

Verillian for Government

FedRAMP-aligned, FISMA

Agencies query their own data in plain language. The data belongs to citizens who entrusted it by law. Verillian governs AI use across the agency with evidence an inspector general can verify.

276

Ransomware attacks on government entities in the first nine months of 2025, up 41 percent year over year, exposing 443,522 records. Comparitech Government Ransomware Roundup, 2025.

WHAT'S AT STAKE

When agency staff query citizen data through AI tools with nothing governing them, no enforced boundary stops sensitive records from leaving the agency, and no verifiable evidence shows what was sent or returned. If an inspector general asks what the AI saw and did, an agency with nothing in place cannot produce a tamper-evident answer.

HOW VERILLIAN ANSWERS

Stop citizen data at the boundary

Sensitive-data detection redacts or blocks SSNs, medical record numbers, and other configured types before a request leaves the agency boundary, so data citizens entrusted by law does not reach an outside provider unreviewed.

Evidence an inspector general can verify

Every AI interaction is signed on the device and hash-chained into an append-only, tamper-evident record, giving oversight bodies non-repudiable proof of what the AI was asked and what it returned.

Deny by default, fail closed

Policy is enforced at the moment of execution. A missing policy or an audit-pipeline failure stops AI traffic rather than letting ungoverned queries through.

Self-hosted, with framework-aligned retention

Built for FedRAMP-aligned and FISMA environments. A zero-knowledge server stores only ciphertext under the agency's own key, Verillian retains none of the agency's interactions, and retention aligns to the governing framework.