



• AUTHORIZATION AT AI RETRIEVAL

Securing AI Decisions, Before Data Is Accessed

A trust layer between your AI applications and your data. For every request, Shield decides whether the person behind the AI is allowed to see the data being fetched — before it is fetched — and records why.

THE BLIND SPOT

At rest

✓ covered

In transit

✓ covered

At login

✓ covered

In use

X unguarded

Legacy security protects data at rest, in transit and at login. It cannot protect data in use — while the AI reasons. That is the gap Shield closes.

63%

of enterprises hit an AI security incident in the past year.
Gartner, 2024

\$4.88M

global average cost of a data breach. IBM, 2024

THE DIFFERENTIATOR

Authorization-Aware Graph RAG

Traditional RAG asks “what is relevant?”. Shield also asks “what is this user allowed to see?” — so only authorized evidence ever reaches the model, based on the policies you configure.

Prevention, not detection. The decision happens before data is fetched — not after the answer is written.

THE PLATFORM

Three components, one governed layer

① Local AI Runtime

A banking-tuned small language model (IBM Granite) that runs entirely on your hardware — no prompt or data leaves your network.

OUTCOME Data sovereignty · air-gapped

② Authorization-Aware Graph RAG

Per-user, relationship-aware retrieval over your entitlement graph — only evidence the user is authorized to see reaches the model.

OUTCOME Explainable, least-privilege retrieval

③ Agentic Guardrails

An AI agent inherits the permissions of the user it acts for; every retrieval and action is validated before it runs.

OUTCOME Safe, policy-validated AI actions

Remove any one and the trust gap re-opens. Together they form a deterministic, auditable AI execution layer — running inside your perimeter, additive to the IAM, DLP and SIEM you already own.

WHAT WE COVER — HONESTLY mapped to the OWASP LLM Top 10

✓ WE OWN (2)

Excessive Agency (the “confused deputy”) and sensitive-data over-retrieval — the two most likely to leak regulated data.

🔍 WE REDUCE (6)

Prompt injection · output handling · tool misuse · prompt-leakage · resource abuse · over-reliance. Contained, additive.

✗ NOT US (2)

Supply-chain and data poisoning — these stay with your AppSec and data-pipeline teams.

WHY BANKS TRUST IT

Additive, not rip-and-replace

Extends your IAM, DLP and SIEM at a point they don't reach.

Fail-closed by design

Hard violations blocked; a timeout fails open only with a flag.

Provable

Every decision policy-cited, in an append-only hash-chained log.

Sovereign

Self-hostable & air-gapped, on open-weight IBM Granite.

Managed — in-region VPC

Sovereign — on-prem / air-gapped

Compliance: helps you meet RBI, DPDPA and PCI-DSS controls — a support for your compliance programme, not a guarantee of compliance.

See it on your own data.

A governed pilot reaches a go / no-go decision in about four weeks — on your infrastructure, with your policies.