes next?