DECISIONFLOW

Controlled Growth Strategy and Portfolio Roadmap

Scaling the validated individual workflow without prematurely becoming an enterprise platform

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-01
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

Strategic Objective

The controlled-growth strategy expands DecisionFlow only where the released core solves a frequent, high-friction decision-documentation job. Growth is judged by retained finalized decisions, paid retention, AI quality, and support-adjusted contribution—not by registrations alone.

Strategic thesis - Win a narrow category of frequent individual decision owners before considering a multi-user platform. Preserve decision accountability as the differentiated value proposition.

Objective	12-month criterion
Retained usage	W24 activated-user retention >=55%; W52 measured and interpreted by segment.
Commercial repeatability	At least two controlled channels with observed payback <=12 months.
AI accountability	Owner/date accuracy >=92%, ECE <=0.08, human finalization retained.
Service economics	Support-adjusted contribution >=60%; contacts/activation <=0.20.
Capacity	80 concurrent AI jobs at p95 <=45 seconds and failure <1%.
Scope discipline	No unqualified collaboration, integration, or enterprise build.

Beachhead and Adjacent Segments

Segment	Core job	Evidence fit	Scale treatment
Product / Program Managers	Turn complex meeting evidence into an accountable decision record.	Strongest repeat usage and templates.	Primary beachhead.
Founders / Operators	Make fast strategic and operational trade-offs with preserved rationale.	Strong value; more variable frequency.	Controlled secondary segment.
Independent Consultants	Create structured, client-ready decision evidence.	Promising willingness to pay; export demand.	Limited channel cohort.
General knowledge workers	Summarize meetings and notes.	Problem broad but differentiation weaker.	Do not target broadly.
Enterprise buying groups	Collaborative governance, approvals, SSO, audit, integrations.	Requirements exceed current product.	Discovery only.

- Segment expansion is allowed only when the job, repeat frequency, and willingness to pay match the core workflow.
- Requests for team features are logged as evidence of a different product surface, not treated as minor enhancements.
- Messaging stays focused on accountable decisions, not generic meeting summarization.

Growth Principles and Boundaries

Principle	Operating implication
Retained value before reach	Do not increase acquisition faster than W24 retention and support economics can be measured.
One channel at a time	Add a second repeatable channel before diversifying further.
No hidden product expansion	A request involving shared access, permissions, external systems, or enterprise controls triggers a separate discovery gate.
Human accountability remains	AI may recommend and structure; users explicitly finalize outcomes.
Capacity is measured	Marketing and sales do not exceed the approved service boundary.
Strategic optionality	Preserve the ability to maintain, partner, integrate capabilities, or retire without an irreversible enterprise rebuild.

Portfolio Roadmap

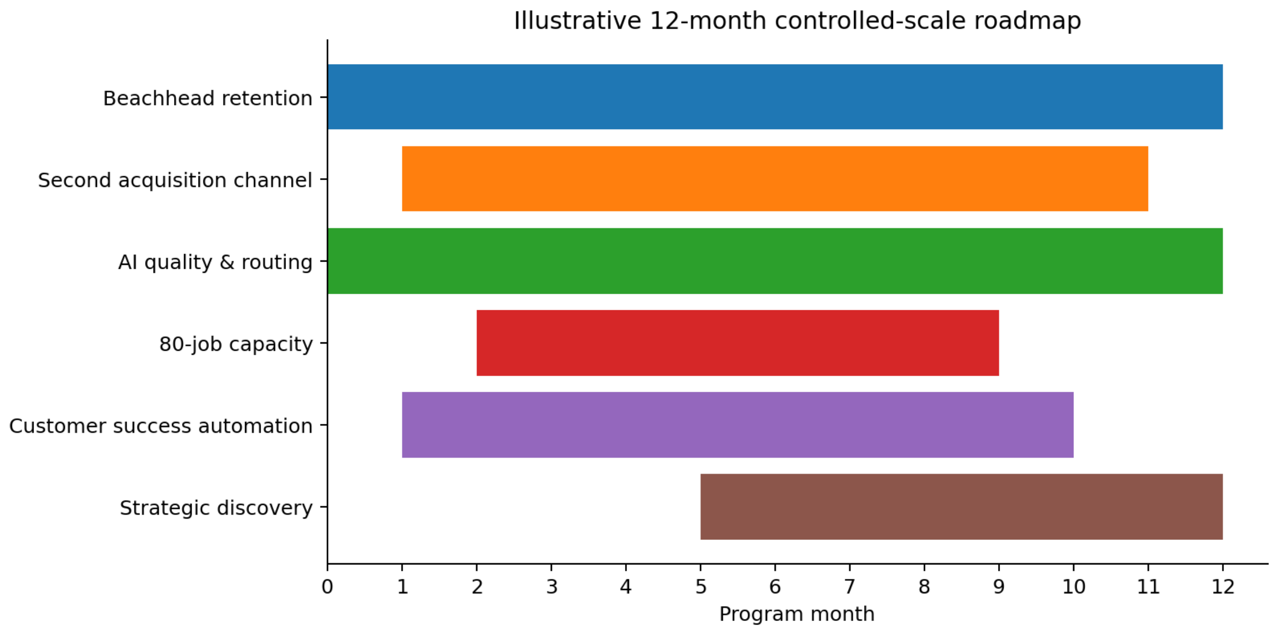

Controlled-scale portfolio roadmap

Horizon	Committed	Discovery / option value	Explicitly deferred
Months 1-3	Role templates, lifecycle instrumentation, partner-channel experiment, AI regression cadence.	Export job and collaboration interviews.	Shared workspace build.
Months 4-6	Retention journeys, support automation, 80-job architecture work, pricing refinement.	Webhook prototype and permission model concept.	Production integrations.
Months 7-9	Second-channel proof, routing economics, capacity qualification, paid-health reviews.	Enterprise procurement interviews.	SSO/admin/compliance build.
Months 10-12	Cohort review, strategic option analysis, roadmap reset.	Partner/license and capability-integration options.	Enterprise-platform commitment.

Experiment Portfolio

Experiment	Hypothesis	Success evidence	Stop condition
Role-specific templates	Frequent decision owners activate and return more often when the first workflow matches their job.	Activation +4pp and W12 +3pp without support increase.	No meaningful lift or trust decline.
Partner referrals	Selected PM/coaching/consulting partners can acquire qualified users efficiently.	Payback <=12m and paid retention >=80%.	Low activation, high support, or weak payback.
Retrieval habit loop	Weekly decision review increases second/third decision completion.	Higher W24/W52 retention in exposed cohort.	Notification fatigue or no retained lift.
Cost-aware routing	Simpler tasks can use lower-cost routes without quality regression.	Cost decrease with all AI gates maintained.	Any accountable-field or groundedness regression.
Export prototype	Consultants pay for a structured export before collaboration is necessary.	Observed repeated use and willingness to pay.	Request remains verbal or usage is one-off.

Capability Decision Gates

Capability	Required evidence before delivery	Current decision
PDF/DOCX ingestion	Security/parser evaluation, accuracy by format, support cost, and repeated customer value.	Deferred.
Collaboration	Validated multi-user job, permission model, willingness to pay, threat model, and data-isolation architecture.	Discovery only.

Capability	Required evidence before delivery	Current decision
Integrations	Repeated workflow interruption, target-system demand, webhook/API security, and economics.	Prototype / discovery only.
Multilingual commercial support	Language-specific evaluation, reviewed UI, support readiness, and provider consistency.	Not marketed.
Enterprise controls	Qualified enterprise segment, procurement evidence, architecture and compliance investment case.	No current build.

Strategic Options

Option	When it becomes credible	Current view
Continue controlled standalone growth	Retention, economics, and support remain healthy in the beachhead.	Viable near-term.
Build enterprise collaboration platform	Strong multi-user demand, willingness to pay, funding, and technical qualification.	Not justified.
Partner or license capabilities	A platform partner values the AI/evaluation/decision-record assets.	Strategic option to investigate.
Integrate capabilities into another product	Standalone growth plateaus but core workflows remain valuable.	Preserve option value.
Maintain / harvest	Product remains profitable but growth is structurally limited.	Possible later state.
Retire responsibly	Standalone investment no longer meets strategic or economic thresholds.	Future option, not current action.

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