

DECISIONFLOW

Develop Phase Overview

From prototypes to a pilot-ready AI product

Document code	DEV-00
Version	2.0
Status	Final Public Portfolio Edition — Implementation Baseline
Baseline date	14 July 2026
Author	Min Thu Kyaw
Role	AI Product Manager

Evidence transparency: Actual implementation evidence is separated from targets, proposed release actions, and the illustrative DEV-06 continuity bridge. Later commercial outcomes are not presented as actual results.

Document Control

Purpose	Document the Develop phase for Develop Phase Overview.
Primary audience	Hiring managers, product leaders, engineering stakeholders, and pilot reviewers.
Evidence classification	Verified implementation fact / implementation-derived assessment / target / proposed next action.
Source limitation	Implementation status is grounded in the authored implementation baseline and public repository references supplied for the case study.
Change control	Publication rule: retain evidence labels unless content is replaced by documented real-world observations.

Evidence Links

- [Low-Fidelity Prototype](#)
- [High-Fidelity Prototype](#)
- [MVP Frontend Repository](#)
- [MVP Backend Repository](#)

Contents

1. Executive Summary
2. Evidence Model
3. Development Journey
4. MVP Scope and Workflow
5. Readiness Assessment
6. Implementation Snapshot
7. Release-Critical Gaps
8. Architecture Summary
9. Development Milestones
10. Actual MVP Readiness Decision
11. Illustrative Release-Candidate Bridge
12. Document Package Map

Executive Summary

DecisionFlow has progressed beyond concept and planning into a working, integrated MVP. The current product combines a Next.js frontend, FastAPI backend, PostgreSQL-oriented persistence, authenticated user ownership, and an AI workflow that extracts structured decision data and generates explainable recommendations. The Develop phase is therefore not a design exercise; it is a release-hardening and evidence-building phase.

Current decision — The MVP demonstrates the core product hypothesis, but it should not be presented as external-pilot ready until the P0 integrity, AI transparency, security, test, analytics and deployment gaps are resolved.

Evidence Model

Evidence class	Meaning	Examples
Verified implementation fact	Reported by the implementation-derived baseline as present in the current code.	Auth, decision CRUD, AI extraction, recommendation history.
Partial implementation	Present, but incomplete, local-only, or not production-ready.	Reviewed edits, weights, comparison, finalization, resume.
Target	A proposed threshold or release expectation, not an achieved result.	AI precision target, latency target, pilot conversion.
Proposed action	Recommended work required for release.	P0 backlog, session hardening, E2E tests.
Public design evidence	Prototype behavior and visual intent visible on public pages.	Low-fi and high-fi prototypes.

Development Journey



Figure 1. DecisionFlow progression from planning to pilot readiness.

Prototype links: [Low-Fi](#) • [High-Fi](#)

Implementation references: [Frontend](#) • [Backend](#)

MVP Scope and Workflow

The current MVP serves an individual decision owner. A user starts with meeting notes or a UTF-8 transcript, reviews AI-extracted structure, adjusts criteria and context, generates an explainable recommendation, compares options, finalizes the outcome, and retrieves the saved decision later.

- 1. Register or log in.
- 2. Create a clean decision record.
- 3. Paste text or upload a TXT transcript.
- 4. Run structured AI extraction.
- 5. Review topic, options, criteria, stakeholders, risks and actions.
- 6. Generate or regenerate an explainable recommendation.
- 7. Compare options using a defined scoring method.
- 8. Explicitly select and finalize the outcome.
- 9. Retrieve the durable record from the decision library.

Readiness Assessment




Figure 2. Implementation-derived assessment used to prioritize work; not a measured product KPI.

Assessment interpretation — The strongest evidence is the connected product workflow. The weakest areas are release operations, frontend/E2E quality evidence, security hardening, and durable workflow state.

Implementation Snapshot

Capability	Status	Current evidence	Release action
Public landing page	Implemented	Explains product and routes to login/dashboard.	Add analytics and final content review.

Capability	Status	Current evidence	Release action
Registration and login	Implemented	Email/password auth; bearer JWT stored in browser storage.	Harden session design and monitor failures.
Dashboard	Implemented	Live decision counts, confidence and recent records.	Add telemetry and pagination threshold.
Paste text	Implemented	Text submitted for AI extraction.	Add privacy and size guidance.
TXT upload	Implemented	UTF-8 TXT is read and processed.	Align client and backend validation.
AI extraction	Implemented	Structured JSON is validated and stored.	Add benchmark dataset and prompt versioning.
Review extracted fields	Partial	Edits exist in browser state.	Persist every reviewed field and reload by decision ID.
Criteria weighting	Partial	Sliders change local display state.	Persist weights and use them in recommendation context.
AI recommendation	Implemented	Recommendation versions are stored.	Expose provenance, alternatives and fallback status.
Option comparison	Partial	Browser-generated deterministic scores.	Use a documented scoring model and persist the snapshot.
Finalize decision	Partial	Status changes to finalized.	Store selected option, accepted version, actor and immutable snapshot.
Decision library	Implemented	Search, filter, open and archive.	Add pagination and reliable resume.
Resume in-progress decision	Not reliable	Single local browser draft is used.	Hydrate working state from server by decision ID.
PDF/DOCX input	Not implemented	UI mentions formats but rejects them.	Hide for MVP or implement secure parsing.
Notifications / privacy / AI preferences	UI only	Controls do not affect backend behavior.	Remove or label unavailable for MVP.
Collaboration and roles	Not implemented	Single-user ownership model.	Defer to post-MVP.

Release-Critical Gaps

ID	Work item	Required outcome
P0-01	Declare release source of truth	Remove ambiguity between prototypes and current repositories.
P0-02	Decision-ID hydration and resume	Allow reliable cross-session resume.
P0-03	Persist reviewed extraction	Make human edits durable and authoritative.
P0-04	Unify criteria weights	Use the user's actual priorities.
P0-05	Fix comparison scoring	Prevent synthetic scoring from being misrepresented.
P0-06	Semantic finalization	Save selected outcome and immutable audit snapshot.
P0-07	Align supported inputs	Ship paste and TXT only unless more parsers are implemented.
P0-08	Remove demo seed data	Prevent inherited or misleading content.
P0-09	Surface AI provenance	Label provider output, model, version and fallback.

ID	Work item	Required outcome
P0-10	Enforce transitions server-side	Protect record integrity.
P0-11	Security hardening	Secrets, session, recovery and rate limits.
P0-12	Frontend/E2E tests and CI	Protect the core workflow.
P0-13	Staging and deployable frontend	Create an integrated release candidate.
P0-14	Analytics and observability	Measure funnel, failures, latency and cost.

Architecture Summary

- Frontend: Next.js 16, React 19, TypeScript, Tailwind CSS and App Router.
- Backend: FastAPI, Pydantic, SQLAlchemy, Alembic and PostgreSQL runtime.
- AI: OpenRouter-compatible chat completion through the OpenAI Python SDK, structured JSON and Pydantic validation.
- Ownership: backend resource checks isolate decisions, transcripts and recommendations by authenticated user.
- Known boundary problem: the backend owns durable records, but the frontend still owns too much active-draft state in local storage.

Development Milestones

Milestone	Primary outcome	Exit criterion
1. Workflow integrity	Resume, reviewed edits, weights, transitions and final snapshot become server-durable.	A decision can be reconstructed entirely from backend data in a new session.
2. AI trust	Provenance, versioning, fallback labels, comparison methodology and evaluation fixtures.	Every AI result and score has a traceable source and version.
3. Release hardening	Security, E2E tests, analytics, staging, monitoring and recovery.	All release gates pass in staging.
4. Controlled pilot	Small user cohort, weekly evidence review and issue triage.	Scale, iterate or stop decision documented.

Actual MVP Readiness Decision

Recommendation: ITERATE before external pilot — The implementation-derived MVP is suitable for continued internal development and demonstration. External access should wait until ownership, persistence, finalization, AI transparency, security and regression-test gates pass.

Illustrative Release-Candidate Bridge

To maintain continuity with the simulated Launch phase, this package now includes DEV-06. The addendum does not rewrite the actual MVP baseline. It documents an illustrative case-study scenario in which the P0 release-

hardening work is assumed complete and a production-ready DecisionFlow v1.0 core release candidate is approved for controlled launch.

Evidence layer	Meaning	Portfolio use
Actual MVP baseline	Verified implementation-derived state and remaining gaps.	Demonstrates honest technical and product assessment.
Illustrative release candidate	Assumed completion of named P0 items in DEV-06.	Creates the explicit bridge to the simulated Launch phase.
Simulated launch	Illustrative customers, metrics, commercial outcomes and incidents.	Demonstrates launch judgment without claiming real traction.

Continuity rule — Launch v1.0 is limited to the hardened individual decision workflow: paste/TXT intake, durable review, criteria, recommendation, transparent comparison, explicit finalization, personal history, secure access, analytics and production operations. Collaboration, integrations, PDF/DOCX, SSO and enterprise administration remain post-launch.

Document Package Map

Code	Document	Primary question
DEV-00	Develop Phase Overview	What exists, what is missing, and can it enter pilot?
DEV-01	As-Built Implementation & Traceability Report	How does the built product compare with plans and prototypes?
DEV-02	AI Engineering & Evaluation Report	How is AI implemented, controlled and evaluated?
DEV-03	Quality, Security & Release Readiness Report	Is the product technically safe and releasable?
DEV-04	Pilot Validation & Measurement Plan	How will real-user value and trust be validated?
DEV-05	Development Decision Log & Release Notes	What implementation decisions and limitations must remain visible?
DEV-06	Productionization & Release Candidate Completion Addendum	What is assumed complete between the actual MVP and simulated launch?

Evidence and References

Source	Use in this document
Implementation-derived baseline	DECISIONFLOW_MVP_DEVELOPMENT_PHASE.md, baseline date 14 July 2026.
Low-fidelity prototype	https://minthukyawdev-droid.github.io/decision-flow/
High-fidelity prototype	https://minthukyawdev-droid.github.io/decisionflow-hifi-prototype/
Frontend repository reference	https://github.com/minthukyawdev-droid/decisionflow-webapp

Source	Use in this document
Backend repository reference	https://github.com/minthukyawdev-droid/decisionflow-api
Plan-phase package	DecisionFlow Plan Phase Document Set, version 1.0.

Interpretation rule — Targets, proposed metrics, future releases, and pilot thresholds are not reported as achieved results. Actual results must be added only after executing the relevant test or study.