

LimFlex for AI Product Vendors

Public, non-confidential technology overview

Pluggable judgment-function wrapper for reviewable decisions and audit-ready outputs.



Can AI vendors reduce customer adoption friction while keeping the model, UI, data workflow, and agent stack intact?



Host product retains release / execution control.

Not only logging: the main value is pre-release judgment output.

1 What it is

- Pluggable, non-invasive wrapper around existing AI/data workflows
- Attaches before responses, recommendations, edits, approvals, routing, or actions
- Keeps the host model, UI, workflow, and agent stack intact

2 What it does

- Improve output quality by sharpening retrieval/evaluation based on intent and context
- Help filter unsupported or hallucination-prone outputs
- Check proposed outputs/actions against permissions, policy, effective access, and audit/logging requirements

3 What it returns

- Allow, defer, restrict, block, or route to human review
- Attach evidence, reason codes, or audit-ready explanation
- Return control before release or execution when needed

4 Where it matters

The strongest fit is high-impact agentic workflows where AI outputs may become:

- user-facing responses or recommendations
- edits, approvals, or routing decisions
- workflow actions or business-state changes

Focus: trust, governance, permissions, auditability, accountability.

5 Why it can stay thin

Accumulated decisions, evidence, deferrals, reason codes, and review outcomes can become reusable governance assets.

Prior cases become reusable governance references; the active wrapper remains thin.

6 Integration and public-use boundary

API-first, low-intrusion path

Host product sends output/action, intent/context, evidence, permission metadata, policy context, and audit/logging requirements.

Enterprise alignment

Deployment depends on IAM, policy decision/enforcement, logging, auditability, and ownership across Security, Governance, and Platform teams.

Public boundary: Not a product sales page. Initial public, non-confidential discussion only before any commercial step.

Do not send: confidential information, paid requests, pricing proposals, licensing terms, or commercial conditions.

Evaluators: Security / Governance / Trust-Safety / Product / Platform / Data Infrastructure / Agent Infrastructure

Public, non-confidential technology overview. Not a product sales page.

No confidential evaluation materials requested.

Technical foundation: eight-layer reference architecture for boundary-aware AI workflow control.

Koji Mochizuki | Independent Researcher, Japan

limflex.org | koji.mochizuki@limflex.org