

DECISIONFLOW

Segment Expansion and Retention Plan

Proving repeated value across larger cohorts without losing the validated beachhead

CONTROLLED SCALE PHASE | Version 1.0 | Illustrative 12-month bounded-growth program

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
Document code	CS-02
Version / status	2.0 — Final Public Portfolio Edition — Simulated Controlled-Scale Scenario
Baseline / author	14 July 2026 - Min Thu Kyaw
Scenario horizon	Illustrative 12-month period after Deliver; outcome data is simulated.
Lifecycle continuity	Actual MVP -> assumed productionization -> simulated Launch -> Qualify -> Deliver -> Controlled Scale.
Scale decision	Grow the individual workflow within tested customer, AI, commercial, support, and capacity boundaries.
Scope boundary	No broad collaboration, integrations, richer ingestion, multilingual promise, SSO, or enterprise-platform claim.
Evidence rule	Actual, assumed, simulated, target, and recommendation statements are separated.

Actual Project Artifacts

Artifact	Reference
Low-Fidelity Prototype	View prototype
High-Fidelity Prototype	View prototype
MVP Frontend Repository	Open repository
MVP Backend Repository	Open repository

Contents

01	Retention objective and metric contracts
02	Cohort architecture
03	Lifecycle and activation system
04	Retention results
05	Customer health and intervention
06	Qualitative research
07	Retention decision

Retention Objective and Metric Contracts

The retention program measures whether customers repeatedly use DecisionFlow for real decisions after the initial novelty period. The unit is an activated user who finalized at least one decision; retained users must perform meaningful decision activity in the defined period.

Metric	Definition
Activation	Registered user finalizes a first decision within 14 days.
Second-decision rate	Activated user finalizes a second distinct decision within 30 days.
W12 retention	Activated user performs a qualifying decision action in days 78-90.
W24 retention	Activated user performs a qualifying decision action in days 162-180.
W52 retention	Activated user performs a qualifying decision action in days 351-365.
Qualifying action	Create, resume, finalize, or retrieve a substantive decision; login alone does not count.

Denominator discipline - All retention percentages use the activated cohort denominator. Cohort start, observation window, exclusions, and late-arriving events are frozen before reporting.

Cohort Architecture

Cohort	Size	Purpose	Guardrail
Product / Program Managers	250 activated	Primary repeatability test for recurring decision jobs.	W24 $\geq 60\%$; trust $\geq 4/5$.
Founders / Operators	170 activated	Test faster, less-structured strategic decisions.	Completion $\geq 80\%$; support ≤ 0.20 contacts/activation.
Consultants	100 activated	Test client-facing evidence and export demand.	Paid retention $\geq 80\%$; no unsupported sharing claim.
Total	520 activated	Controlled-scale denominator.	No acquisition beyond service and support boundary.
Cohort control		Practice	
Acquisition tag		Source and offer captured without decision content.	
Role / job tag		Self-reported role and primary decision job.	
Exposure logging		Templates, lifecycle prompts, routing, and pricing experiment exposure recorded.	
Privacy		No transcript text or rationale content in product analytics.	
Comparability		Metrics reported by cohort age, not calendar totals alone.	


Lifecycle and Activation System

Lifecycle moment	Product / service intervention	Purpose
Before first decision	Role-specific example, privacy notice, supported-format guidance.	Reduce uncertainty and setup errors.
During extraction	Progress, fallback/provenance disclosure, editable-field guidance.	Preserve trust and control.
Before finalization	Accountability confirmation and decision-quality checklist.	Improve durable record quality.
Days 3-14	Retrieval prompt and second-decision template suggestion.	Build a repeated workflow, not a one-time demo.
Days 21-45	Value recap using counts/timing only; no sensitive content.	Reinforce utility and identify inactive users.
At risk	Contextual help and human support based on failure or inactivity signals.	Recover value without generic spam.

- Lifecycle messaging is capped and preference-aware.
- No message exposes decision content in email subject lines or analytics.

- Support interventions prioritize blocked workflows, AI trust questions, and repeated fallback—not simple inactivity alone.

Retention Results



Measure	Simulated result	Interpretation
Activation	520 / 620 = 83.9%	Held near Deliver baseline despite higher volume.
Second decision within 30 days	392 / 520 = 75.4%	Strong evidence that value is not limited to first use.
W12 retained	353 / 520 = 67.9%	Stable repeat value.
W24 retained	306 / 520 = 58.8%	Passes controlled-scale gate.
W52 retained	244 / 520 = 46.9%	Meaningful annual use, but also a clear long-horizon ceiling.
PMF survey very disappointed	48.6%	Strong in the beachhead; sample remains self-selected.

Customer Health and Intervention

Health state	Signals	Action
Healthy	Second decision, recent finalization/retrieval, low failure, positive trust.	Education and optional advanced templates.
Promising	First decision complete; no second decision yet.	Role-based second-decision guidance.
At risk	Repeated extraction error, fallback, abandoned review, or 30-day inactivity.	Contextual recovery and optional support.
Commercial risk	Payment issue, high support use, weak value evidence.	Health review before renewal or expansion.
Trust risk	High correction, low confidence understanding, unsupported-output report.	AI-quality review and direct follow-up.
Service indicator	Deliver result	Scaled result
Self-service activation	72%	80%
Support contacts per activation	0.24	0.16

Health state	Signals	Action
Median first value	6m 55s	6m 20s
Unresolved high-severity customer issue	0	0 at close

Qualitative Research

Theme	Simulated evidence	Product implication
Durable rationale matters	Users retrieve records before planning, reviews, or revisiting trade-offs.	Keep history and retrieval central.
Accountability creates value	Owner/date confirmation is seen as useful despite added friction.	Do not remove deliberate confirmation.
Recurring templates help	Frequent roles create similar decision structures repeatedly.	Invest in role-specific templates.
Collaboration requests vary	Some need comments; others need formal approval, SSO, or integrations.	Treat as multiple problem spaces, not one feature.
Consultants need portable evidence	Export demand is stronger than shared workspace demand in this segment.	Prototype export before collaboration.
Long-term use has a ceiling	Some decisions are episodic; not every activated user has weekly demand.	Measure value per decision and segment, not DAU obsession.

Retention Decision

Decision - RETENTION IS QUALIFIED FOR CONTINUED BEACHHEAD SCALE, NOT BROAD HORIZONTAL EXPANSION. W24 improves to 58.8% and W52 reaches 46.9%, demonstrating durable value while revealing a finite usage ceiling.

- Continue role-specific templates and retrieval journeys.
- Report retention by job and cohort age; avoid blending segments.
- Do not force artificial engagement for an episodic decision product.
- Use the W52 ceiling in the strategic review and standalone-product economics.

Evidence discipline: replace all simulated controlled-scale values with observed, reproducible results before presenting them as actual product performance.