DECISIONFLOW

Qualify Phase Overview

Evidence synthesis, qualification gates, and conditional scale decision

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-00
Version / status	2.0 — Final Public Portfolio Edition — Simulated Qualification Scenario
Baseline / author	14 July 2026 — Min Thu Kyaw
Role / purpose	AI Product Manager — Qualify Phase Overview
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

Executive Qualification Result

Phase decision — CONDITIONAL SCALE. DecisionFlow demonstrates sufficient simulated evidence to continue investing in the production-ready v1.0 individual decision workflow and to expand carefully inside the beachhead segment. It is not yet qualified for unrestricted acquisition, broad multilingual AI claims, collaboration, integrations, or enterprise packaging.

The qualification review combines the simulated 12-week launch evidence with an illustrative 12-week follow-up period. It evaluates desirability, usability, AI trust, business viability, technical feasibility, retention, and scalability. The strongest signals are activation, workflow completion, durable decision retrieval, AI groundedness, trust, reliability, and early renewal. The weakest signals are long-term habit formation, due-date accuracy, confidence calibration, support-adjusted contribution margin, and performance beyond the tested concurrency range.

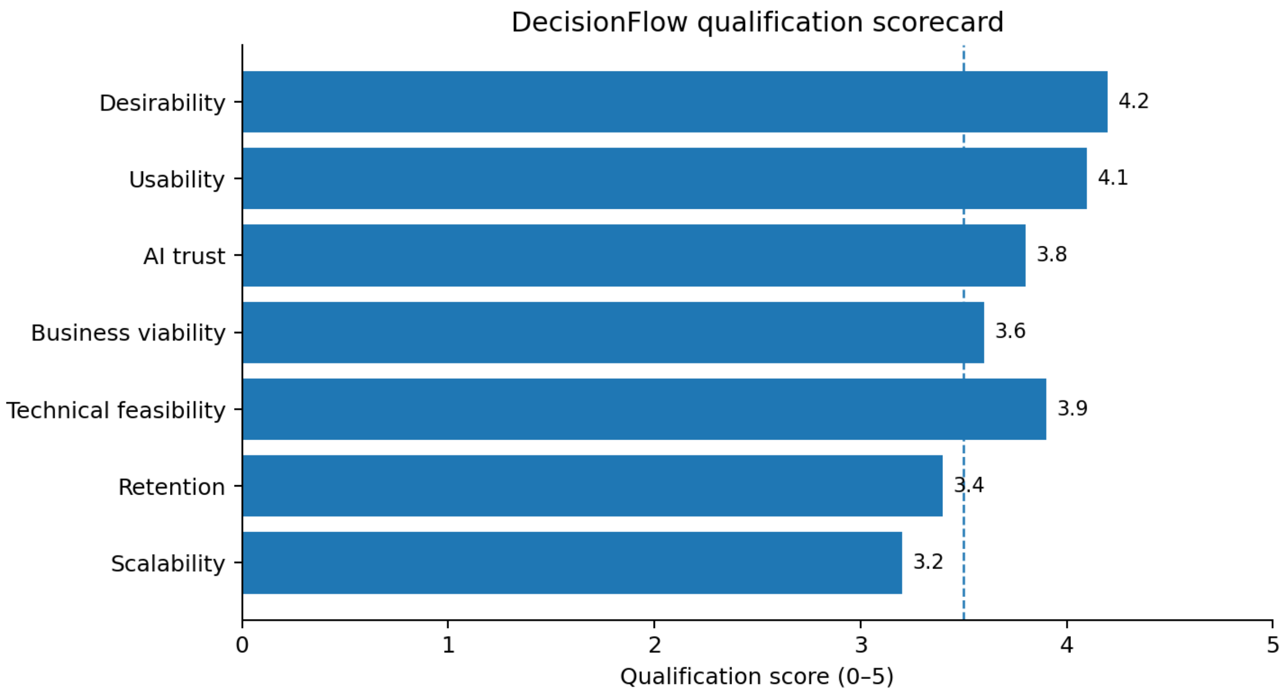

Qualification scorecard

Evidence Model and Metric Reconciliation

Evidence class	Meaning	Use in this phase
Actual artifact	Prototypes, repositories, uploaded MVP baseline, and prior phase documents.	Establishes what was genuinely designed and implemented.
Assumed capability	The release-candidate capabilities explicitly documented in DEV-06.	Defines the simulated v1.0 launch product boundary.
Simulated observation	Customer, AI, financial, operational, and research results.	Demonstrates qualification analysis without claiming real traction.
Recommendation	Product decision derived from the illustrative evidence.	Shows how evidence changes investment and roadmap choices.

Metric audit finding — The launch package reported 43 Week-12 retained users and also stated 49% retention. Using the documented activated cohort of 73, the reproducible activated-user retention is 58.9%. Using registered users as the denominator gives 45.7%. The 49% label is therefore retired. The launch "team participation" metric is also excluded because v1.0 is an individual workflow without implemented collaboration.

Qualification Framework

Dimension	Gate question	Simulated conclusion
Desirability	Does the product solve a high-value, recurring problem?	Qualified in product-management and founder segments; weaker in lower-frequency roles.
Usability	Can users complete the workflow and reach value quickly?	Qualified; 77.7% activation, 75.4% completion, median 8m 42s.
AI trust	Are outputs grounded, transparent, and safely overseen?	Conditionally qualified; core quality passes, due dates and calibration need work.
Viability	Can the commercial model produce sustainable contribution?	Promising but small-sample; renewal and NRR positive, scaled CAC unproven.
Feasibility	Can the product operate reliably and securely?	Qualified within tested load and single-user scope.
Retention	Does the product become a repeated habit?	Moderate; 58.9% Week-12 and 49.3% Week-24 activated-user retention.
Scalability	Can growth increase without unacceptable risk or cost?	Not broadly qualified; async AI work and higher-load validation remain conditions.

Customer Value and Retention

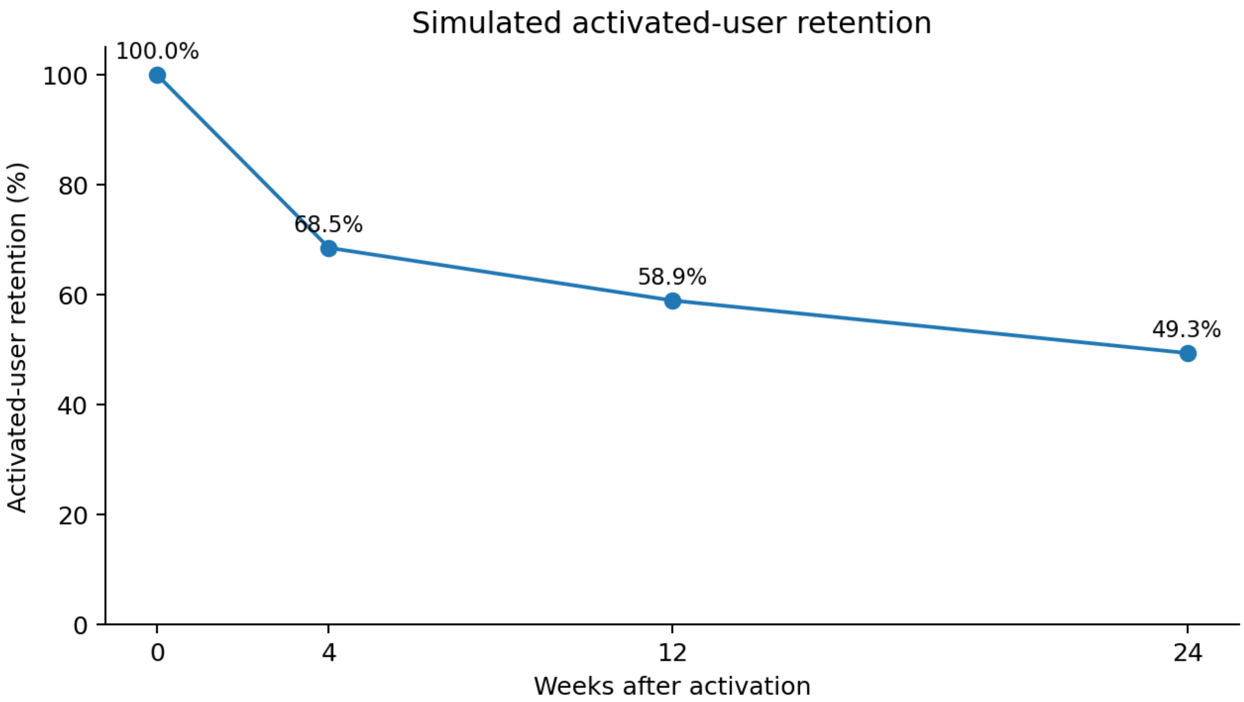

Activated-user retention curve

Signal	Simulated result	Qualification reading
24-hour activation	77.7% (73/94)	Strong first-value signal.
Start-to-finalize completion	75.4% (291/386)	Core workflow is usable and relevant.
Week-12 activated retention	58.9% (43/73)	Moderate repeat value after denominator correction.
Week-24 activated retention	49.3% (36/73)	Meaningful durability, but insufficient for aggressive scaling.
PMF survey	40.5% very disappointed (15/37)	Passes the illustrative 40% gate narrowly; segment concentration matters.
Retrieval adoption	44% of active users weekly in Weeks 8–12	Supports durable-record value beyond initial AI output.

AI Quality and Trust

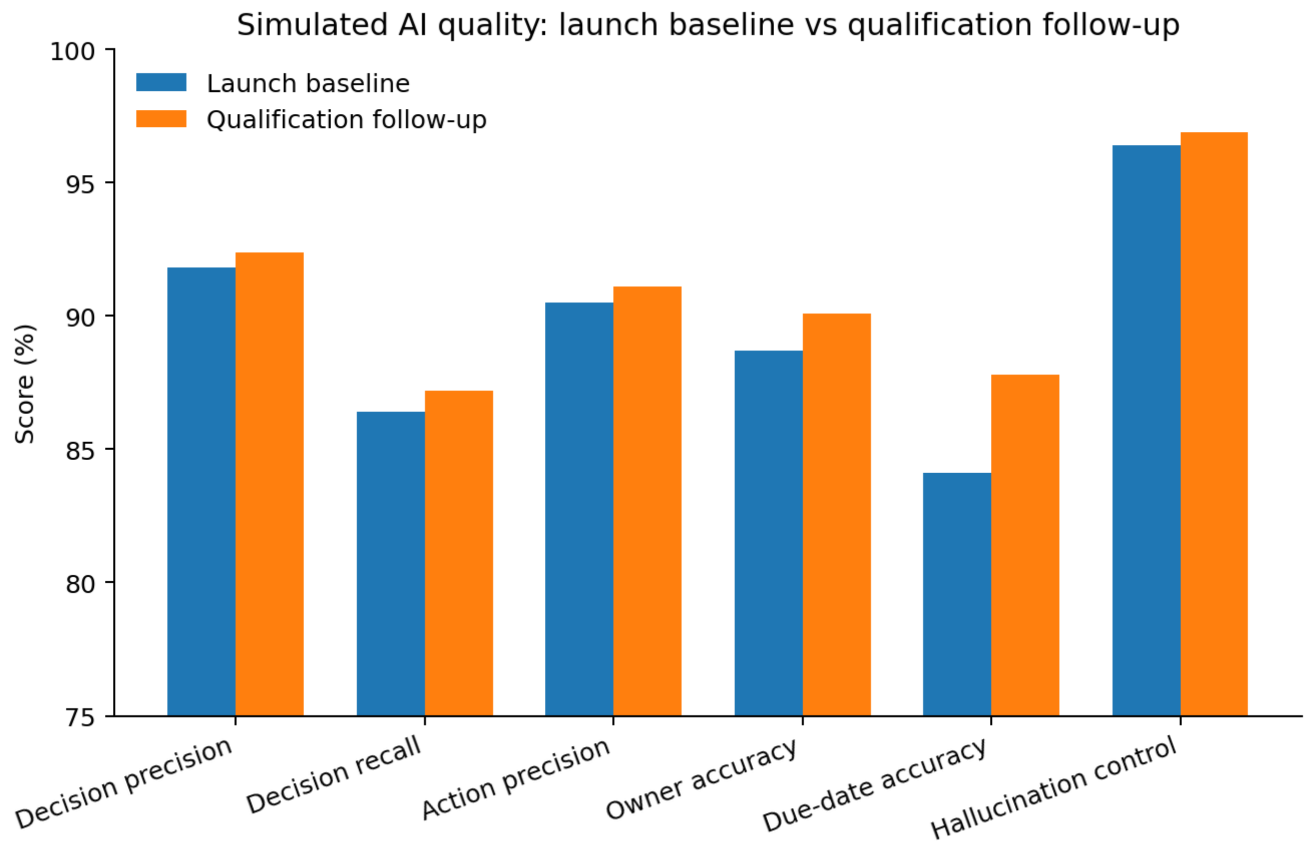

AI quality comparison

Metric	Follow-up	Gate	Decision
Decision precision	92.4%	≥90%	Pass
Decision recall	87.2%	≥85%	Pass
Owner accuracy	90.1%	≥90%	Pass narrowly
Due-date accuracy	87.8%	≥90%	Fail; confirmation remains mandatory
Material hallucination rate	3.1%	≤5%	Pass
Accepted without material edit	78%	≥70%	Pass
Trust rating	4.3/5	≥4.0	Pass
Expected calibration error	0.12	≤0.10	Fail; confidence should be treated cautiously

Business Viability

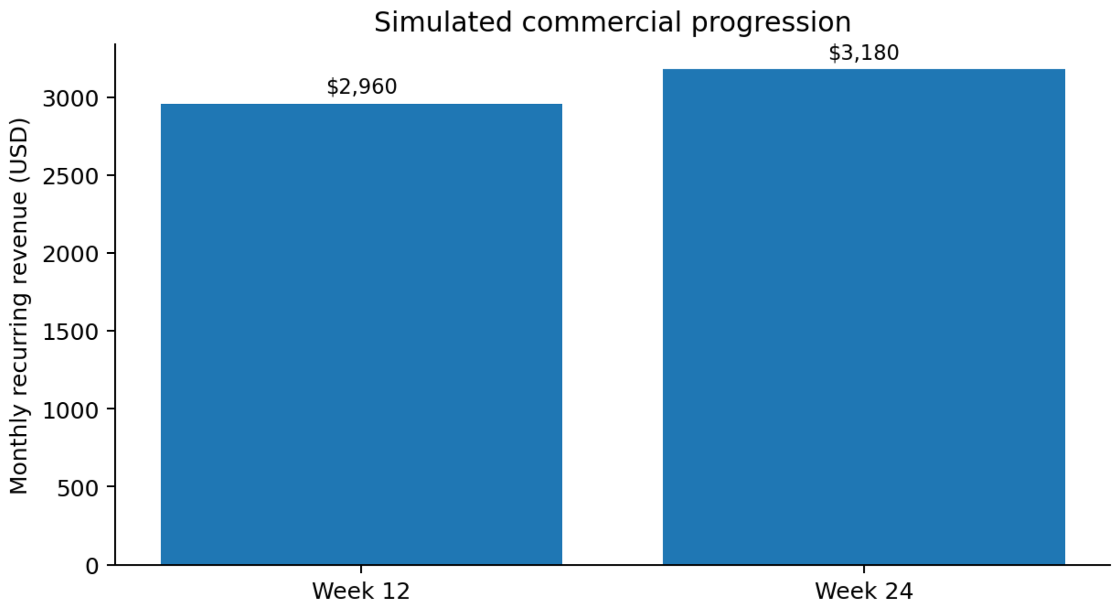

Commercial progression

Measure	Simulated result	Qualification
Original paid organizations retained	7 of 8 by Week 24	Positive renewal signal; sample too small for generalization.
Gross revenue retention	92.6%	Passes a 90% early-cohort floor.
Net revenue retention	107.4%	Expansion offsets one churn; encouraging but concentrated.
Gross margin	82%	Passes initial software margin gate.
Support-adjusted contribution margin	47%	Below 55% scale gate due to guided onboarding and support.
Modeled blended CAC	$2,100	Modelled, not observed at scale.
Modeled gross-profit payback	6.9 months	Acceptable for controlled expansion; channel economics remain unproven.

Technical and Operational Qualification

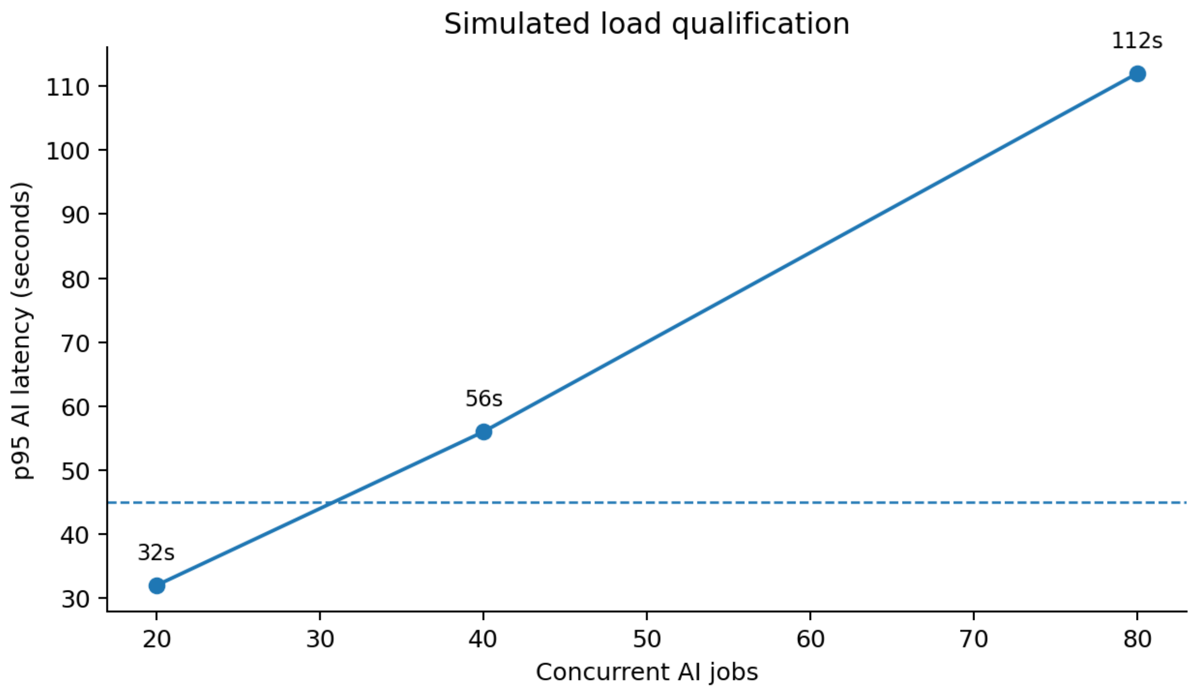

Load qualification

Area	Simulated result	Conclusion
Availability	99.95% during qualification period	Pass
p95 non-AI API latency	390 ms at normal qualification load	Pass
20 concurrent AI jobs	32 sec p95; 0.8% failures	Pass
40 concurrent AI jobs	56 sec p95; 1.9% failures	Fail growth gate
Backup restore	42-minute RTO; 4-minute RPO	Pass
Security	No confirmed unauthorized access; two medium findings closed	Qualified within simulated scope
Provider concentration	Single primary AI provider remains material	Residual high-impact dependency

Residual Risks and Conditions

- Due-date extraction must reach 90% accuracy or remain a suggestion that always requires explicit confirmation.
- Confidence calibration must improve before displaying confidence as a strong decision-quality signal.
- Week-24 retention must improve from 49.3% toward at least 55% in the next qualified cohort.
- Support-adjusted contribution margin must exceed 55% before broad self-service acquisition.
- The platform must pass the 40-concurrent-AI-job load gate with p95 latency below 45 seconds and failures below 1%.
- English is the only commercially qualified AI-content language; other interface languages remain beta until benchmarked.
- Collaboration, integrations, PDF/DOCX ingestion, SSO, and enterprise administration remain evidence-gated roadmap items.

Decision and Next Phase

Investment decision — Scale the individual v1.0 workflow inside the validated beachhead, but cap acquisition, preserve human approval, and fund a 90-day qualification backlog focused on retention, due-date accuracy, calibration, support efficiency, and asynchronous AI processing.

1. Expand controlled acquisition from 15 to a maximum of 30 qualified organizations, not an unrestricted public campaign.
2. Prioritize Product Managers, founders, and functional decision owners with at least two material decisions per month.

3. Keep Starter and Professional packaging; do not market Team or Enterprise capabilities as available.
4. Run a second qualification gate after the 90-day roadmap conditions are measured.

Document Set Map

Code	Document	Primary question
QUAL-00	Qualify Phase Overview	What evidence supports the investment decision?
QUAL-01	Product-Market Fit and Customer Validation Report	Who receives durable value and why?
QUAL-02	AI Quality, Trust and Human Oversight Qualification Report	Is the AI useful, grounded, calibrated, and safely overseen?
QUAL-03	Business Viability and Unit Economics Report	Can the commercial model become sustainable?
QUAL-04	Technical Scalability and Operational Qualification Report	How far can the product safely scale?
QUAL-05	Qualification Risk, Evidence and Assumption Register	What remains uncertain and how will it be tested?
QUAL-06	Qualification Decision Memo and Growth Roadmap	Should the product scale, iterate, pivot, or stop?

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