DECISIONFLOW

Launch Results and Retrospective

Simulated 12-week outcomes, learning, and post-launch strategy

Document code	LAUNCH-06
Version	2.0
Status	Final Public Portfolio Edition — Simulated Launch Scenario
Baseline date	14 July 2026
Author	Min Thu Kyaw
Role	AI Product Manager
Evidence rule	Actual artifacts are identified separately from assumed capabilities and simulated outcomes.

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
Purpose	Document the Launch phase for Launch Results and Retrospective.
Primary audience	Hiring managers, product leaders, launch stakeholders, engineering, customer success, and operations.
Evidence classes	Actual artifact / assumed post-development capability / simulated launch outcome / recommendation.
Scenario assumption	All development P0 release gates are assumed complete before the simulated controlled launch begins.
Ethical use	Do not present simulated customer evidence, revenue, testimonials, or performance figures as real-world results.
Change control	Publication rule: retain evidence labels unless content is replaced by documented real-world observations.

Actual Project Artifacts

- Low-Fidelity Prototype
- High-Fidelity Prototype
- MVP Frontend Repository
- MVP Backend Repository

Contents

Executive Result

Simulated case-study result - DecisionFlow completed an illustrative 12-week controlled commercial launch. The scenario produced positive activation, workflow completion, AI quality, reliability, trust, and paid-conversion signals. The correct decision is to scale carefully, with a stronger focus on repeat use, owner/due-date accuracy, explainability, and evidence-gated post-launch expansion.

Evidence classification - Actual: public prototypes, repository references, uploaded MVP development baseline, and prior Product/Plan/Develop artifacts. Assumed: the production-ready v1.0 core capabilities documented in DEV-06 as completed for the launch scenario. Simulated: user behavior, commercial results, AI benchmarks, incidents, quotations, and launch outcomes.

Scenario Design

Element	Illustrative scenario
Launch model	Controlled cohort followed by incremental expansion.
Sponsoring customer accounts invited	20 technology organizations.
Sponsoring customer accounts onboarded	15 organizations with named champions and approved use cases.
Individual users invited	118 across product, engineering, project, operations, and leadership roles.
Observation period	12 launch weeks after readiness and onboarding.
Data sources	Simulated analytics, AI evaluation, interviews, support, incidents, and commercial pipeline.
Ethical disclosure	All results, quotations, revenue, incidents, and customer evidence are simulated.

Customer and Usage Results

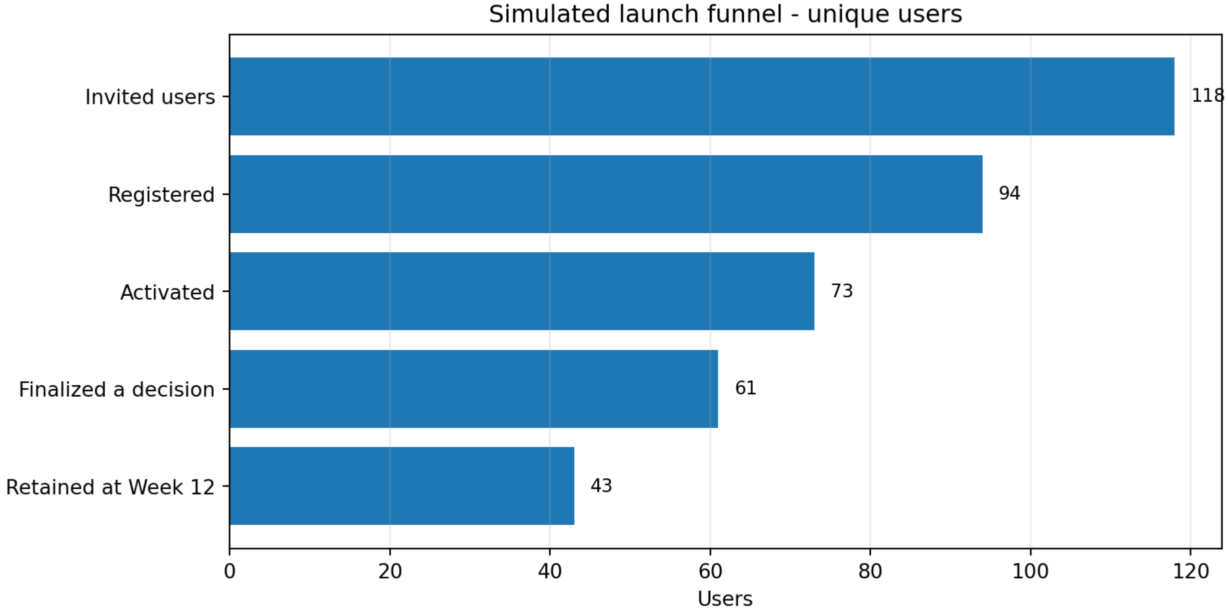

Simulated unique-user funnel from invitation through Week-12 retention.

Metric	Target	Simulated result	Assessment
24-hour activation	>= 70%	77.7% (73 / 94)	Pass
Start-to-finalize completion	>= 70%	75.4% (291 / 386)	Pass
Median time to finalization	<= 12 minutes	8m 42s	Pass
Week-4 activated-user retention	>= 60%	68.5% (50 / 73)	Pass
Week-12 activated-user retention	Exploratory	58.9% (43 / 73)	Needs improvement
Second finalized decision by Week 12	>= 55%	61.6% (45 / 73)	Pass
Search/retrieval usage	>= 40% of active users weekly	44% in Week 8-12	Pass

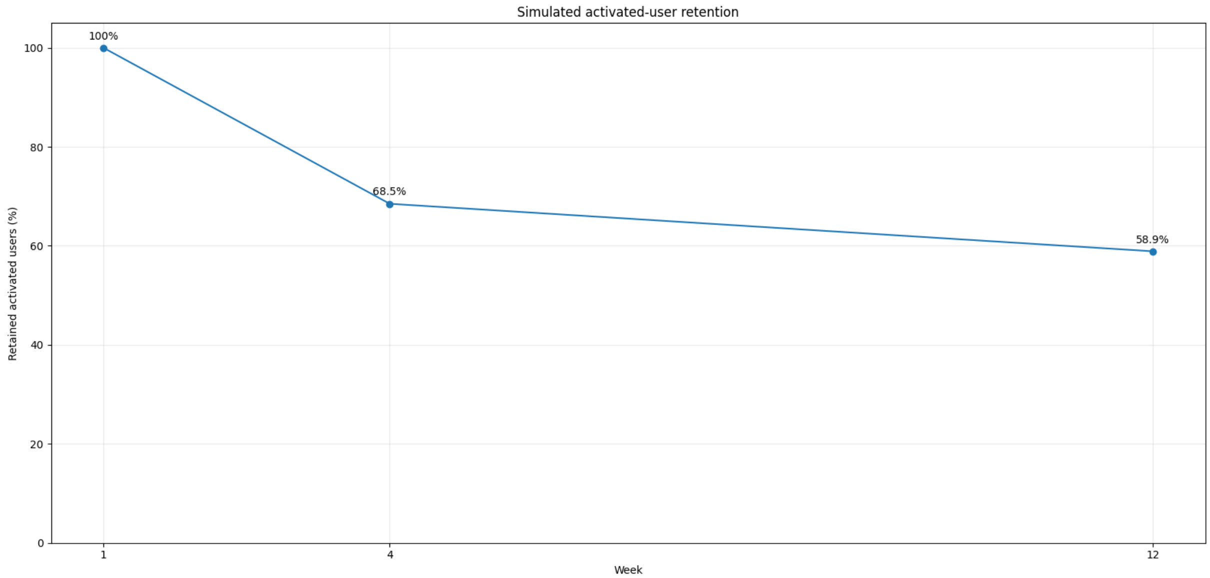

Simulated activated-user retention indicates durable value for about half the cohort by Week 12.

AI Quality and Trust

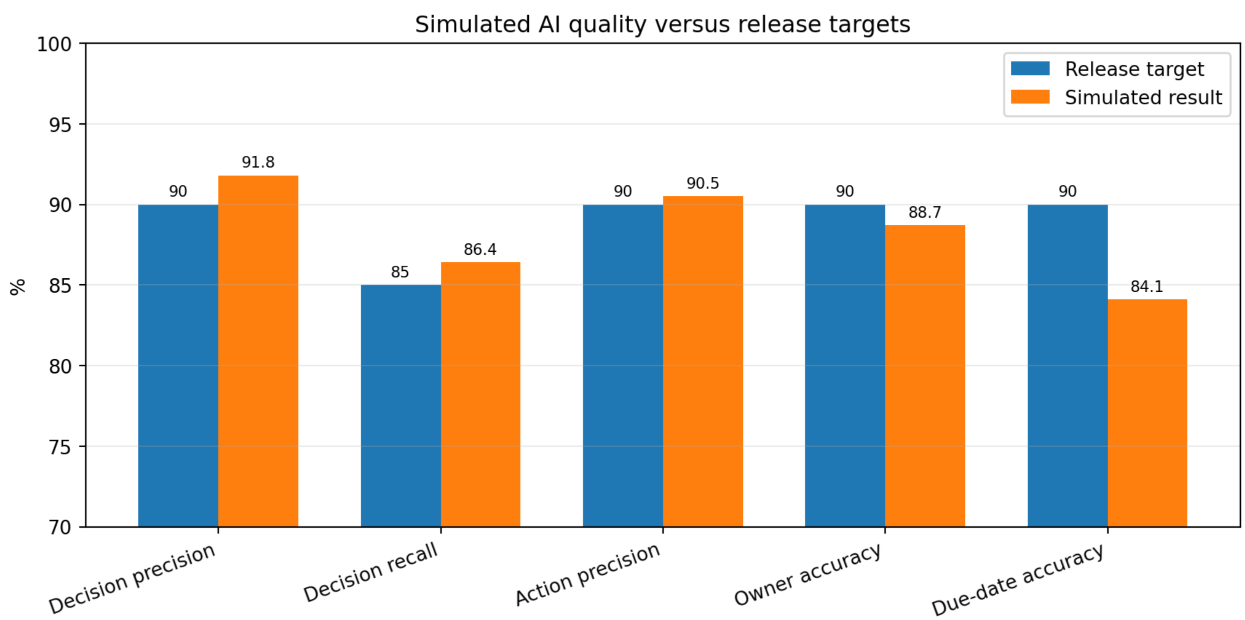

Simulated AI evaluation results. Owner and due-date accuracy miss the original 90% ambition.

Metric	Target	Simulated result	Interpretation
Decision extraction precision	>= 90%	91.8%	Meets release gate.
Decision extraction recall	>= 85%	86.4%	Meets release gate; omissions remain important.
Action-item precision	>= 90%	90.5%	Meets release gate.
Owner accuracy	>= 90%	88.7%	Below target; require confirmation and improve speaker context.
Due-date accuracy	>= 90%	84.1%	Below target; relative dates and milestones are primary failure mode.

Metric	Target	Simulated result	Interpretation
Material hallucination rate	<= 5%	3.6%	Meets gate; severity review remains mandatory.
Accepted without material edit	>= 70%	76%	Strong signal of usefulness.
User trust rating	>= 4.0 / 5	4.2 / 5	Meets gate.
System Usability Scale	>= 80	82	Meets usability target.

Operational Results

Metric	Target	Simulated result	Assessment
Platform availability	>= 99.9%	99.93%	Pass
p95 non-AI API latency	< 500 ms	420 ms	Pass
p95 AI processing latency	<= 35 sec	28 sec	Pass
Failed processing rate	< 1%	0.7%	Pass
Cost per finalized decision	<= $0.60	$0.38	Pass
SEV-1 incidents	0	0	Pass
SEV-2 incidents	Declining trend	4 total; none after Week 7	Acceptable with corrective actions
Backup/restore test	Pass	Passed in scenario	Pass

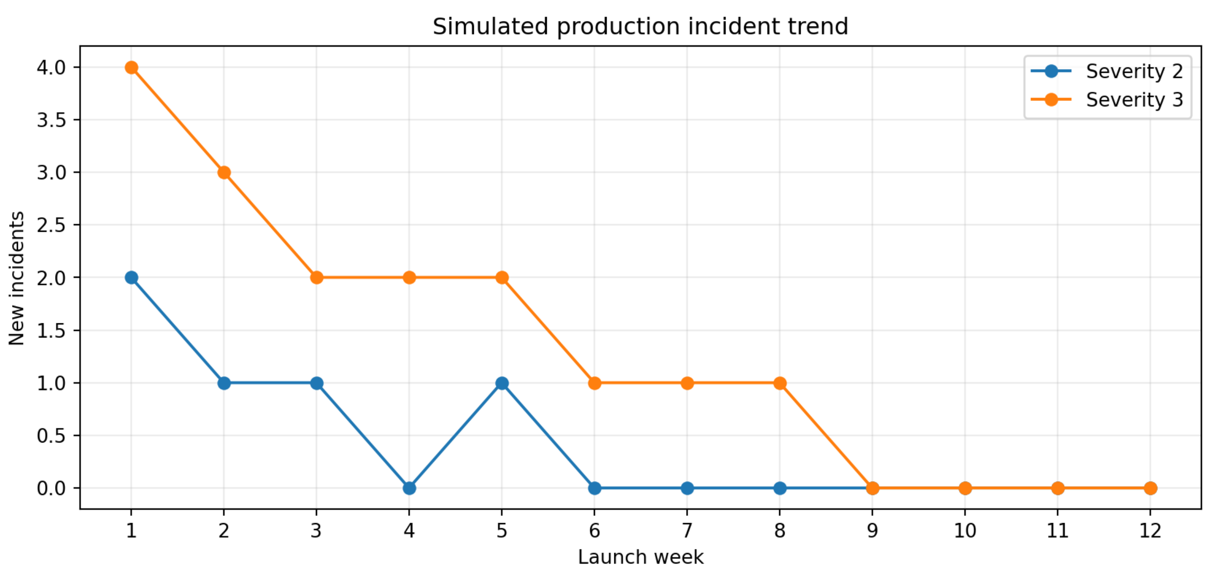

Simulated incident trend shows stabilization after early launch fixes.

Commercial Results

Metric	Simulated result	Interpretation

Metric	Simulated result	Interpretation
Sponsoring customer accounts onboarded	15	Qualified launch cohort.
Organizations converted to paid	8	53% of onboarded organizations; strong but small-sample signal.
Week-12 MRR	$2,960	Illustrative early commercial signal, not actual revenue.
Average paid users	1.6 named users per converted customer account	Users keep separate private libraries; the result does not imply a shared workspace or collaboration capability.
Median procurement cycle	24 days	Security and AI provider review are the main delay.
Expansion interest	5 of 8 paid customer accounts	Requests center on integrations, retention controls, and analytics.
Primary non-conversion reasons	Integration gap, security timing, insufficient repeat use	Roadmap and ICP implications.

Qualitative Feedback

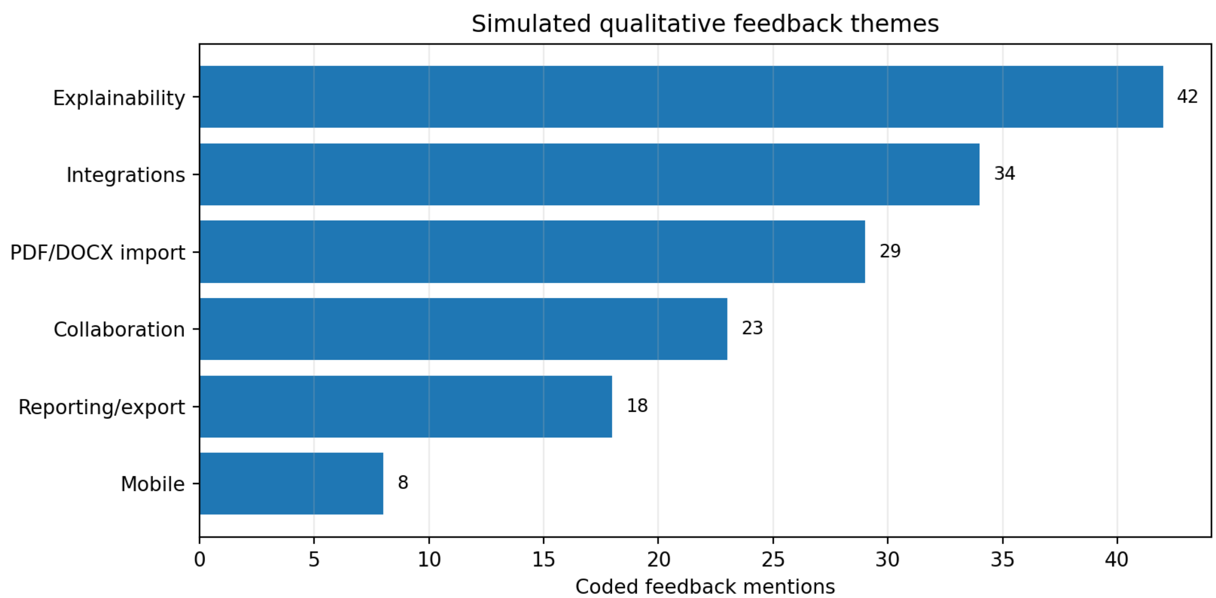

Simulated feedback coding across onboarding, interviews, support, and success reviews.

Illustrative customer quotations

Simulated quote - Product Manager - “The value is not the summary. It is being able to reopen the decision two weeks later and see the evidence, criteria, and owner without reconstructing the meeting.”

Simulated quote - Engineering Manager - “I trust the recommendation more when the system shows what it used and lets me change the weighting. Owner suggestions still need careful review.”

Simulated quote - Founder - “The first decision was easy. The reason I kept using it was the searchable history and the ability to revisit the evidence before the next important choice.”

What Worked

- Decision-first positioning differentiated the product from meeting summaries and created clearer demos.
- Guided onboarding produced a fast first outcome and exposed workflow defects early.
- Explicit human review, source evidence, provenance, and version history increased trust.
- The durable decision record and later retrieval became more valuable than the initial AI output alone.
- Controlled cohorts allowed the team to respond to provider, idempotency, localization, parsing, and indexing issues without broad customer impact.
- Completing a second decision and retrieving earlier records correlated with stronger retention than one-time usage.

What Did Not Work

- Owner and due-date extraction did not meet the original 90% accuracy ambition; ambiguous language and relative dates were difficult.
- Self-service onboarding underperformed guided onboarding for customers without a prepared decision use case.
- Customers expected workplace integrations earlier than the product strategy initially planned.
- Week-12 activated-user retention declined to 58.9% (43/73), indicating that a valuable first decision does not automatically create a recurring habit.
- Security and AI-provider reviews extended commercial cycles for larger accounts.
- The individual package generated the launch evidence; collaboration requests were meaningful but remained unvalidated roadmap demand.

Roadmap Changes After Launch

Priority	Change	Evidence	Expected outcome
P0	Improve owner and due-date extraction with speaker context, relative-date resolution, and mandatory confirmation.	Accuracy misses and user edits.	Higher accountability quality and trust.
P0	Deepen explainability with evidence-first layout and clear score provenance.	Top feedback theme and trust interviews.	Faster verification and fewer inappropriate acceptances.
P1	Validate Jira, Linear, Slack/Teams, and calendar integration demand before building a measured sequence.	Conversion blockers and expansion requests.	Higher repeat use and workflow fit.
P1	Add reusable decision templates and recurring review workflows.	Onboarding friction and retention decline.	More frequent use and faster time to value.
P1	Design organization-level retention, admin, audit, and provider controls for future packages.	Security review and business-package demand.	Shorter enterprise evaluation and higher confidence.
P2	Build executive reporting and cross-project decision health.	Leadership requests; needs discovery to avoid vanity metrics.	Broader organizational value.
Deprioritize	Native mobile application.	Low feedback frequency and desktop-centered work.	Focus resources on higher-value workflow.

Scale Decision

Illustrative Week-12 decision: Scale with safeguards - Expand to additional qualified individual decision owners and carefully tested acquisition channels, but keep account-level onboarding, quality monitoring, and release controls. Do not make aggressive enterprise promises until owner/date accuracy, retention controls, integrations, and security review materials improve.

Next 90 Days

Period	Primary outcomes
Days 1-30	Resolve AI field accuracy gaps; deliver evidence-first recommendation; improve onboarding templates; close incident actions.
Days 31-60	Prototype collaboration and integration concepts; strengthen account security; test second-decision retention.
Days 61-90	Expand controlled acquisition, validate package/pricing, measure cohort retention, and prepare next phase gate.

Product Manager Reflection

Core lesson - AI output is not the product outcome. The outcome is a decision that its owner trusts, finalizes, acts on, and can retrieve later.

Dimension	Reflection
Product	The launch scenario supports the decision-intelligence thesis, but a valuable first workflow is not enough. Repeat use depends on retrieval, a second-decision habit, explainability, and later evidence-gated workflow expansion.
AI	Precision and recall matter, but evidence, uncertainty, user correction, provenance, and safe fallback determine whether quality becomes trust.
Launch	Controlled cohorts, explicit gates, incident discipline, customer education, and measurable learning are part of the product—not administrative work around it.
Leadership	Positive results should not erase unresolved risk. The right decision is controlled scaling, not an unsupported declaration of product-market fit.

This portfolio case study demonstrates how I would connect product strategy, responsible AI, operational readiness, go-to-market execution, customer learning, and post-launch prioritization. In a real launch, every simulated metric, quotation, incident, and commercial result would be replaced with observed, reproducible evidence.