DECISIONFLOW

Product Operations and Governance Playbook

Evidence standards, operating cadence, decision rights, change control, experiment governance, and escalation

DELIVER PHASE | Version 1.0 | Illustrative 90-day controlled-scale 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	DEL-03
Version / status	2.0 — Final Public Portfolio Edition — Simulated Delivery Scenario
Baseline / author	14 July 2026 - Min Thu Kyaw
Lifecycle continuity	Actual MVP -> DEV-06 assumed productionization -> simulated launch -> Conditional Scale -> Deliver phase.
Delivery decision	Execute a 90-day controlled-scale program; keep collaboration, integrations, richer ingestion, and enterprise features evidence-gated.
Evidence rule	Actual, assumed, simulated, target, and recommendation statements are explicitly 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

Governance Purpose and Principles

The Product Operations model ensures that controlled scale remains controlled. Governance is designed to prevent three common failure modes: feature expansion without evidence, AI changes without regression proof, and metric changes that make results appear stronger than they are.

- Evidence before expansion.
- Human authority over final decisions.
- Reversible releases and explicit rollback ownership.
- No customer-content leakage into analytics or support telemetry.
- Metric definitions and denominators are product contracts.
- Discovery is not a promise; simulation is not actual performance.

Operating Cadence

Cadence	Participants	Inputs	Decisions / outputs
Daily delivery sync	Product, engineering, design, QA	Blockers, release health, incidents	Ownership, sequencing, escalation.
Weekly product review	Product, design, engineering, CS	Activation, retention, support, experiments	Backlog and experiment decisions.
Weekly AI quality review	Product, AI/Backend, QA	Evaluation deltas, edits, failures, cost, provenance	Prompt/model promotion, freeze, rollback.
Biweekly risk and assumption review	Product, technical, security	Risk register, scope requests, evidence gaps	Risk treatment and discovery decisions.
Monthly business review	Product, commercial, finance/owner	MRR, retention, margin, support, CAC evidence	Growth cap, pricing, channel and spend decisions.
Quarterly investment gate	Product and accountable sponsor	Benefits, unresolved conditions, roadmap options	Continue, accelerate, correct, pause, pivot, or stop.

Decision Rights and RACI

Decision	Product	Engineering	AI/Backend	Security	Customer Success	Sponsor
Roadmap priority	A/R	C	C	C	C	I
Release go/no-go	A	R	R	C	C	I
Model/prompt promotion	A	C	R	C	I	I
Security incident response	C	R	C	A/R	I	I
Customer communication	A	C	C	C	R	I
Growth-cap change	R	C	C	C	C	A
Scope claim / launch copy	A/R	C	C	C	C	I

R = Responsible; A = Accountable; C = Consulted; I = Informed. A single accountable role must exist for each decision.

Evidence Standards

Evidence level	Definition	Permitted decision
Actual	Public prototype/repository or implementation-derived baseline that can be inspected.	Describe implemented evidence within its observed boundary.
Assumed	DEV-06 case-study capability declared complete for lifecycle continuity.	Use only inside the illustrative scenario with disclosure.
Simulated	Invented cohort, financial, AI, delivery, or operational result.	Demonstrate decision logic; never claim as real performance.
Strong observed evidence	Reproducible, longitudinal, representative, and independently reviewable.	Support scale within tested boundary.
Moderate evidence	Consistent but limited by sample, duration, or concentration.	Support controlled investment and further

Evidence level	Definition	Permitted decision
		testing.
Weak / missing	Modeled, anecdotal, invalid, or unavailable.	Discovery only, reduce claim, or stop.

Intake and Prioritization

1. Capture the request, requestor, customer job, segment, and current workaround.
2. Classify the request as defect, compliance/security need, qualification condition, optimization, experiment, or discovery.
3. Identify expected metric movement, affected data, AI behavior, operating cost, and dependency risk.
4. Score evidence strength and reversibility before estimating delivery.
5. Choose deliver, experiment, discover, defer, or reject and record the reason.
6. Publish the decision and review date in the decision log.

Example request	Classification	Governance decision
Add PDF and DOCX import	Discovery / new parsing and privacy surface	Do not deliver in the program; research demand and safe parsing.
Add shared workspaces	Strategic discovery / multi-user authorization	Interview and prototype only; no release claim.
Improve due-date accuracy	Qualification condition / AI accountability	P0 delivery.
Add another dashboard chart	Optimization with weak evidence	Defer unless linked to a decision or support need.

Experiment and Release Governance

Control	Required definition
Hypothesis	Specific behavior or business outcome expected to change and why.
Population	Eligible segment, exclusions, allocation and sample limitations.
Primary metric	Single decision metric with numerator, denominator, time window and data source.
Guardrails	AI quality, privacy, failure, support, cost and user-control metrics that may not deteriorate.
Stop condition	Safety, quality, cost, operational or evidence threshold that ends the test.
Decision rule	Ship, iterate, revert, or learn without shipping.
Release evidence	Tests, monitoring, runbook, rollback, communication and owner.

Risk and Escalation

Severity	Example	Response expectation
Critical	Confirmed unauthorized access, secret exposure, high-severity harmful AI output.	Stop affected processing, activate incident command, notify accountable owner immediately.
High	Systemic lost jobs, materially incorrect billing, sustained quality regression.	Freeze release or model, assign owner, communicate workaround, review within same business day.
Medium	Localized workflow failure, elevated support demand, non-critical metric loss.	Triage in weekly operating cycle with dated corrective action.
Low	Copy issue, cosmetic defect, low-impact enhancement.	Backlog and prioritize by evidence.

Escalation rule - Any event that challenges the current scale boundary, evidence integrity, customer data protection, or human control is treated as a product-governance issue, not only a technical defect.

Documentation and Metric Control

- Each release updates product scope, architecture, AI model/prompt versions, analytics contracts, runbooks, limitations, and customer communications.
- A metric definition includes purpose, event source, numerator, denominator, window, exclusions, owner, and known limitations.
- Historical results are never silently recalculated. Corrections publish both values and the reason.
- Customer content is excluded from analytics and standard support logs.
- Portfolio documents retain the case-study disclosure and separate actual, assumed, and simulated evidence.

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