

DECISIONFLOW

Productionization & Release Candidate Completion Addendum

Illustrative bridge from the actual MVP baseline to a production-ready v1.0 core release candidate

Document code	DEV-06
Version	2.0
Status	Final Public Portfolio Edition — Illustrative Productionization Bridge
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 Productionization & Release Candidate Completion Addendum.
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. Purpose and evidence boundary
2. Continuity model
3. v1.0 core scope freeze
4. Assumed P0 completion register
5. Durable workflow and data changes
6. AI trust and evaluation completion
7. Security, quality and operations completion
8. Illustrative release evidence
9. Deferred capabilities
10. Release-candidate decision
11. Traceability to Launch phase

Purpose and Evidence Boundary

This addendum closes the narrative gap between the actual integrated MVP and the simulated Launch phase. It does not claim that the listed work exists in the public repositories. It records the portfolio assumption that the release-critical items identified in DEV-00 through DEV-05 were completed before the illustrative commercial launch began.

Evidence classification — Actual: implementation-derived MVP baseline and public prototypes/repository references. Assumed: the release-hardening completion in this addendum. Simulated: benchmark values, launch customers, revenue, incidents, feedback and results in the Launch package.

Continuity Model

Stage	Evidence status	Outcome
Integrated MVP v0.5	Actual implementation-derived baseline	Core workflow exists; persistence, finalization, AI transparency, quality, security and deployment gaps remain.
Release hardening v0.6-v0.8	Illustrative assumed work	P0 data integrity, AI trust, tests, security, analytics and operational controls are completed.
Release candidate v0.9	Illustrative assumed state	Production-ready core individual workflow passes release gates in staging.
Controlled launch v1.0	Simulated Launch phase	Qualified individual decision owners enter a measured commercial release.

v1.0 Core Scope Freeze

Included at launch	Explicitly deferred
Pasted text and UTF-8 TXT intake	PDF/DOCX ingestion and live meeting capture
Structured extraction and durable human-reviewed version	Shared workspaces, invitations, reviewers, comments and team permissions
Criteria and integer percentage weights totaling 100	Slack, Teams, Jira, Linear, calendar and meeting-platform integrations
Explainable recommendation with immutable versions and provenance	SSO, advanced RBAC, enterprise audit administration and regional processing
Transparent weighted comparison and user-selected final outcome	Native mobile applications and predictive decision intelligence
Personal decision library, reliable resume, search, archive and history	Broad multilingual AI claims without benchmark evidence
Secure individual account, monitoring, analytics, backups and support	Organization-wide knowledge graph and cross-project analytics

Product definition — DecisionFlow v1.0 is a production-ready personal decision intelligence workspace for knowledge workers. It is not yet a collaborative organizational decision-memory platform.

Assumed P0 Completion Register

P0 item	Assumed completion	Release evidence expected
Decision-ID hydration and resume	All in-progress records load from backend by ID in any authorized session.	E2E resume tests and new-session smoke test.
Reviewed extraction persistence	Original AI output and human-reviewed version are stored separately; reviewed data drives recommendations.	API tests, migration and reload verification.
Criteria contract	Integer percentages total 100 across UI, API, database and AI context.	Validation and mathematical tests.
Comparison integrity	Weighted matrix has documented source, user-overridable scores and persisted snapshot.	Calculation tests and UI provenance labels.
Semantic finalization	Selected option, accepted recommendation ID, actor, timestamp and immutable snapshot are stored.	Finalization API/E2E/audit tests.
AI provenance	Provider, model, prompt version, usage and fallback state are visible and stored.	Response contract and UI checks.
Workflow state machine	Server enforces legal transitions and idempotent commands.	Transition and retry tests.
Security hardening	Production secrets, secure session approach, rate limits, password recovery and privacy disclosure are complete.	Security checklist and negative tests.
Quality evidence	Frontend build, component, integration, E2E, PostgreSQL migration and smoke tests pass in CI.	Green release pipeline.
Operations and learning	Staging/production deployment, logs, alerts, analytics, backup/restore and rollback are verified.	Runbook sign-off and dashboards.

Durable Workflow and Data Changes

- Every workflow route resolves a decision ID and hydrates authoritative backend state before rendering.
- Original extraction, reviewed extraction, recommendation versions, comparison snapshot and final snapshot remain separately traceable.
- Finalized records are immutable; later change creates a revision or duplicate rather than silently rewriting history.
- Analytics capture metadata and workflow events only, never transcript or decision content.
- Archive behavior is reversible under policy; destructive deletion follows an explicit retention and audit decision.

AI Trust and Evaluation Completion

Measure	Illustrative RC result	Gate	Interpretation
Decision extraction precision	91.8%	>= 90%	Pass
Decision extraction recall	86.4%	>= 85%	Pass

Measure	Illustrative RC result	Gate	Interpretation
Action-item precision	90.5%	>= 90%	Pass
Owner accuracy	88.7%	Advisory; confirmation required	Below ambition but safe with mandatory review
Due-date accuracy	84.1%	Advisory; confirmation required	Below ambition; relative dates remain a known limitation
Material hallucination rate	3.6%	<= 5%	Pass
p95 AI latency	28 seconds	<= 35 seconds	Pass

Simulation disclosure — These benchmark values are illustrative and are reused in the Launch retrospective for continuity. They are not results from an executed public-repository benchmark.

Security, Quality and Operations Completion

Area	Assumed release-candidate state
Session and identity	External-user session design approved; password reset and expiry/revocation behavior tested.
Authorization	100% owner isolation verified across decisions, sources, recommendations, history and final snapshots.
Privacy	AI-provider disclosure, source retention, deletion path and consent language published.
Abuse and cost	Authentication and AI endpoints rate-limited; idempotency and cost budgets active.
Quality	Backend, frontend build, component, integration, E2E, migration, accessibility and browser tests pass.
Operations	Production-like staging, monitored deployment, alerts, backups, restore drill, rollback and incident ownership pass.
Measurement	Activation, workflow, AI, latency, cost, return and support events are visible without sensitive content.

Deferred Capabilities

- PDF and DOCX parsers are not part of v1.0 and remain a post-launch controlled experiment.
- Collaboration, shared workspaces, reviewers, permissions and comments remain post-launch discovery and development.
- Jira, Linear, Slack, Teams, calendar and meeting-platform integrations are roadmap items, not launch promises.
- SSO, enterprise administration, advanced retention, regional processing and executive analytics are not generally available.
- Organization-wide decision memory and knowledge graphs remain the longer-term product vision.

Release-Candidate Decision

Illustrative decision: GO for controlled v1.0 launch — The assumed release candidate passes the integrity, AI transparency, security, quality, operations and measurement gates for a controlled commercial launch of the individual core workflow. It is not approved for enterprise or team-collaboration promises.

Traceability to Launch Phase

Launch package statement	Required interpretation
"Production-ready product"	Means production-ready v1.0 core individual workflow, not every long-term roadmap capability.
Commercial accounts	Organizations may sponsor individual seats; shared team workflows are not implied.
Team/collaboration requests	Simulated post-launch feedback and roadmap evidence, not launch functionality.
PDF/DOCX and integrations	Post-launch experiments or roadmap, not v1.0 scope.
Business/enterprise packages	Planned or design-partner offers until their features are built and evidenced.

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
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.