

# Adam

Runtime governance for LLMs

MrF X Nexus

Riyadh, Saudi Arabia

info@mrf103.com · www.mrf103.com/adam

A model-agnostic runtime governance layer for any organization that deploys LLMs and needs to prove which model answered, that it wasn't swapped, and that it stayed within policy.

## THE PROBLEM

Organizations putting LLMs into production often can't prove which model actually answered a request, whether it was swapped or altered behind the endpoint, or whether it stayed within policy at runtime. Auditors and risk teams ask for that proof; application logs alone don't provide it.

## HOW IT WORKS

|                                                                                                                    |                                                                                                                    |                                                                                                                          |                                                                                                         |
|--------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| 01<br><b>Seal</b><br>Every weight file is fingerprinted into a sealed identity; the running process must match it. | 02<br><b>Gate</b><br>Each request passes access, identity, isolation and egress checks before the model is called. | 03<br><b>Record</b><br>Delivered, withheld and blocked exchanges go to an append-only ledger and a signed evidence pack. | 04<br><b>Verdict</b><br>Nine indicators roll into one daily go / no-go; one red blocker closes the day. |
|--------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|

## CAPABILITIES SHOWN ON THE RUNNING SYSTEM

- **Sealed model identity** — fingerprint over all weight files; confirms the live process is the sealed model, sandboxed, with zero outbound connections.
- **Runtime gate** on every request, with each stage timed and its reason stated.
- **Claim check** — an answer contradicted by live evidence is withheld from the caller and kept in the ledger.
- **Swap detection** — altering one byte of the weights blocks the next request before the model is called.
- **Evidence pack** — SHA-256 and ed25519 signature, verified and exported; mapped to EU AI Act record-keeping articles (a mapping, not a conformity assessment).
- **Release monitor** — one daily verdict across nine indicators, three of them blockers.

## DEPLOYMENT

- Runs in your own environment, next to your inference stack.
- Governs open-weight models served behind OpenAI-compatible endpoints.
- **Wired today:** MLX on Apple silicon, on-premises — AceGPT-v2-32B (4-bit, with a fine-tuning adapter) cleared for service; ALLaM-7B sealed and registered, not cleared.

## STATUS

### PRODUCT

Running system; the demo is a screen recording of it (2026-10-20).

### IP

Patent-pending (filed with SAIP): SA 1020268315 · SA 1020268588.

### MEASUREMENT

Self-measured; no third-party audit yet.

## BUILT FOR

### Regulated enterprises

Finance, healthcare and telecom teams that must show auditors which model answered and that it stayed within policy.

### Public-sector IT

Departments that need a signed, append-only record of every model exchange, including the blocked ones.

### Teams shipping LLM features

Product teams whose LLM features have to pass an audit, with evidence mapped to record-keeping requirements.

## CONTACT

Request a demo at [www.mrf103.com/adam](http://www.mrf103.com/adam) or write to [info@mrf103.com](mailto:info@mrf103.com).

MrF X Nexus · Riyadh, Saudi Arabia

### Watch the 2-minute demo

Scan to open the product page and the recording of Adam running.

