

LimFlex

AI Vendor Adoption Support

Keep your roadmap. Move adoption forward.

An external wrapper concept for more adoption-ready AI responses while keeping the existing AI product intact.

Adoption support without rebuilding.
LimFlex is positioned as an external layer around existing AI workflows, not as a core-product replacement or core-model rebuild.

For AI vendors	External layer	Non-confidential start
Reduce enterprise adoption hesitation without forcing a rebuild of the core AI product.	Support response fit, reviewability, escalation, evidence, and return-to-baseline behavior.	First contact should stay public and non-confidential; deeper review can proceed separately.

Where it fits

Core product remains the baseline	LimFlex is evaluated as an outer wrapper, not as a replacement for model roadmap, product roadmap, or existing differentiation.
Adoption-critical outputs get extra structure	Candidate outputs can be routed toward proceed, re-check, HOLD, escalation, or return before operational commitment.
Evidence supports later review	The wrapper concept emphasizes reviewability and reconstruction, helping teams explain why an output was used, held, or escalated.

Public prototype observations

These numbers are not presented as product claims. They are evaluation-entry signals only: not production benchmarks, guarantees, or safety claims.

99.98%	3.9x
Observed issue-screening signal in a public prototype.	Observed review-loop convergence signal in a public prototype.

Suggested first question

Could an external adoption-support wrapper help our product become more reviewable, usable, and enterprise-fit without changing our core roadmap?

Contact: koji.mochizuki@limflex.org | Please keep first contact non-confidential.

Public overview only. Patent-pending technology. Independent Researcher, Japan.