

DecisionFlow

Controlled Scale Phase Overview

Expanding DecisionFlow within validated customer, AI, commercial, and operational boundaries

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-00
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	Executive scale decision
02	Lifecycle continuity and purpose
03	Controlled-scale thesis and boundaries
04	Twelve-month scorecard
05	Integrated workstreams
06	Evidence and governance model
07	Strategic review outcome
08	Document set map

Executive Scale Decision

Phase mandate - Scale the production-ready individual decision workflow through bounded cohorts, repeatable channels, stronger AI operations, and higher tested capacity. Growth is permitted only while retention, safety, contribution, and support thresholds remain healthy.

DecisionFlow enters this phase after the Deliver review concluded Continue Controlled Scale. The purpose is not to maximize sign-ups or feature breadth. It is to determine whether the released core can grow repeatedly without eroding customer value, AI accountability, service economics, or operational control.

Decision	Meaning
Approved	Increase customer volume, paid accounts, and tested infrastructure capacity within explicit limits.
Not approved	Unrestricted acquisition, enterprise repositioning, collaboration build, broad integrations, or richer-ingestion marketing.
Review trigger	Any failed safety, retention, unit-economics, data-isolation, or capacity gate pauses expansion.
End-of-phase decision	Maintain controlled scale and begin a standalone-product strategic review; do not fund an enterprise rebuild on current evidence.

Lifecycle Continuity and Purpose

Lifecycle stage	Evidence status	Contribution to this phase
Concept and Plan	Actual case-study artifacts	Problem, users, product strategy, requirements, metrics, and scope.
Develop MVP	Actual implementation-derived evidence	Working frontend/backend/AI baseline and honest gaps.
DEV-06 productionization	Assumed scenario bridge	Release-hardening capabilities required for v1.0 continuity.
Launch and Qualify	Simulated	Controlled commercial evidence and Conditional Scale decision.
Deliver	Simulated	Closed the first retention, AI, support, and 40-job capacity conditions.
Controlled Scale	Simulated	Tests repeatability across larger cohorts, a second channel, 80-job capacity, and longer retention.


Continuity rule - The actual MVP is never retroactively described as the scaled product. Productionization and all post-MVP outcomes remain explicitly assumed or simulated.

Controlled-Scale Thesis and Boundaries

Scale dimension	Approved boundary	Excluded claim
Customer segment	Frequent individual decision owners: product/program managers, founders/operators, and selected consultants.	Broad horizontal knowledge-worker market.
Product	Paste/TXT intake, human-reviewed extraction, criteria, recommendations, comparison, finalization, history, and resume.	Collaboration, PDF/DOCX, integrations, or enterprise administration.
Commercial	Starter and Professional packages; two controlled acquisition channels.	Enterprise contract or platform pricing.
AI	English-language core use cases with model routing, fixed evaluation, provenance, and human finalization.	Autonomous decisions or broad multilingual quality claims.
Operations	80 concurrent AI jobs at the approved latency/failure gate.	Unlimited capacity or 150-job production claim.

- Scale is an evidence-controlled change in volume, not a declaration of product-market fit.
- No metric can override a critical safety, privacy, authorization, or data-integrity failure.
- Requested features remain discovery inputs until value, willingness to pay, risk, and operating feasibility are quantified.


Twelve-Month Scorecard



Controlled-scale KPI movement

Dimension	Deliver result	Controlled-scale result	Assessment
Activated users	110	520	Volume expanded while activation remained stable.
Start-to-finalize	82.0%	84.1%	Core workflow quality held.
Week-24 retention	56.4%	58.8%	Improved; supports beachhead repeatability.
Week-52 retention	Not measured	46.9%	Useful but shows a meaningful long-horizon ceiling.
Paying accounts	15	78	Commercial proof strengthened, still segment-concentrated.
MRR	\$6,840	\$31,420	Growth is meaningful but not enterprise-scale.
AI variable cost	\$0.29	\$0.23	Routing and prompt efficiency improved.
Support contribution	58%	65%	Self-service economics strengthened.
80-job load	68s / 1.4%	43s / 0.8%	Passes controlled-scale gate.

Integrated Workstreams




Controlled-scale roadmap

Workstream	Outcome
Beachhead retention	Role-specific templates, retrieval habits, recurring decision journeys, and lifecycle interventions.
Channel repeatability	Validate one additional channel without exceeding payback or support thresholds.
AI quality and economics	Maintain fixed quality gates while reducing latency, fallback, and cost.
Infrastructure capacity	Qualify 80 concurrent AI jobs and keep 150 jobs outside the approved boundary.
Customer operations	Increase self-service activation and lower support contacts per activation.
Strategic discovery	Quantify the value and risk of collaboration, integrations, and enterprise expectations without building them.

Evidence and Governance Model

Evidence class	Use in decision-making
Actual	Prototype, repository, and MVP implementation evidence only.
Assumed	DEV-06 release-hardening capabilities used to connect the actual MVP to the simulated v1.0 product.
Simulated	Cohorts, customer outcomes, AI benchmarks, financials, incidents, and scale results in this package.
Target	A threshold not yet achieved or an approved operating boundary.
Recommendation	A proposed action requiring owner approval and future evidence.
Gate	Pause / escalation condition
Customer	W24 retention <55%, significant trust decline, or rising unresolved workflow failure.
AI	Quality regression, ECE >0.10, unsupported-output increase, provenance gap, or silent fallback.
Commercial	Observed payback >12 months, contribution <55%, or support demand invalidates packaging.
Operations	Availability <99.9%, failed isolation/security gate, or capacity outside the tested boundary.
Strategy	Growth requires an unapproved collaboration or enterprise rebuild.

Strategic Review Outcome



Controlled-scale evidence scorecard

End-of-phase decision - MAINTAIN CONTROLLED SCALE AND INITIATE A STANDALONE-PRODUCT STRATEGIC REVIEW. The core is valuable and economically supportable within its beachhead, but evidence does not justify a major enterprise-platform rebuild.

- Continue serving and improving the validated individual workflow.
- Keep acquisition and capacity within tested economics and service boundaries.
- Evaluate maintain, partner/license, capability integration, and eventual retirement options over the next review horizon.
- Do not market collaboration, integrations, richer ingestion, or enterprise controls as current functionality.

Document Set Map

Code	Document
CS-00	Controlled Scale Phase Overview
CS-01	Controlled Growth Strategy and Portfolio Roadmap
CS-02	Segment Expansion and Retention Plan
CS-03	Commercial Scaling and Channel Economics Report
CS-04	AI Scale, Quality and Model Operations Plan
CS-05	Technical Capacity, Reliability and Cost Expansion Plan
CS-06	Operating Model, Governance and Risk Playbook
CS-07	Controlled Scale Results and Strategic Review

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