DECISIONFLOW

Qualification Decision Memo and Growth Roadmap

Formal phase-gate recommendation, investment boundaries, and next 90 days

QUALIFY PHASE | Version 1.0 | 14 July 2026

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	QUAL-06
Version / status	2.0 — Final Public Portfolio Edition — Simulated Qualification Scenario
Baseline / author	14 July 2026 — Min Thu Kyaw
Role / purpose	AI Product Manager — Qualification Decision Memo and Growth Roadmap
Lifecycle continuity	Actual MVP → DEV-06 illustrative productionization → simulated v1.0 launch → qualification decision.
Evidence rule	Actual artifacts, assumed capabilities, simulated outcomes, and recommendations are labelled separately.

Actual Project Artifacts

Artifact	Reference
Low-Fidelity Prototype	View low-fidelity prototype
High-Fidelity Prototype	View high-fidelity prototype
MVP Frontend Repository	Open frontend repository
MVP Backend Repository	Open backend repository

Contents

Decision Requested

Should DecisionFlow scale, iterate, pivot, or stop after the simulated v1.0 launch and qualification review? The decision must account for customer value, retention, AI trust, business economics, operational limits, and the evidence quality behind every conclusion.

Brief Answer

Recommendation — CONDITIONAL SCALE. Continue investing in the individual DecisionFlow v1.0 workflow for high-frequency decision owners, expand only through controlled cohorts, and require a second qualification gate before broad paid acquisition or scope expansion.

DecisionFlow should not pivot or stop because the simulated evidence shows a coherent value proposition, strong activation and completion, meaningful Week-24 retention, narrow product-market-fit evidence, trusted AI under human oversight, early renewal, and reliable operation at qualified load. It should not scale without conditions because retention is segment-dependent, due-date accuracy and confidence calibration miss gates, unit economics depend on high-touch support, and the architecture fails the higher-concurrency test.

Evidence Summary

Dimension	Key evidence	Decision
Customer value	77.7% activation; 75.4% completion; 40.5% PMF; retrieval is top theme.	Scale in the beachhead.
Retention	58.9% W12 and 49.3% W24 activated-user retention.	Improve habit loop before broad expansion.
AI	Precision 92.4%; recall 87.2%; hallucination 3.1%; trust 4.3/5.	Qualified with human review.
AI gaps	Due-date 87.8%; ECE 0.12; one primary provider.	Mandatory confirmation and calibration work.
Business	7/8 retained paid orgs; 107.4% NRR; 82% gross margin.	Promising.
Business gap	47% support-adjusted contribution; CAC modeled.	Do not scale paid acquisition yet.
Technical	99.95% availability; 20-job load pass; restore pass.	Controlled growth.
Technical gap	40-job and 80-job stress profiles fail.	Async architecture is a scale condition.

Alternatives Considered

Option	Why considered	Why not selected
Scale immediately	Core launch gates and early commercial signals are positive.	Evidence is small-sample; margin, retention, calibration, and capacity remain limiting.
Conditional scale	Preserves momentum while tying investment to explicit evidence gates.	Selected.
Iterate without growth	Would reduce operational risk.	Too conservative given qualified customer and AI value.
Pivot to meeting summarization	Could broaden the market.	Weakens differentiation and durable-decision value.
Stop	Avoids future AI and operational risk.	Not supported by the positive value and renewal evidence.

Scale Boundaries

Scale now	Do not scale yet

Scale now	Do not scale yet
Individual Starter and Professional packages.	Team, Business, or Enterprise packages.
English-language AI-content workflows.	Commercial multilingual quality claims.
Paste and UTF-8 TXT intake.	PDF/DOCX or meeting-platform ingestion.
Product Managers, founders, and high-frequency decision owners.	Broad horizontal knowledge-worker acquisition.
Controlled cohort of up to 30 qualified organizations.	Unrestricted public signup or large paid campaigns.
Private personal decision library and history.	Shared workspaces, comments, approvals, and roles.

90-Day Growth Roadmap

Priority	Workstream	Outcome metric	Gate
P0	Retention: recurring decision templates, second-decision journey, retrieval onboarding.	Next-cohort W24 retention.	≥55%
P0	AI accountability: relative dates, speaker context, mandatory owner/date confirmation.	Due-date and owner accuracy.	≥90% each or remove claim
P0	Confidence: calibration, copy, evidence-first display.	ECE and comprehension.	≤0.10; ≥85%
P0	Scale architecture: async jobs, idempotency, queue/back-pressure, provider health.	40 concurrent AI jobs.	p95 ≤45s; failures <1%
P0	Support efficiency: self-service onboarding and privacy-safe diagnostics.	Contribution margin.	≥55%
P1	Commercial evidence: second controlled acquisition channel.	Observed CAC and payback.	≤12 months
P1	Integration discovery: export/webhook prototype before native integrations.	Retention or conversion lift.	Predefined experiment pass
P2	Richer ingestion and collaboration discovery.	Validated demand and safe design.	No build commitment before evidence

Investment and Ownership

Area	Recommended allocation	Accountable role
Product and UX	Retention, onboarding, templates, confidence comprehension.	Product Manager
AI engineering	Due-date/owner quality, calibration, evaluation, provider abstraction.	AI/Backend Lead
Platform	Async generation, capacity, observability, failover.	Platform/Backend Lead
Customer success	Controlled onboarding, recurring-use plans, feedback quality.	Customer Success Lead
GTM	Narrow ICP acquisition, pricing evidence, renewal.	GTM Lead
Security/privacy	Quarterly review, provider and retention controls.	Security Owner

Next Qualification Gate

Gate area	Minimum evidence for broader scale
Customer	At least 30 activated users in second cohort; W24 retention ≥55%; PMF ≥40% in the primary segment.
AI	Due-date/owner gate achieved or product claim reduced; ECE ≤0.10; hallucination ≤5%.
Business	≥15 paying organizations; 90-day logo retention ≥80%; support-adjusted contribution ≥55%; observed CAC payback ≤12 months.
Technical	40 concurrent AI jobs with p95 ≤45 sec and failures <1%; successful provider and restore game days.
Security	No open critical/high finding; privacy, retention, and provider disclosures reviewed.

Gate area	Minimum evidence for broader scale
Scope	No collaboration, enterprise, integrations, or richer-ingestion launch without separate evidence and release gates.

Product Manager Reflection

The important product decision is not whether the case-study metrics look positive. It is whether the evidence is strong enough for the next irreversible investment. DecisionFlow qualifies for continued investment because the core workflow, human-controlled AI, durable record, and focused customer segment reinforce one another. The same evidence also sets boundaries: the product should not use a successful individual workflow to imply team collaboration, multilingual quality, enterprise readiness, or unlimited technical scale.

In a real product, this memo would be replaced by an accountable phase-gate meeting using observed analytics, customer research, audited AI evaluation, financial records, security evidence, and operating data. The decision would remain conditional until those artifacts were independently reproducible.

Final phase-gate decision — CONDITIONAL SCALE — approve the next 90-day investment within the stated boundaries and reconvene after the qualification gates are measured.

Evidence discipline: replace all simulated values with observed, reproducible evidence before using this material as a real product-performance claim.