DECISIONFLOW

Operating Model, Governance and Risk Playbook

Keeping customer growth, AI changes, commercial commitments, and service capacity under one evidence system

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-06
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

Operating Model

Controlled scale is governed as one system. Product growth, AI quality, pricing, support, security, infrastructure, and strategic scope are reviewed together so that one team cannot create commitments another team cannot safely deliver.

Operating domain	Accountability
Product	Customer value, scope, retention, roadmap, and metric contracts.
Engineering / Platform	Reliability, capacity, data integrity, security implementation, and recovery.
AI / Backend	Prompt/model routes, evaluation, provenance, provider operations, and cost.
Customer success	Onboarding, support, health, education, and feedback synthesis.
Commercial	Packaging, channel evidence, pricing, and promise discipline.
Security / Privacy owner	Risk review, sensitive-data handling, incident response, and approvals.
Executive sponsor	Investment boundary and strategic option decision.

Decision Rights and RACI

Decision	A	R	C	I
Product scope / roadmap	Product lead	Product	Engineering, AI, CS, Commercial	Executive
AI model/prompt/routing release	AI/Technical owner	AI + Engineering	Product, Security	CS, Commercial
Capacity ceiling increase	Technical owner	Platform	AI, Product, Finance	Commercial, CS
Pricing / package change	Product/Commercial owner	Commercial + Product	Finance, CS, Engineering	Executive
Security/privacy incident	Security owner	Incident team	Product, Legal/Privacy, Engineering	Executive, affected customers
Enterprise/collaboration investment	Executive sponsor	Product discovery	Engineering, Security, Commercial	All functions
Strategic maintain/partner/retire option	Executive sponsor	Product + Finance	All accountable owners	Stakeholders

Management Cadence

Cadence	Agenda	Output
Daily operational review	Incidents, queue/provider state, billing/access failures, critical support.	Operational actions and owner.
Weekly growth review	Activation, retention leading indicators, channel cohorts, support and capacity.	Experiment and acquisition decisions.
Weekly AI quality review	Evaluation, edit rates, fallbacks, cost, provider/model changes, incidents.	Approve, block, canary, or rollback.
Monthly business review	Paid retention, MRR, payback, contribution, segment/channel mix.	Budget and packaging decisions.
Quarterly risk / roadmap review	Evidence register, security/privacy, architecture boundary, capability discovery.	Roadmap reset and investment gates.
End-of-phase strategic review	Standalone ceiling, options, asset value, and future investment.	Maintain, invest, partner, integrate, or retire path.

Growth and Change Controls

Change	Required evidence
Acquisition increase	Current channel retention/payback, service capacity, support headroom, and budget.
New package / price	Deployed entitlement, value evidence, margin/support model, and clear terms.
AI route change	Fixed evaluation, cost/latency, canary, provenance, and rollback.
Capacity increase	Load/failure tests, provider quota, database/queue headroom, recovery, and cost.

Change	Required evidence
New input / language	Security/parser/language evaluation, support readiness, and accurate marketing.
Collaboration / enterprise capability	Separate discovery, multi-user threat model, architecture, willingness to pay, and investment approval.

Stop-the-line authority - Security, privacy, data-isolation, and critical AI-accountability failures can pause releases and acquisition without waiting for the next business review.

Risk and Assumption Register

Risk / assumption	Evidence strength	Impact	Next control
Beachhead retention generalizes only partly	Moderate	High	Report by role, channel, and cohort age.
W52 retention creates a standalone ceiling	Moderate	High	Longer cohorts and strategic economics.
Two channels remain repeatable	Moderate	High	Continue channel-specific payback and paid-retention reporting.
Primary provider remains available/economic	Moderate	High	Game days, alternate route, contract and cost review.
80-job capacity is sufficient near term	Strong within test	Medium	Monitor observed peaks; requalify before raising cap.
Collaboration demand can justify a platform rebuild	Weak / mixed	Critical investment risk	Jobs, willingness to pay, threat model, prototype, financial case.
Enterprise claims can be deferred without harming core	Moderate	Medium	Segment and lost-deal analysis.

Strategic Review Process

1. Freeze metric definitions and evidence window.
2. Review customer value, retention, AI quality, economics, capacity, risks, and product-boundary integrity.
3. Estimate standalone maintenance/growth cost and the incremental investment for collaboration or enterprise expansion.
4. Value reusable assets: prompts, evaluation suite, decision schema, workflows, repositories, customer learning, and brand/domain.
5. Compare continue, invest, partner/license, capability integration, maintain/harvest, and retire options.
6. Record assumptions, dissent, decision owner, decision date, and next review trigger.

Strategic option	Key question
Continue controlled scale	Can the current core keep creating retained value and acceptable contribution?
Invest in broader platform	Is multi-user/enterprise value large enough to justify architecture, security, and GTM investment?
Partner / license	Can another platform commercialize the capability more efficiently?
Integrate capability	Should the AI/evaluation/decision-record assets become part of another product?
Maintain / harvest	Can a stable product serve customers responsibly with lower investment?
Retire	Has standalone value fallen below strategic/economic thresholds, and can customers transition safely?

Governance Decision

Decision - THE CONTROLLED-SCALE OPERATING MODEL IS APPROVED. Growth remains bounded by shared evidence gates, and a formal standalone-product strategic review is initiated at phase close.

- No function can sell, build, or operate outside the agreed product and capacity boundary unilaterally.
- All metrics and outcomes retain evidence classification.
- Strategic optionality is treated as a managed product asset, not a sign of failure.