

PRD Summary: Adaptive Workplace AI Retention Layer

A concise product-management summary of the Adobe Life AI concept: the problem, users, solution, and technical product-management relationship behind the Active-Shift retention system.

Product area	Adaptive workplace AI assistant for task support, learning, and workplace flow.
Core feature	Active-Shift logic: user-declared modes plus behavior-based adaptation.
Primary constraint	Design within strict legal and ethical boundaries, including avoiding emotion inference.
TPM focus	Translate ambiguous AI strategy into scoped signals, launch phases, metrics, and compliance-safe implementation logic.

Problem

- Workplace AI can become either too generic to retain users or too invasive when it tries to infer internal states such as emotion, stress, or motivation.
- The product challenge was to design adaptive assistance that could improve activation and retention without relying on prohibited or high-risk inference patterns.
- The PM challenge was turning broad AI-personalization ambition into measurable, legally constrained product behavior.

Users

- Primary users: knowledge workers who need AI support for task execution, prioritization, learning, and workplace navigation.
- Secondary users: team leads and enterprise buyers who care about adoption, trust, compliance, and repeat usage.
- Internal users: product, legal, data, and engineering teams who need clear signal boundaries before building adaptive AI features.

What We Did

- Defined activation metrics and user-declared interaction modes so personalization starts from explicit user choice rather than hidden inference.
- Designed Active-Shift logic, a retention mechanism that adapts the product based on behavioral signals like usage patterns, mode selection, and task context.
- Sequenced the feature launch into a three-part rollout so the product could start with safe onboarding, then add adaptive behavior, then expand retention loops after validation.

Solution

- Use declared modes as the personalization anchor: users choose how they want the AI to support them, and the system adapts within those stated boundaries.

- Use behavioral signals instead of emotion inference: the system watches interaction patterns and task flow, not sensitive emotional or psychological states.
- Pair the retention layer with compliance-aware metrics: activation, mode adoption, repeat use, and task-completion signals become the product learning loop.

TPM Relationship

- Product strategy: clarified the user problem, adoption motion, retention mechanics, and success metrics for an AI workplace product.
- Technical scoping: converted the strategy into implementation-friendly signal categories, system boundaries, and launch phases.
- Risk management: aligned adaptive AI behavior with legal constraints by excluding emotion inference and keeping personalization explainable.
- Cross-functional execution: bridged PM, legal, and engineering needs through a PRD-style framework that defines users, signals, guardrails, and rollout logic.

PRD Notes

Success should be evaluated through explicit-mode adoption, task-completion improvements, repeat usage, user trust signals, and compliance review readiness. The recommended MVP keeps the adaptive system narrow: declared modes, behavioral telemetry, transparent user controls, and clear launch gates.