AI SECURITY AUDITOR:

Real-Time Threat Detection for OpenEdge Applications

Carla-Noemi Spáczai
Developer

Wayfare



HI!
I'M CARLA.

Developer at Wayfare.

Wayfare

ABOUT WAYFARE

We've supported businesses of different sizes and across multiple industries to achieve their goals.

TOP TALENT, SHARED KNOWLEDGE, CLOSE COLLABORATION

Our cost-effective solutions and active community engagement foster innovation and efficiency.

We offer tailored services including Modernization, New Development, Maintenance & Support, Quality Assurance, UI/UX Design, Consultancy and Project Delivery.



OUR EXPERTISE

PROGRESS
OpenEdge, CorticonJS, Sitefinity

MOBILE React Native

DEVOPS Docker, Kubernetes, Jenkins

MANUAL AND AUTOMATION TESTING

WEB TECHNOLOGIES
Angular, React, Vue.JS, NodeJS, JS, TypeScript, Java, .Net

CLOUD AWS, Azure, Google Cloud Platform

UI/UX Figma, XD, Zeplin, Illustrator

LOW CODE/ NO-CODE FRAMEWORKS

Wayfare

AGENDA

The Challenge: Security in OE Applications
Solution Architecture
Collecting Events from ABL
Machine Learning Detection
Explainability with Local AI (LLM)
Live Dashboard & Compliance Reporting
Live Demo & Code Walkthrough
Results & Conclusions
Q&A

Wayfare

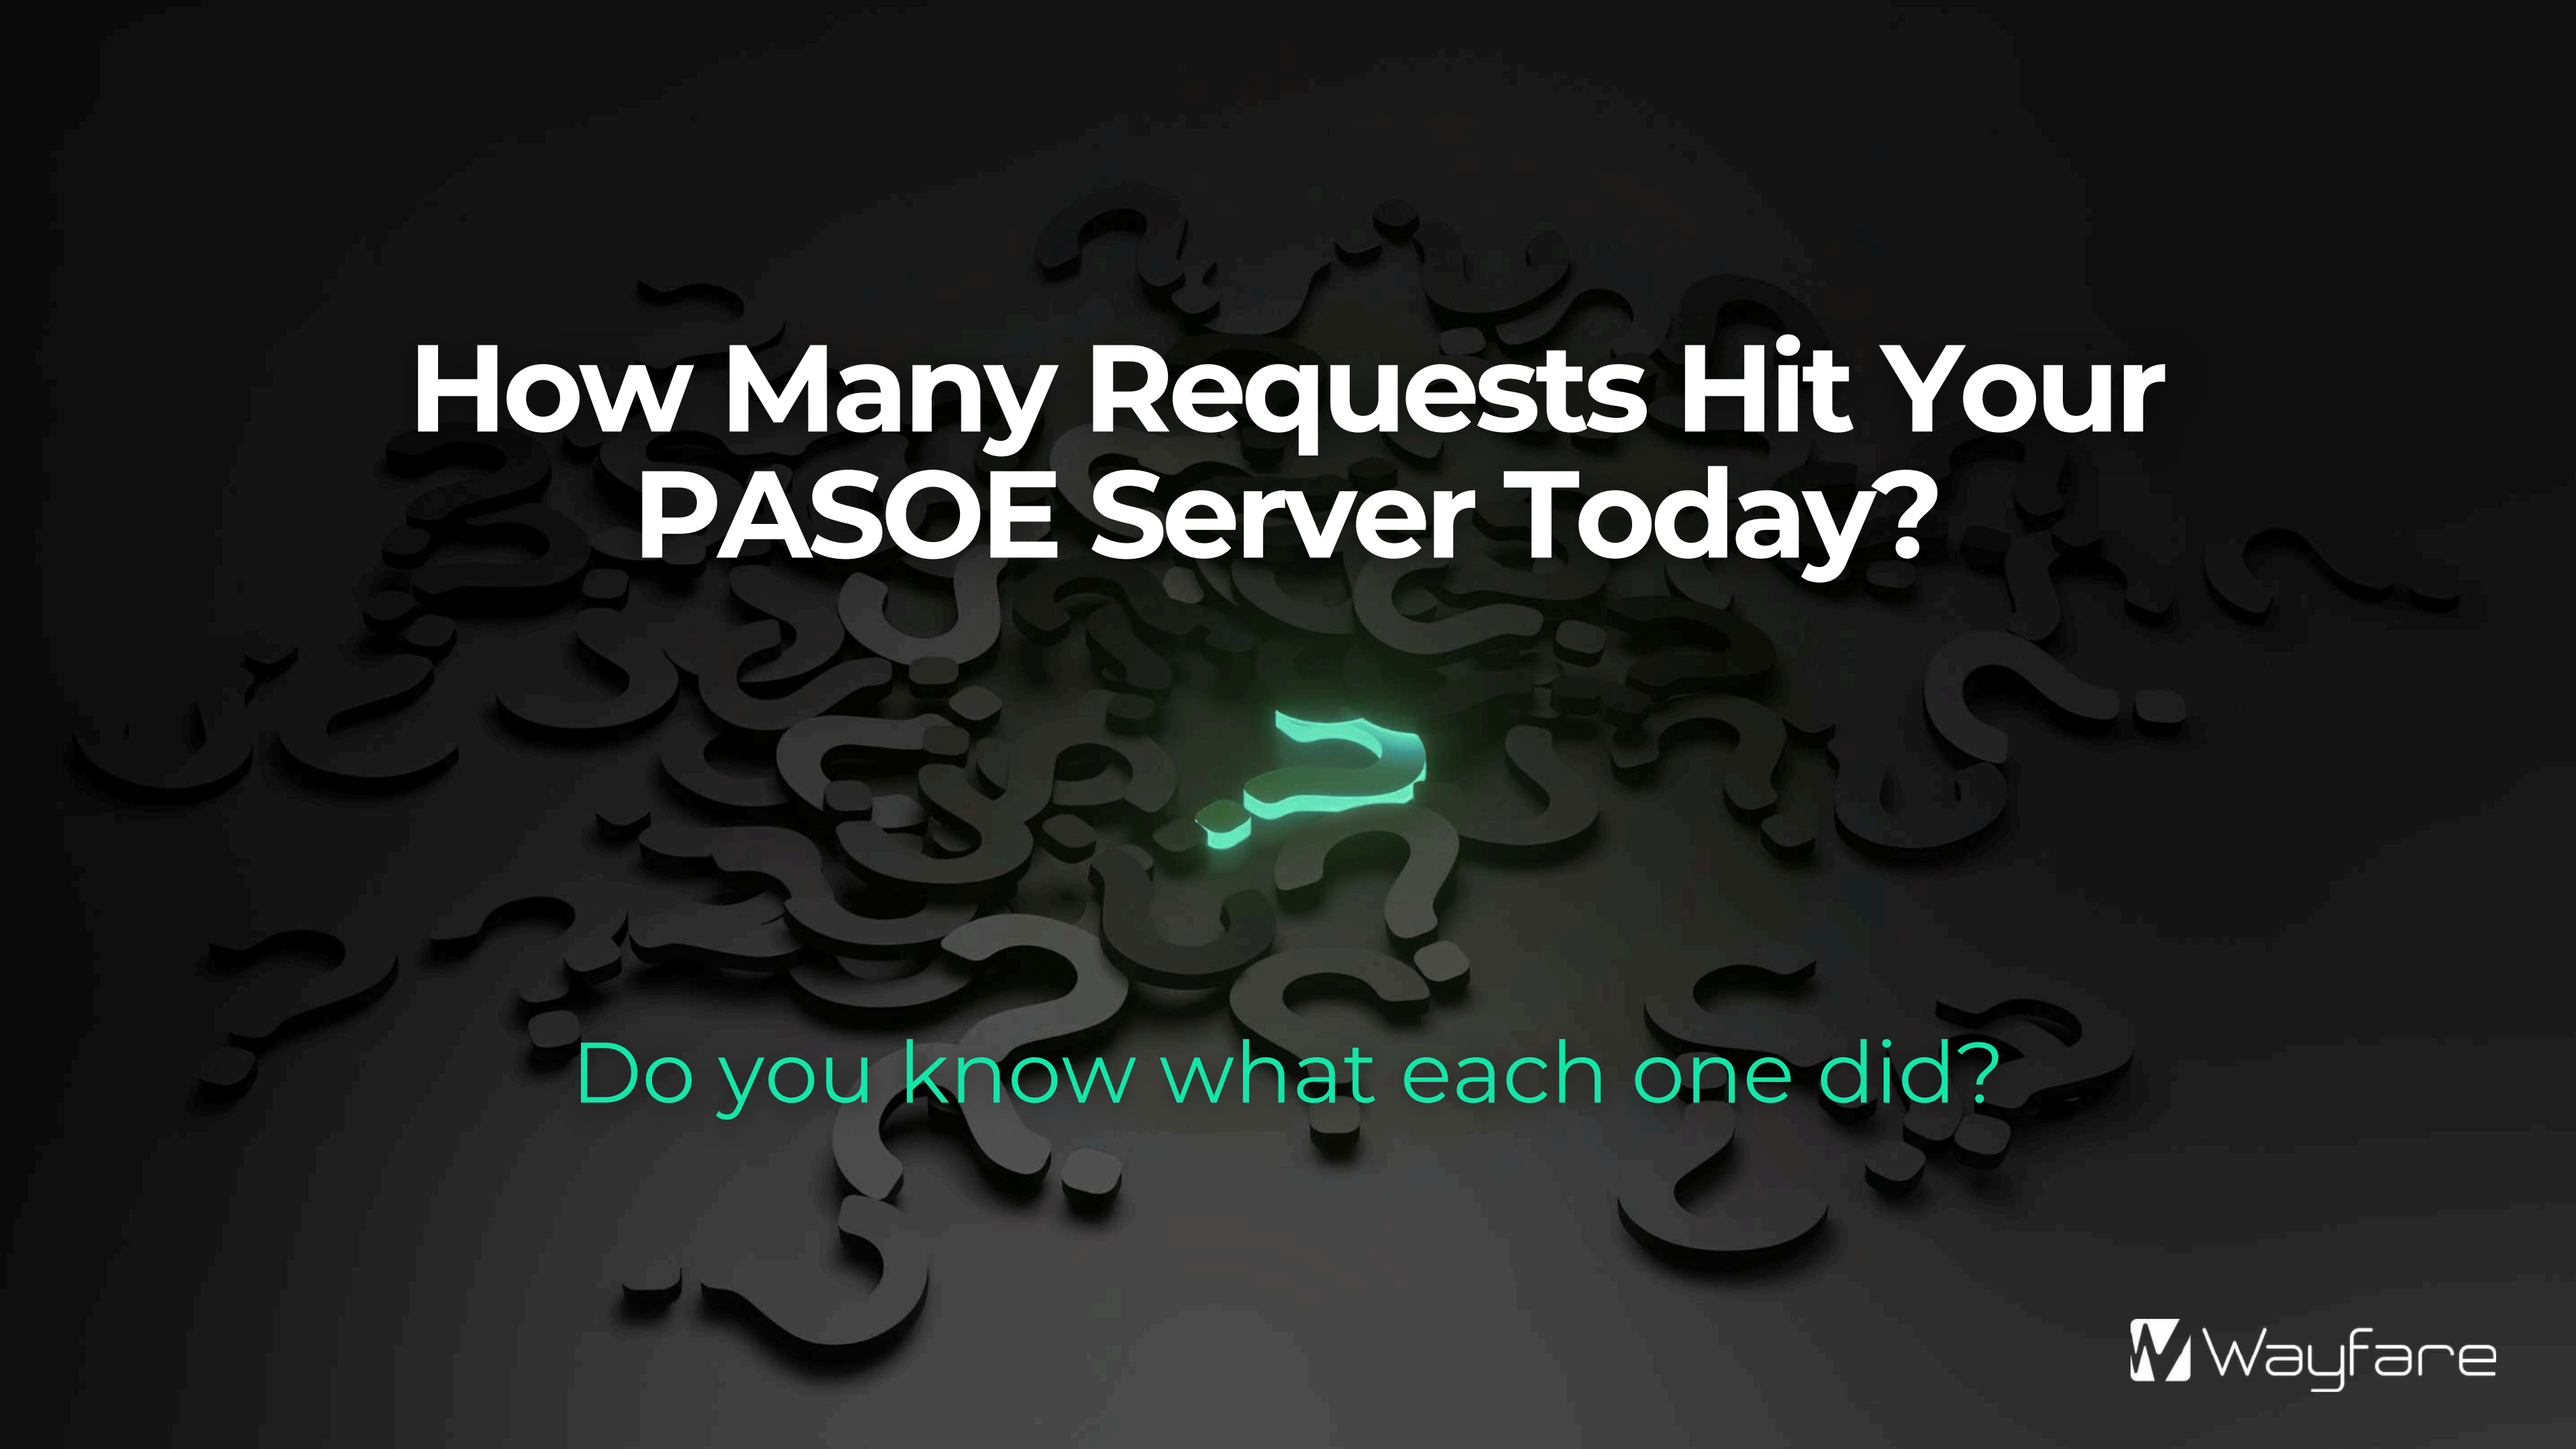
THE CHALLENGE

Security in OpenEdge Applications



Wayfare

How Many Requests Hit Your PASOE Server Today?

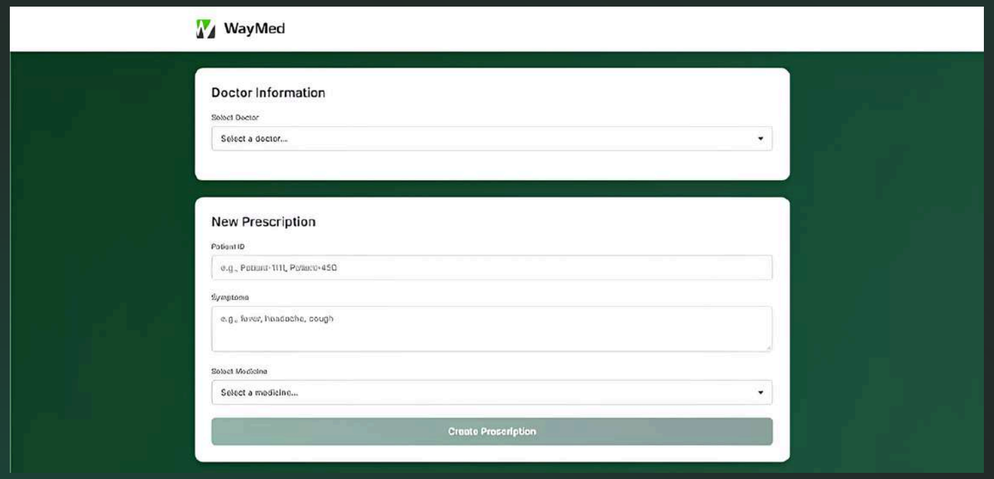


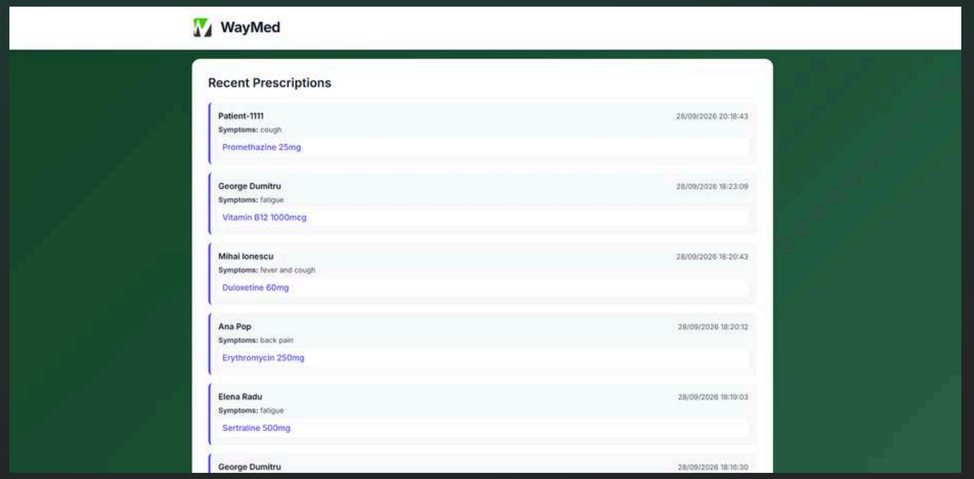
Do you know what each one did?

Wayfare

Meet the Application

A model of a real OpenEdge system: doctors, medicines, prescriptions



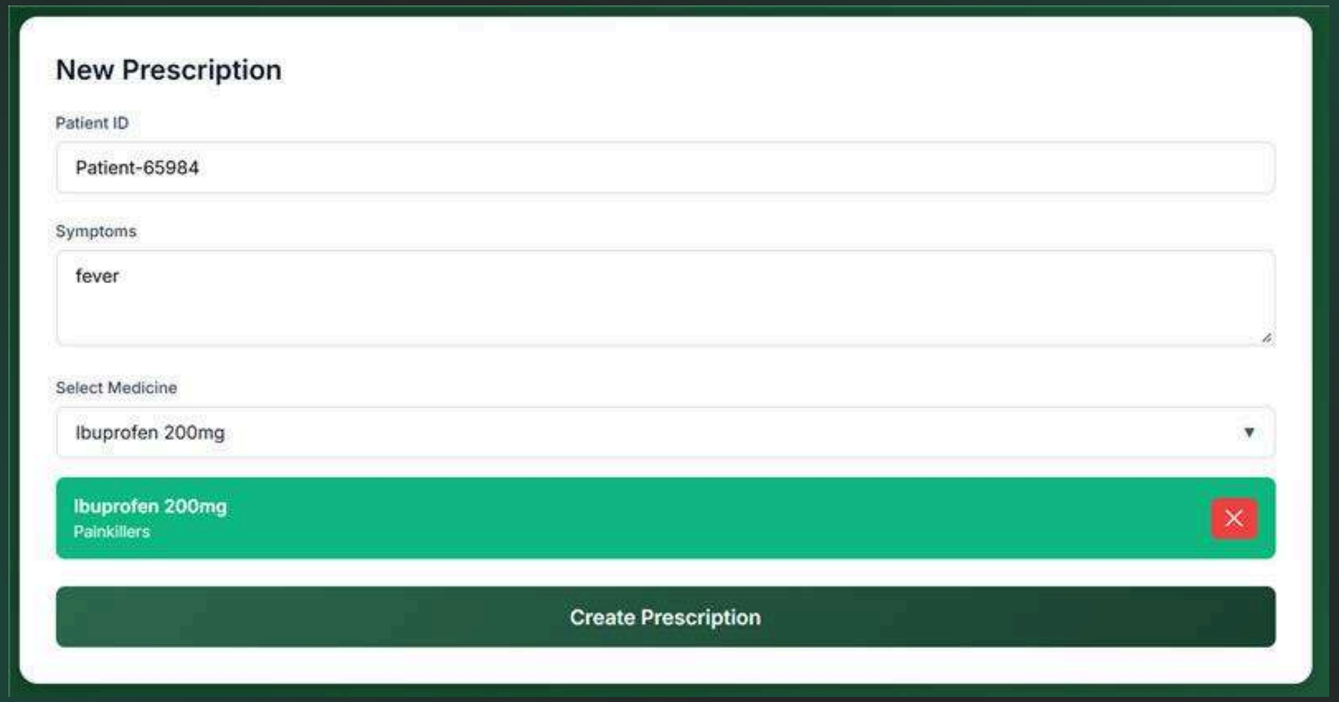


Built entirely in ABL, exposed through PASOE WebHandlers

Everything you're about to see runs on this app

Wayfare

A Normal Day



A doctor searches and prescribes. A WebHandler responds.

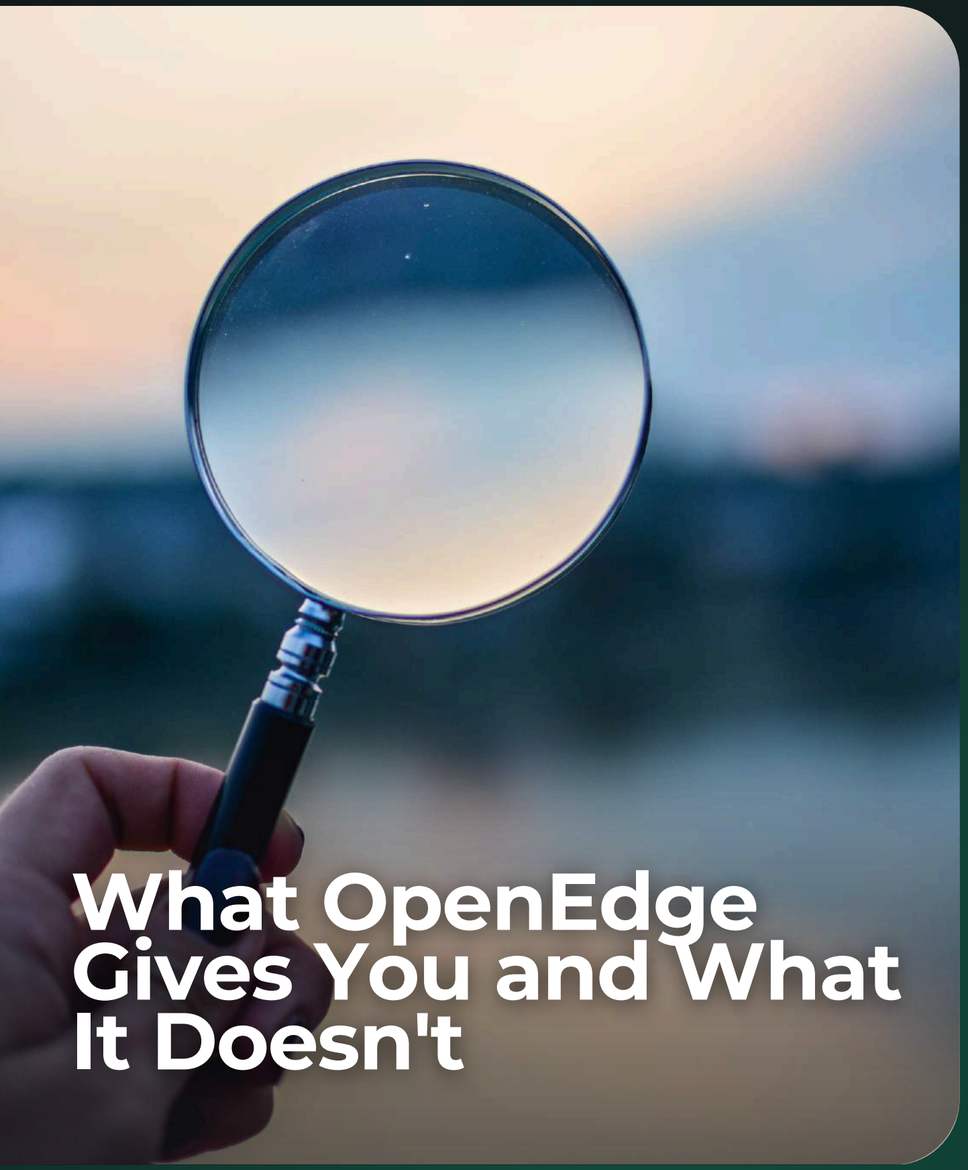
By default, that request either goes unlogged, or lands in an audit table nobody's watching in real time

Wayfare

Now Imagine...
Same app, same endpoints, but different intent!

- One session writes 40 prescription records in two minutes
- Even with Auditing enabled, nothing flags this as unusual - it's just another row in the audit table
- It looks exactly like traffic you already trust

Wayfare



What OpenEdge Gives You and What It Doesn't

ACID transactions, REST APIs, decades of proven business logic

It even has a native Auditing feature - logs who, what, and when changed

But that's a change log, not a behavioral system - it doesn't learn what's "normal" and it doesn't alert in real time

The gap isn't "no logging" - it's no detection

Wayfare



Why This Gap Is Expensive

- Users trust our OpenEdge applications with their personal data
- Compliance asks "prove nothing bad happened" - a request log alone doesn't do that
- Manual review doesn't keep up past a handful of alerts a week

Wayfare

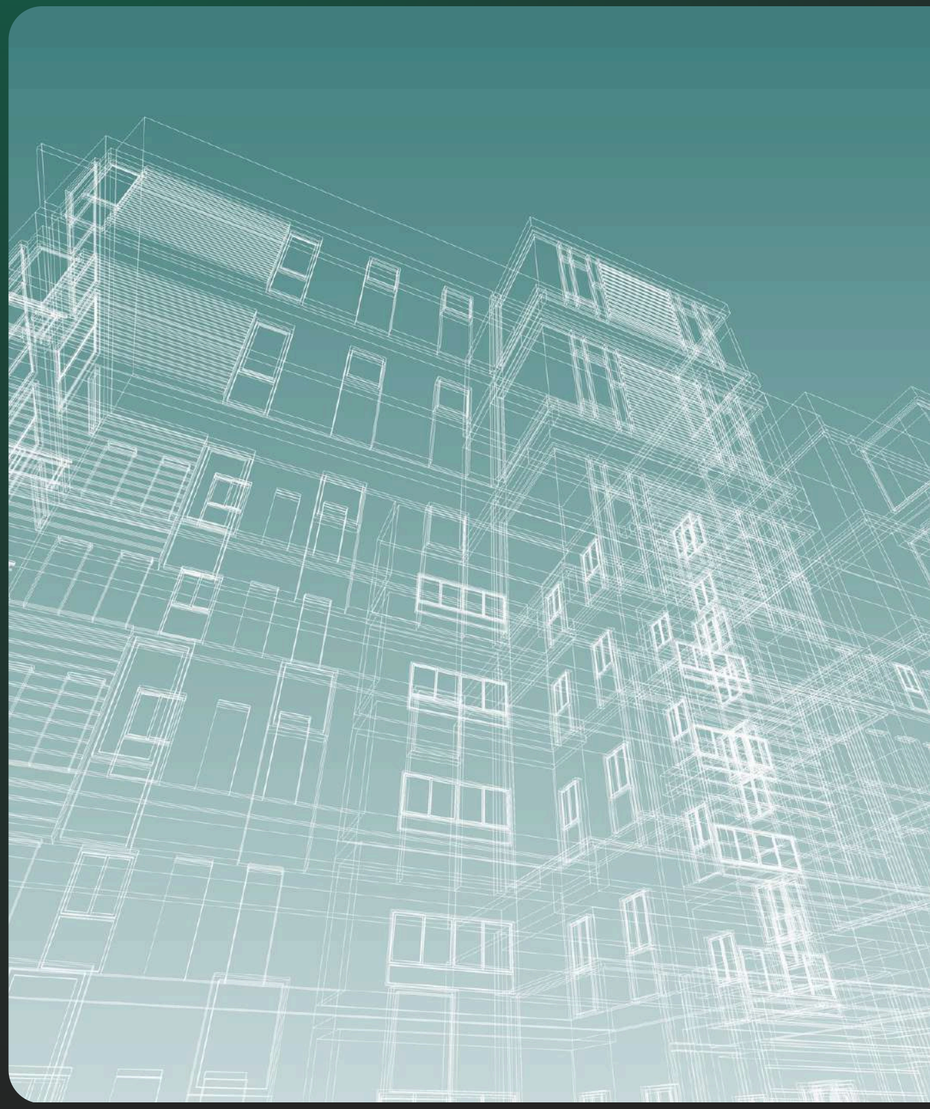
Closing the gap

without leaving ABL

Real-time detection, instrumented directly in the application you already have!

Wayfare

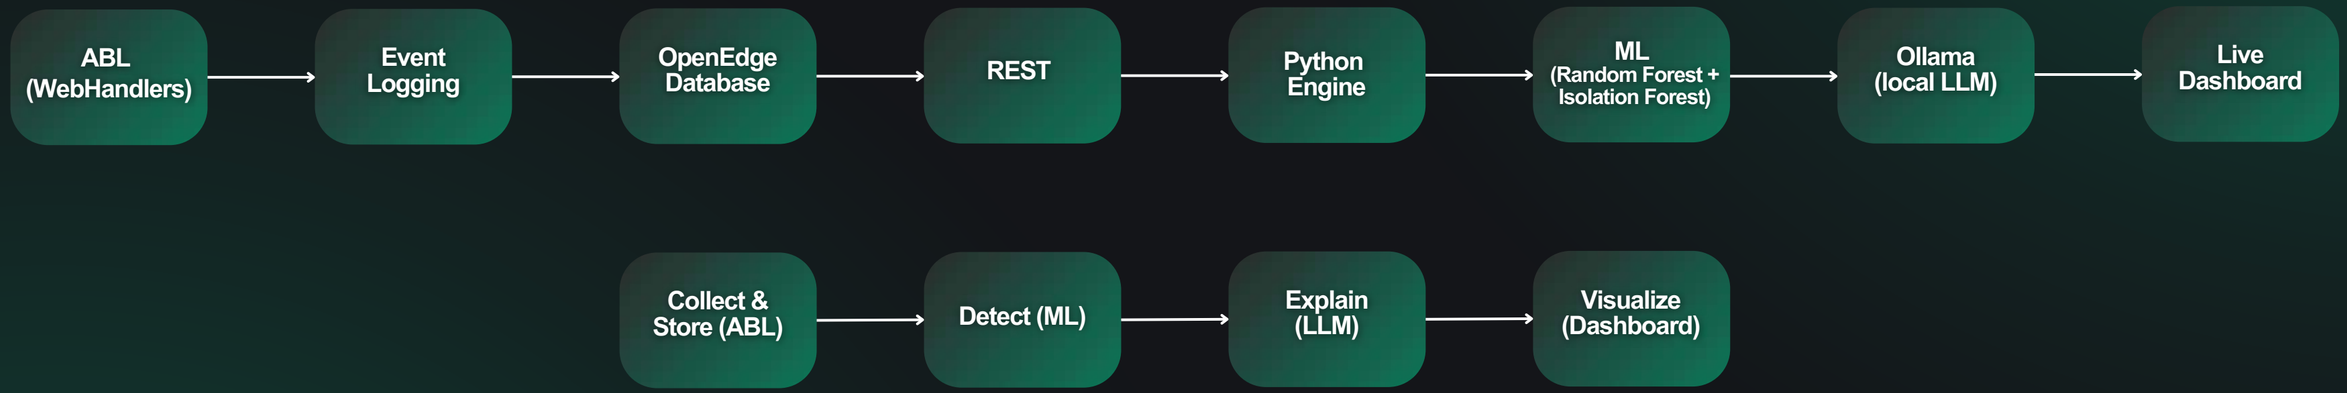
SOLUTIONS ARCHITECTURE



Wayfare

ONE PIPELINE

Four Main Jobs



Wayfare

WHY
it starts in ABL

- Every request already passes through a WebHandler
- Instrumentation = one added call per endpoint
- No agent, no proxy, no change to the PASOE installation itself

Wayfare

THE STACK
Layer by Layer

ABL

OpenEdge 12.8, PASOE, WebHandlers, REST (storage: one new table, directly in the OpenEdge database - no separate store)

Detection

Python 3.10+, scikit-learn (Random Forest + Isolation Forest)

Explainability

Ollama, local LLM

Visualization

a lightweight live dashboard - zero external dependencies

Wayfare

WHY
keep OpenEdge at the Center

- No database migration
- No rewrite - existing business logic stays untouched
- We only add a logging call and a read-only export path via REST
- Everything that follows is additive, not a replacement

Wayfare

Every Click Already Leaves a Trace in ABL,

before a single line of Python runs.



Wayfare

COLLECTING EVENTS FROM ABL



Wayfare



WHAT GETS LOGGED

- Every read and write the app already performs
- Who did it, from where, and exactly what was touched - the same fields, every time
- A free-text detail field for context (what was searched, how many records)

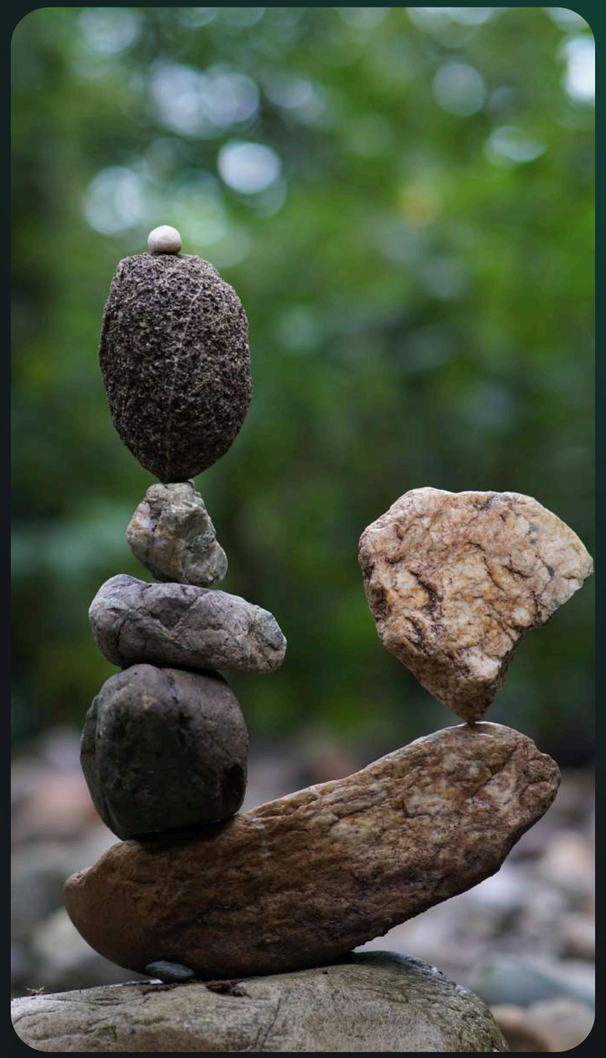
Wayfare

One Event, Start to Finish

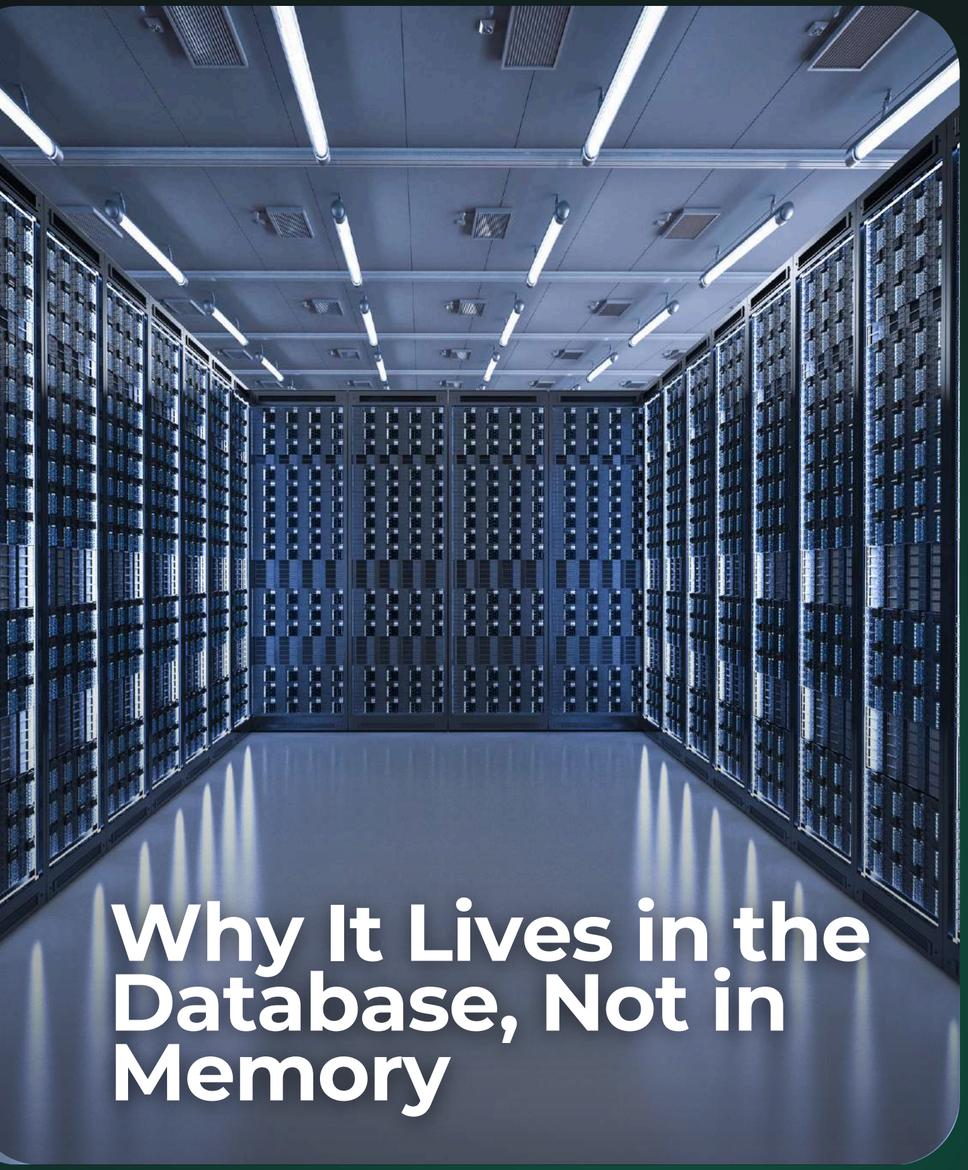
A doctor issues a new prescription

The write is captured in the exact same transaction, the instant it happens - no extra round trip, nothing to coordinate separately

From there it's just another row, waiting to be picked up - nothing special, nothing slow



Wayfare



Why It Lives in the Database, Not in Memory

PASOE runs multiple sessions in parallel to serve concurrent requests - each with its own isolated memory

An in-memory log would only ever see a fraction of the real traffic

The database is the one thing every session actually shares

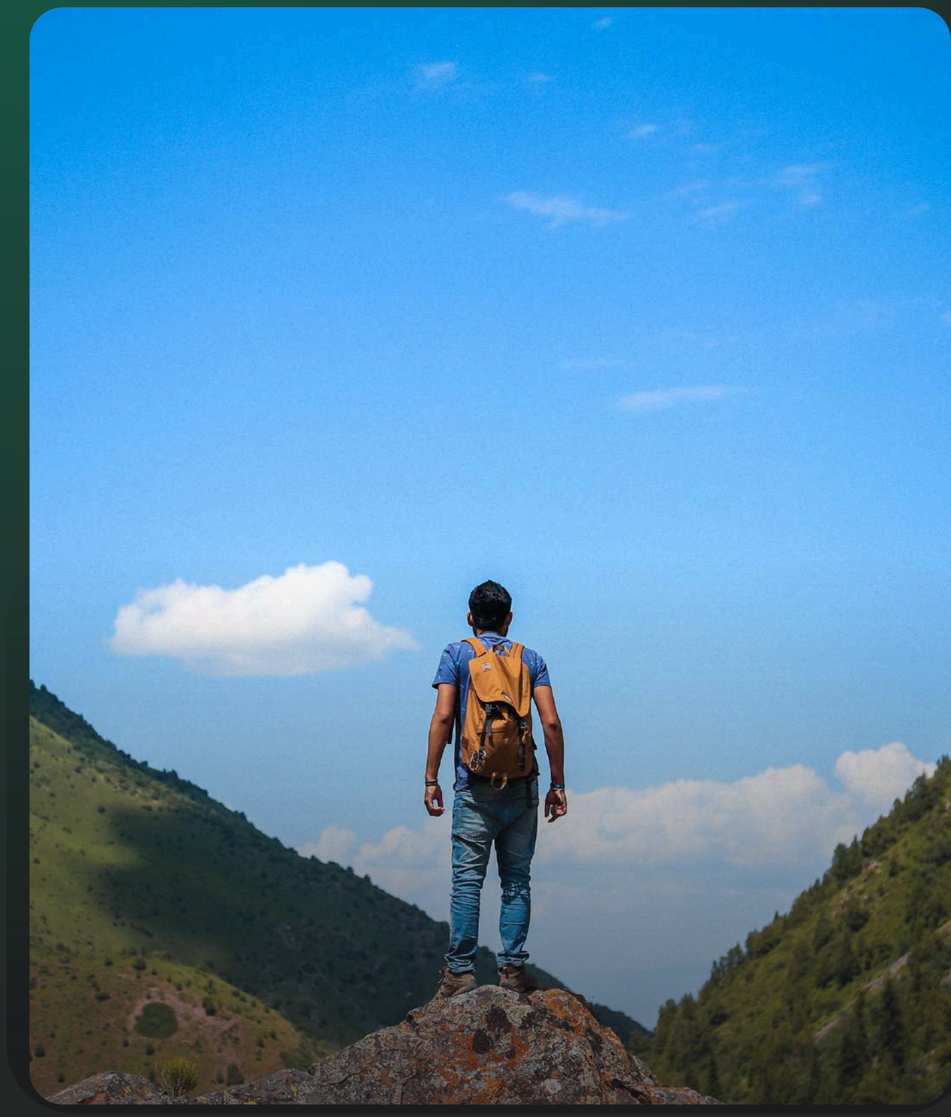
Wayfare

Race-Safe by Construction

Two requests firing at the same instant never collide.

Each event gets an atomic, strictly increasing number straight from the database.

No coordination code needed - the database guarantees it.



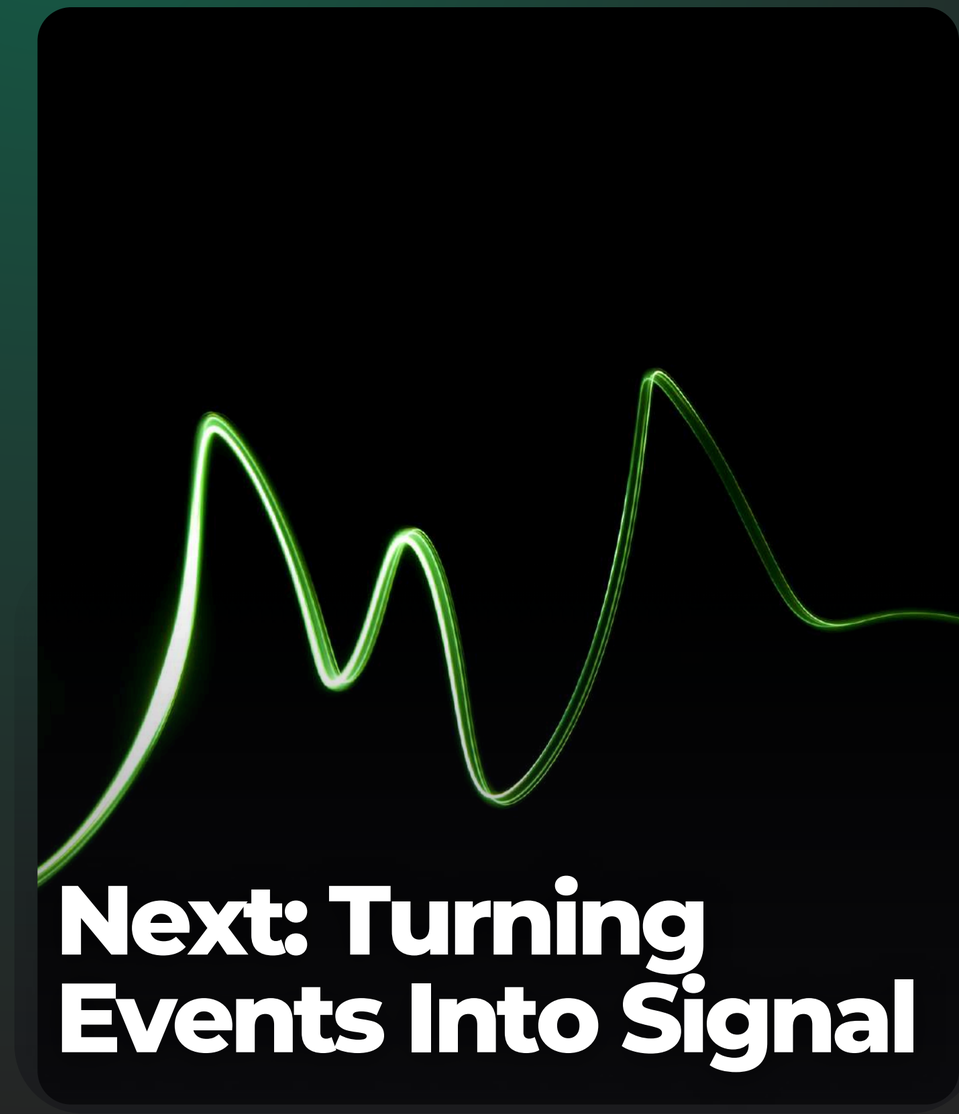
Wayfare

Getting Events Out, Safely

A dedicated, read-only endpoint - nothing it exposes can modify data

Python asks for "everything since sequence N" and polls every two seconds

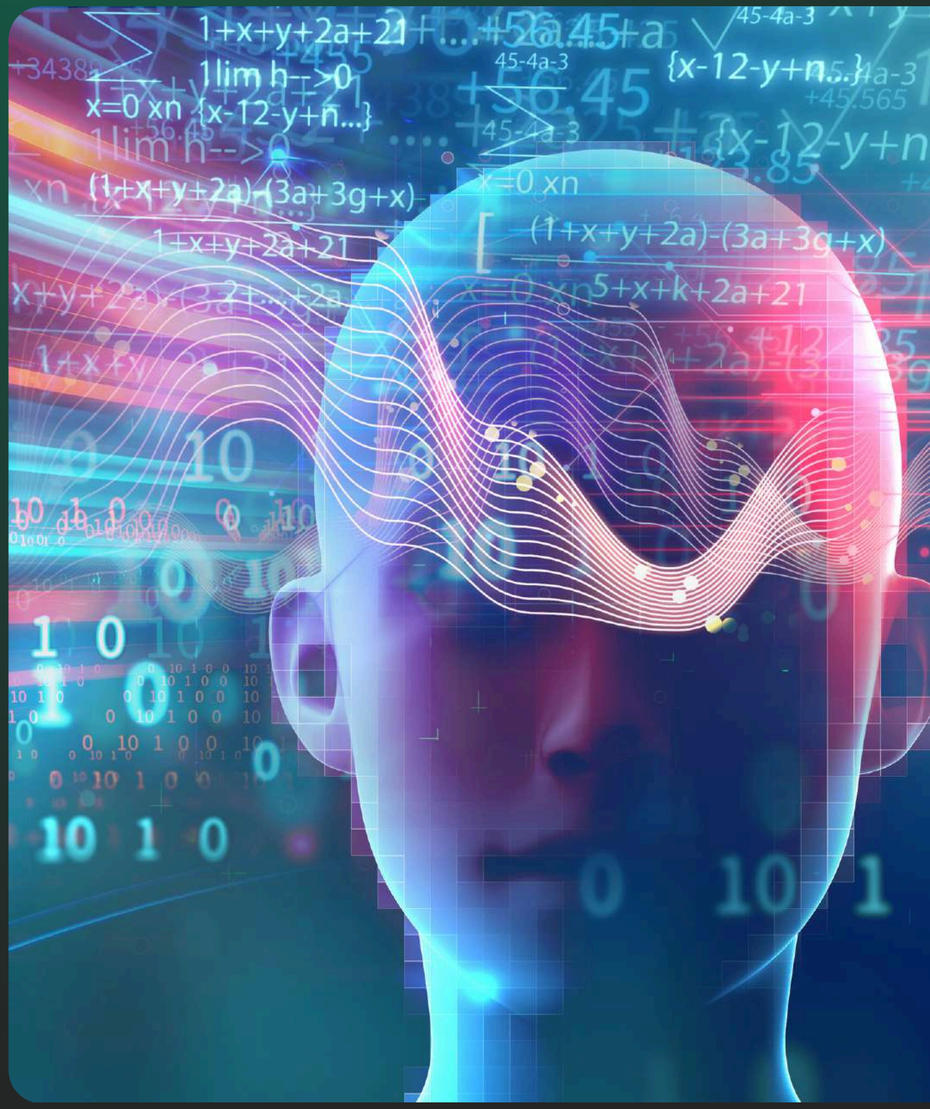
A watermark, not a full scan - no missed events, no reprocessing



Next: Turning Events Into Signal

Wayfare

MACHINE LEARNING DETECTION



Wayfare



From Events to Signal

Every session's last 30 seconds of activity becomes one feature snapshot

How many actions, how varied they are - different kinds of records, different individual records, different identities - what share were writes, how much time since the last one

Training and live scoring compute it exactly the same way - no drift between what the model learned and what it sees live

Wayfare

Two Layers, Two Jobs

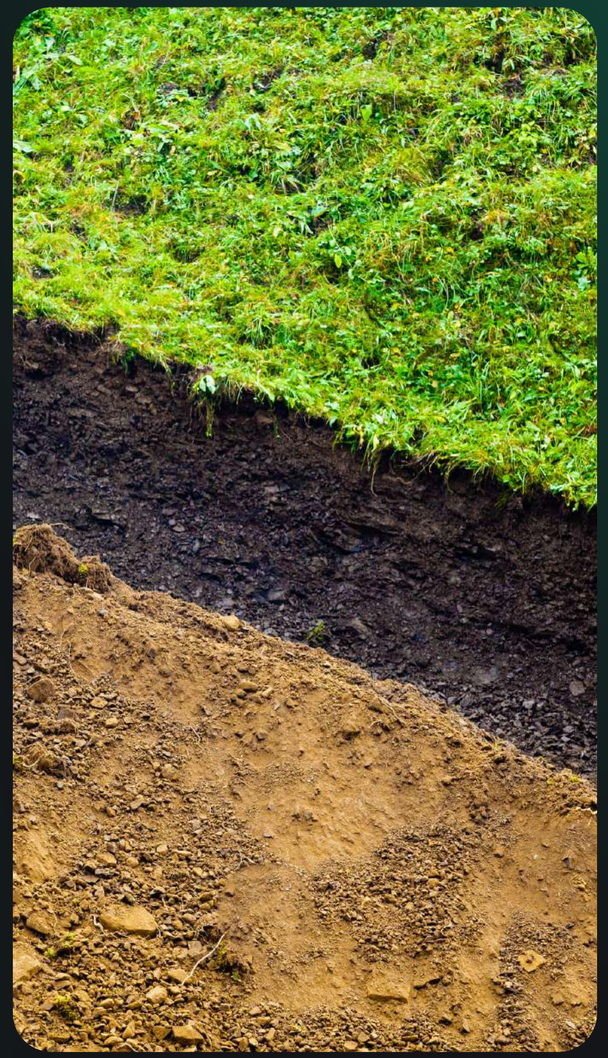
RANDOM FOREST:

trained once, offline, on known-good vs. known-attack examples - recognizes familiar attack shapes with high confidence, never drifts

ISOLATION FOREST:

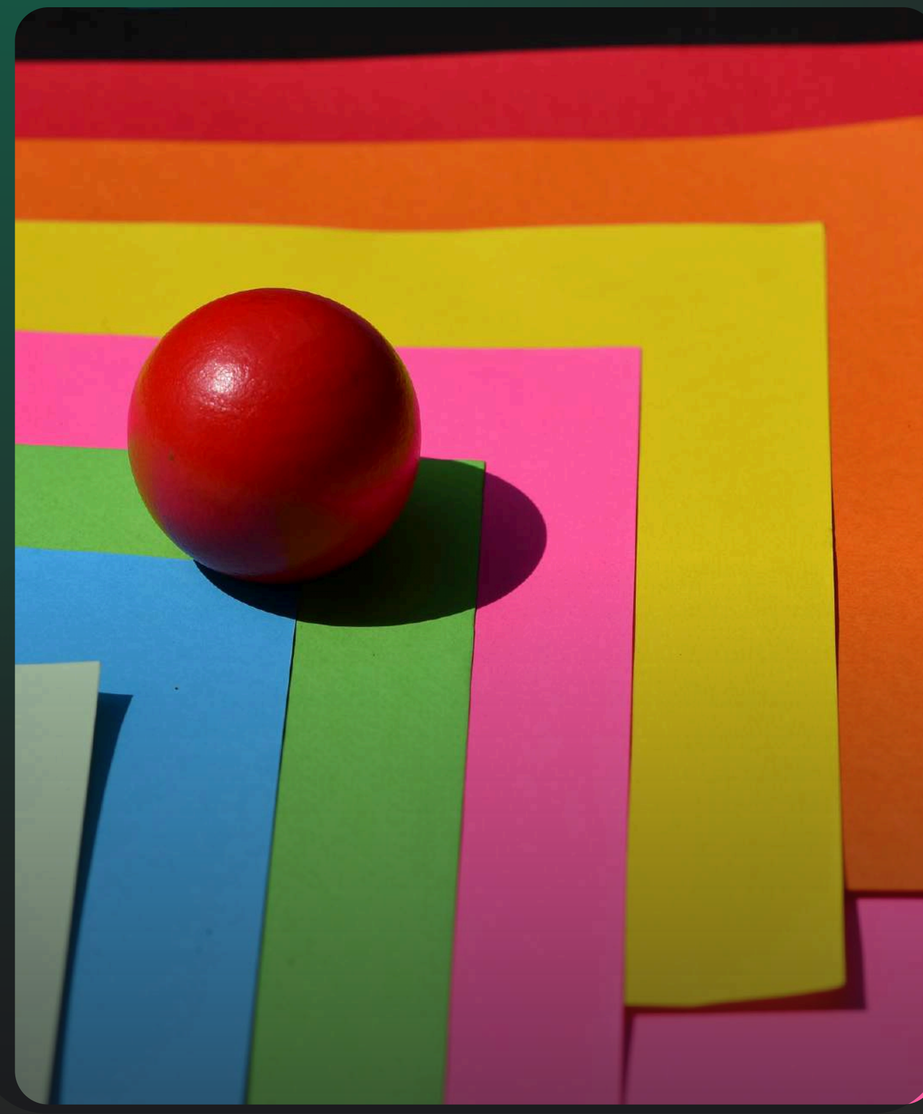
keeps adapting to what "normal" looks like from live traffic - catches patterns nobody labeled in advance

AN EVENT IS FLAGGED THE MOMENT EITHER ONE DISAGREES WITH IT



Wayfare

Why These Two, Not Something Heavier

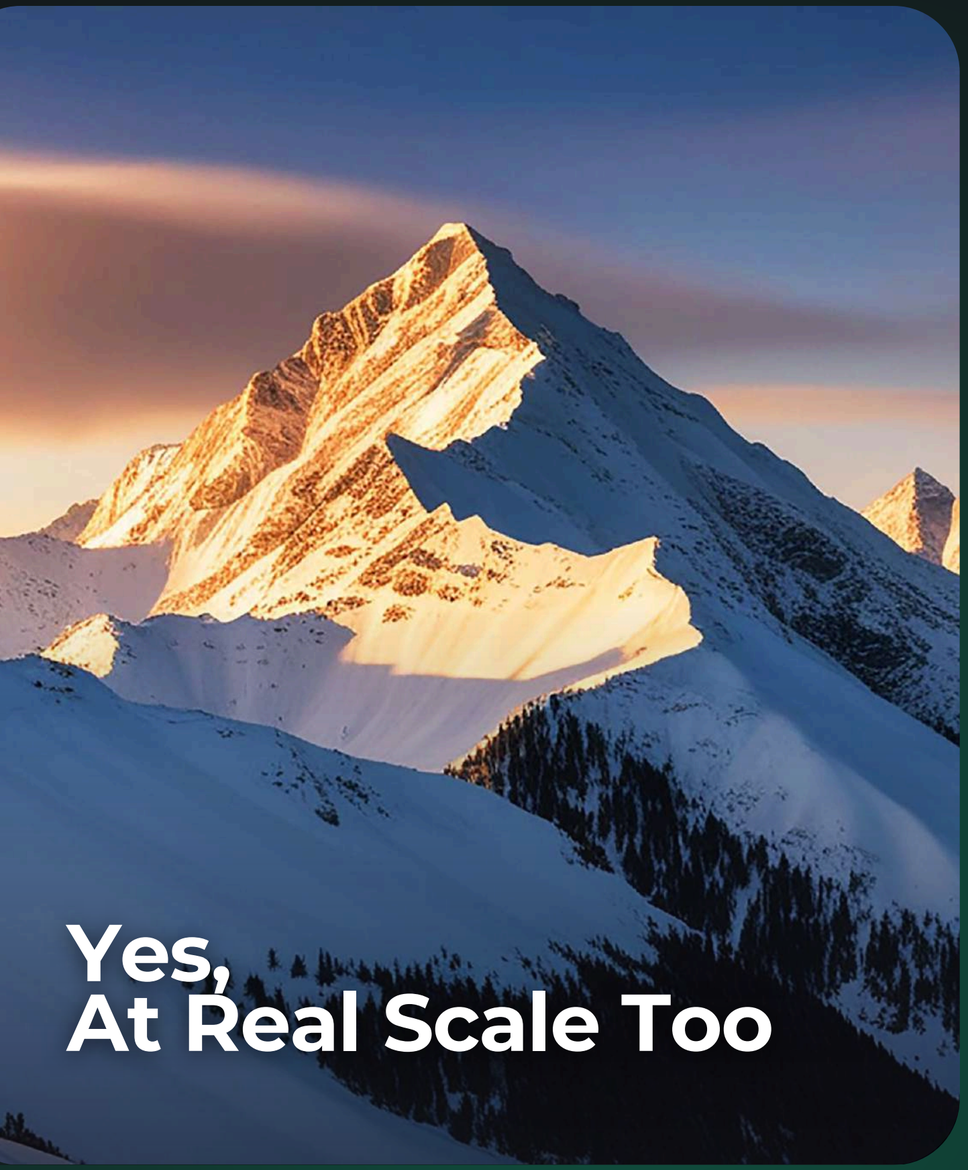


The data is small, tabular, six numbers per event - not images or raw text; tree-based models are proven to match or beat deep learning at exactly this scale and shape

Isolation Forest doesn't need labeled attacks to train - essential, since nobody can pre-label every future attack shape

Both run almost instantly on ordinary hardware, no GPU, no massive training set - fits a real-time poll loop, not a research lab

Wayfare



Yes,
At Real Scale Too

Isolation Forest was built for this: it trains on small sub-samples, not the whole dataset - proven to make it both faster and more accurate at scale, not less

Random Forest's trees train independently - scaling up means more cores, not a different algorithm

What changes at production scale isn't the algorithm - it's splitting one global model into several (per role or workflow) and re-tuning thresholds against real traffic instead of synthetic

Wayfare

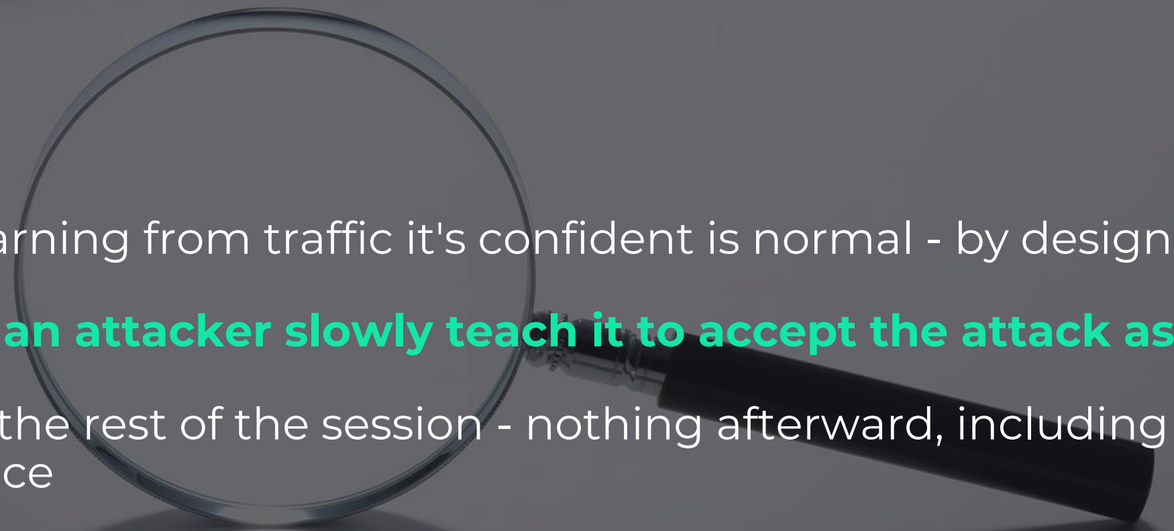
Catching What Nobody Taught It



- One attack pattern was deliberately held back from Random Forest's training data
- Only Isolation Forest has a chance at it - no known-pattern match to rely on
- It catches it anyway, on its own - proof the system isn't just memorizing a list of known attacks

Wayfare

Why It Can't Be Talked Out of a Detection



- Isolation Forest keeps learning from traffic it's confident is normal - by design
- Obvious question: **could an attacker slowly teach it to accept the attack as normal?**
- One switch freezes it for the rest of the session - nothing afterward, including a live attack, can retrain it into tolerance

Wayfare

Turning a Score Into a Sentence

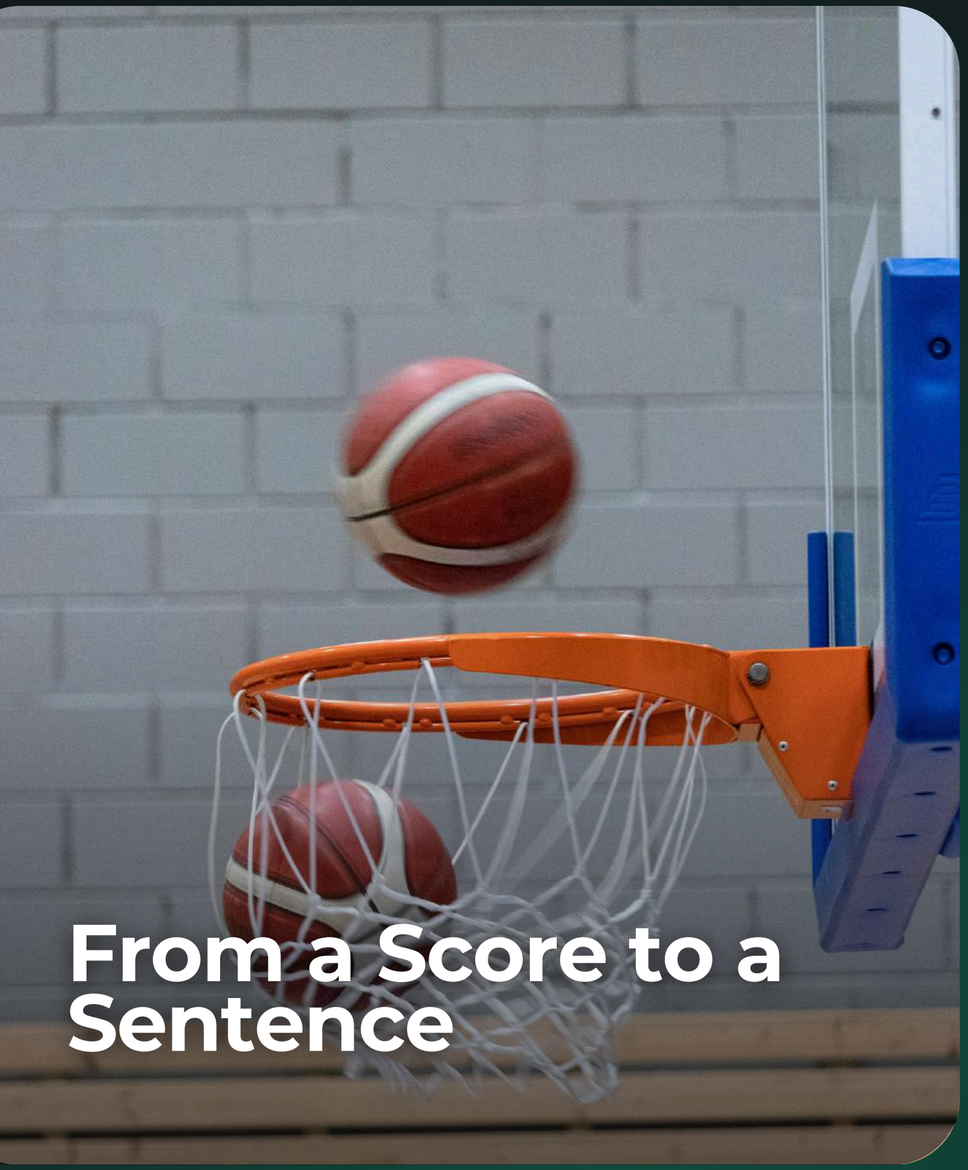


Wayfare

EXPLAINABILITY WITH LOCAL AI



Wayfare



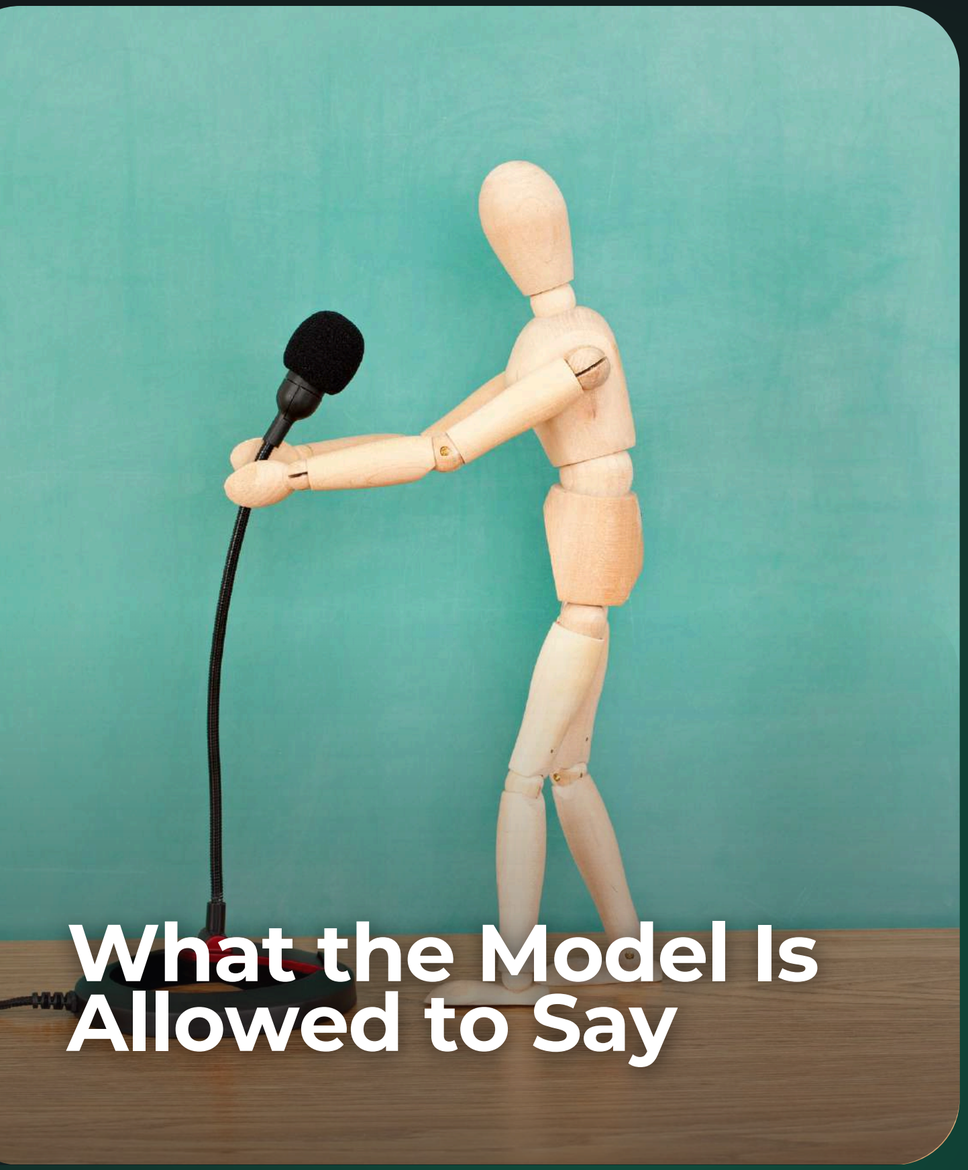
From a Score to a Sentence

A flagged event comes in as numbers - a probability, a handful of counts

The model turns that into two or three plain-English sentences, for a reader who isn't technical - a compliance officer, a hospital administrator

Only the first flagged event per session gets explained - not the same story 45 times during one burst

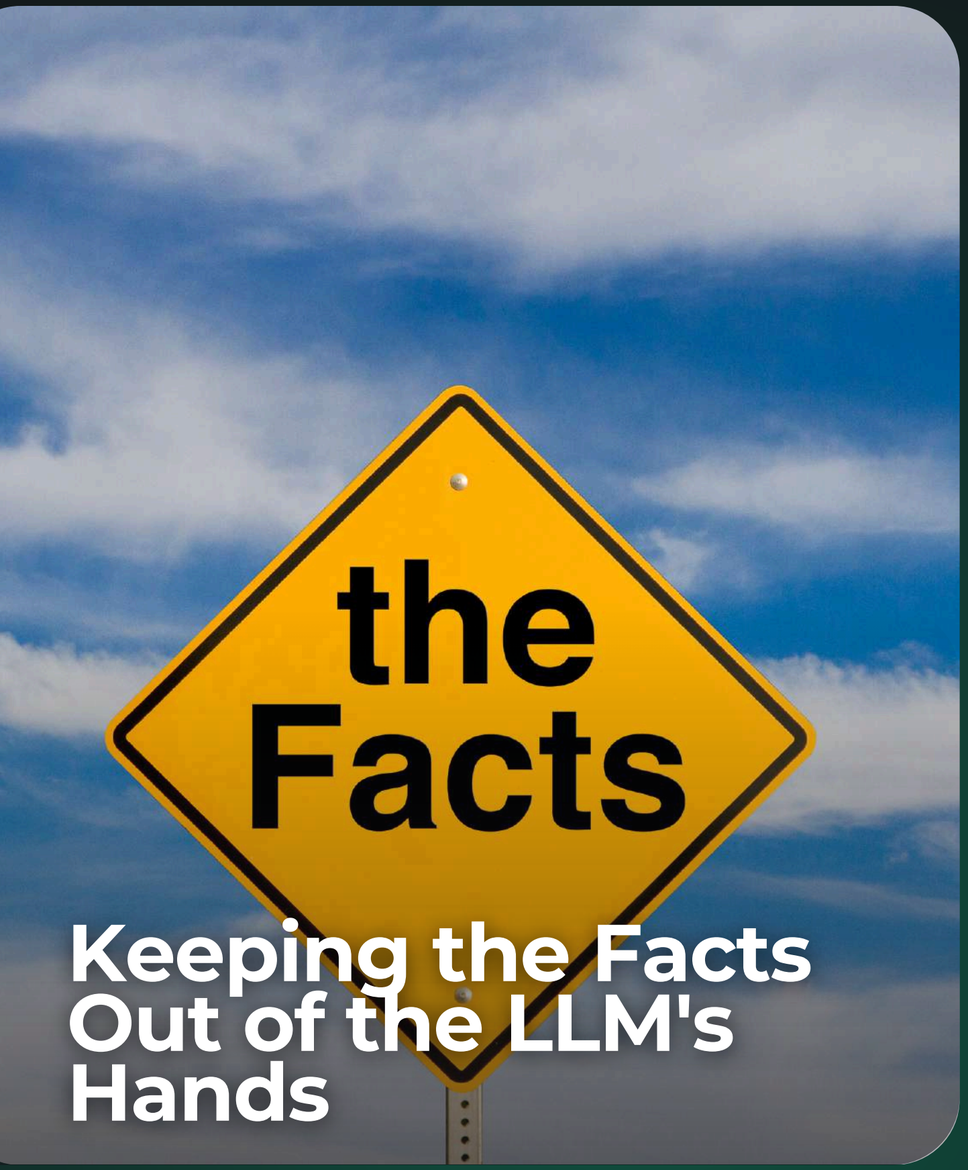
Wayfare



What the Model Is Allowed to Say

- Only the numbers it was actually given - never an invented action, role, or detail
- No speculation about intent or state of mind - never "may be under duress," never "verify their intention"
- Always ends with one concrete action: review the access, confirm it was authorized

Wayfare



Keeping the Facts Out of the LLM's Hands

Which detector flagged the event is computed in plain code, stated as fact - never left for the model to recall

A demo can't afford a small local model occasionally saying "both agree" when only one did

The one sentence that has to be exactly right, every time, is the one sentence the model never writes

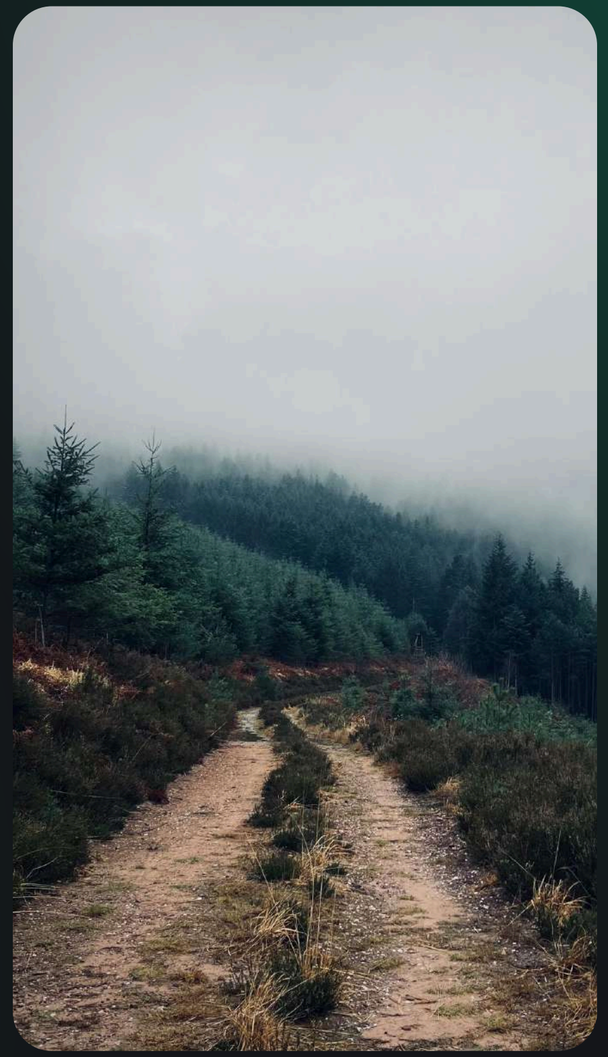
Wayfare

Tuned Against Real Failures, Not Guesses

Run at low temperature - a factual-description task, not a creative one

Found by testing: at default settings, the model embellished - once inventing "potentially altered records" when the data said zero writes

A second safety net catches a specific habit - defaulting to "prescription" even when told a different record type - and corrects it before it's shown



Wayfare

Fails Safely, Costs Nothing

If the local model isn't running or times out, the system says so instead of crashing - detection resumes right after, it's only the explanation that's delayed.

No API key, no cloud call - nothing about the event ever leaves the machine.



Wayfare

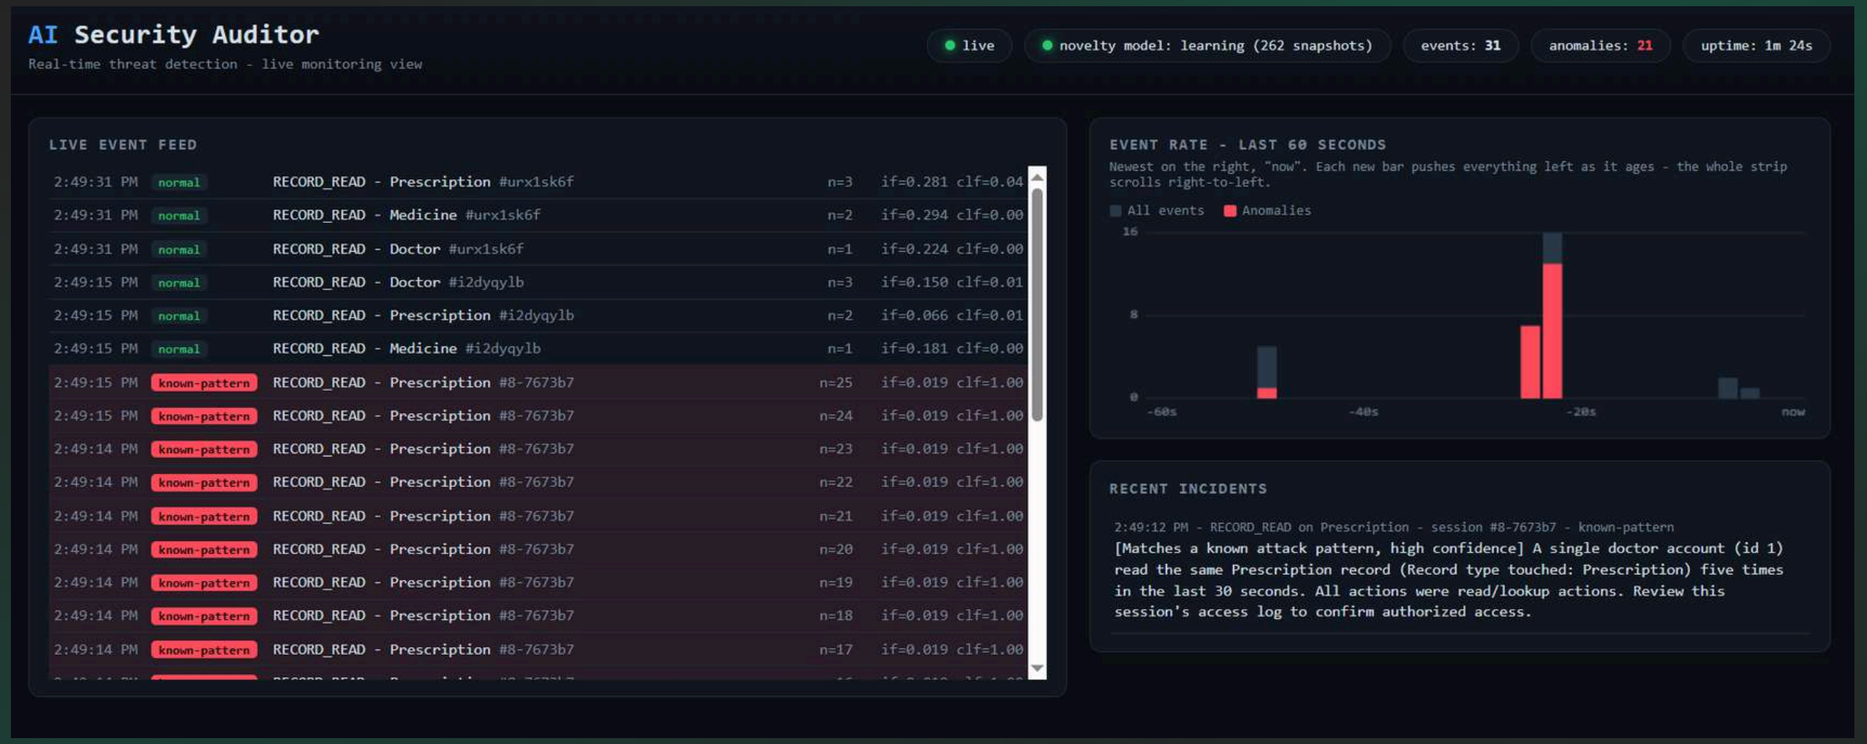

From One Sentence to a Reportable Trail
Wayfare

LIVE DASHBOARD & COMPLIANCE REPORTING



Wayfare

Inside the Live Dashboard



A live feed of every event, normal or flagged, updating every second

A rolling 60-second traffic chart, with anomalies highlighted in the same view

A banner that surfaces the plain-English explanation the moment a new incident lands

Wayfare

One Explanation, Reused Everywhere

- Each incident's explanation is generated once, live, the moment it's flagged
- The same text feeds the dashboard banner, the incident list, and the compliance report later
- Nothing is regenerated for the report - stays fast even with dozens of incidents, and the explanation can't drift between what you saw live and what's on paper

Wayfare

From Incidents to a Report

- Every explained incident is already being written to a plain log as it happens
- At any point, that log becomes a compliance report - a short executive summary plus a full table, one row per incident
- Session, event, which detector caught it, confidence, explanation - already there, not reconstructed afterward

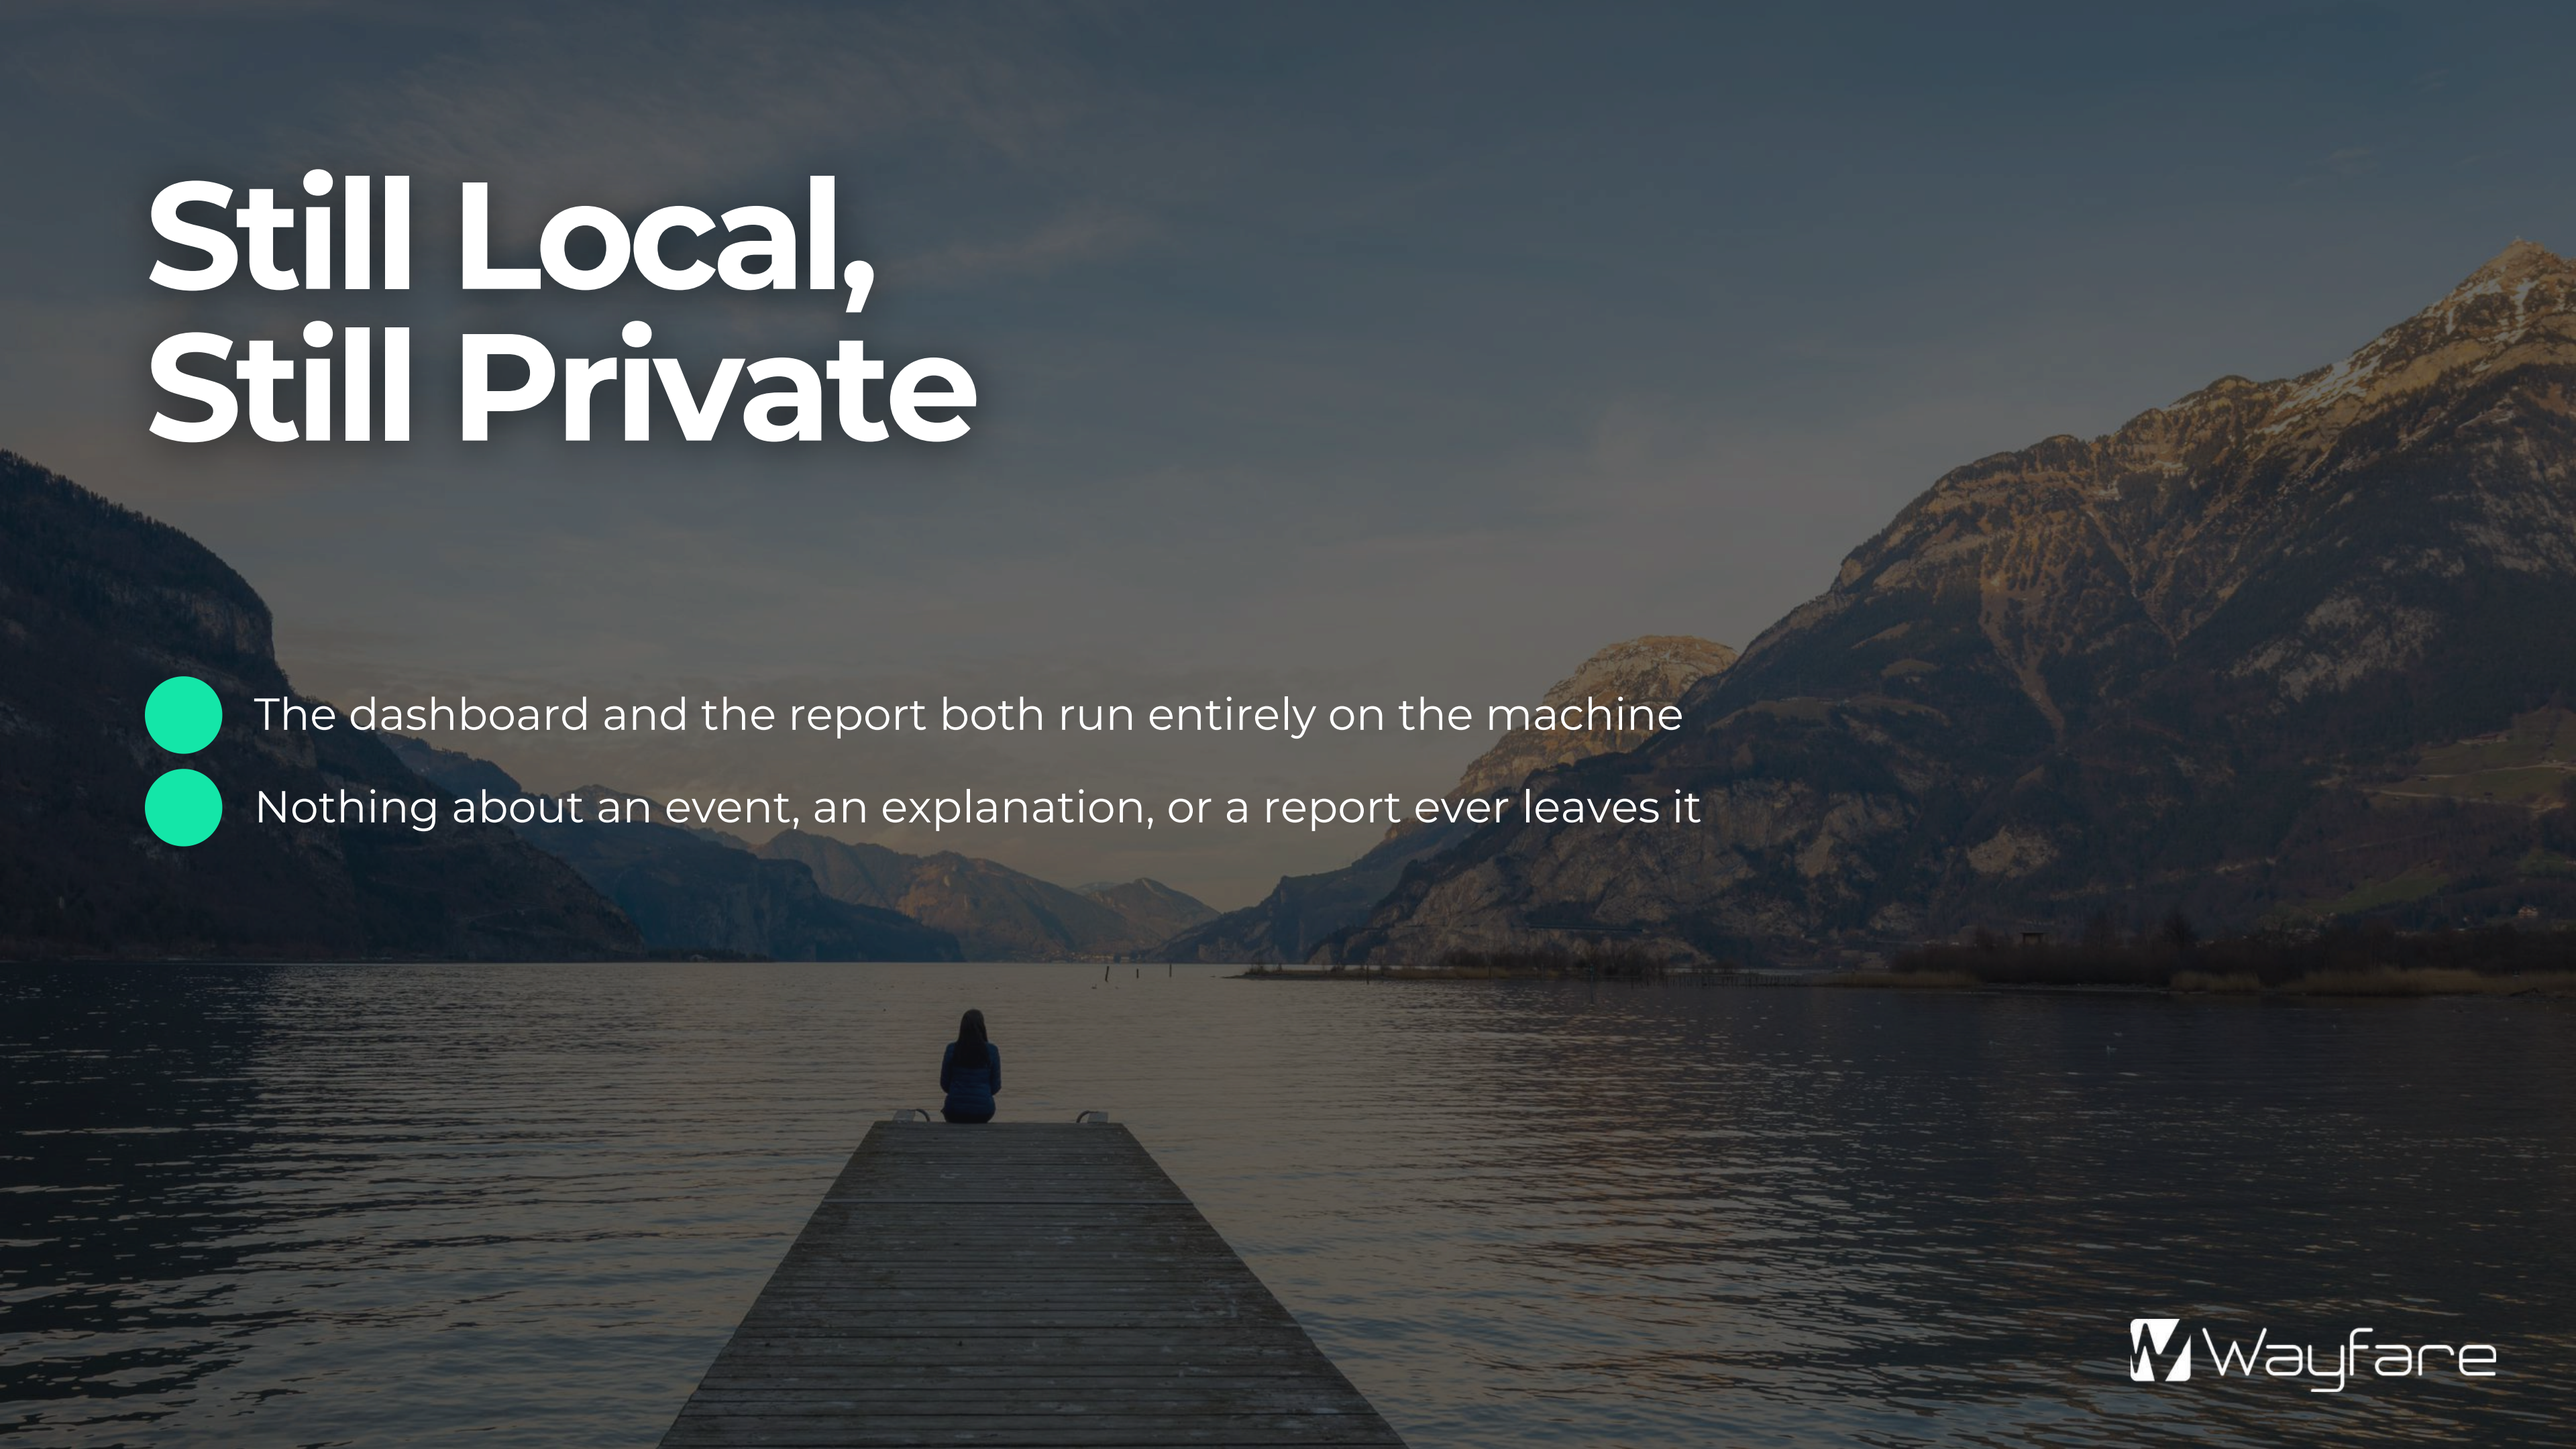


Wayfare

Same Discipline, One Level Up

- The executive summary follows the same discipline as every individual explanation - only the real incidents, no invented pattern, no model internals, same low temperature
- It describes the shape of the whole period - how many incidents, how many sessions, what kinds of behaviour - not a copy-paste of each one

Wayfare



Still Local, Still Private

- The dashboard and the report both run entirely on the machine
- Nothing about an event, an explanation, or a report ever leaves it

Wayfare


NOW
LET'S SEE IT RUN

LIVE DEMO & CODE WALKTHROUGH



Wayfare

RESULTS & CONCLUSIONS



Wayfare

Trained on Real, Labeled Traffic



550 labeled feature snapshots from this machine's own run - 238 normal (from 40 realistic browsing sessions) and 312 attack examples across all three trained shapes, after trimming each attack's ambiguous first few events.

What the classifier leaned on hardest: how many actions landed in 30 seconds and how fast they came - not which record type was touched, and barely the write ratio at all.

Wayfare

Known Attacks Confirmed in Under 5 Events



All three trained attack shapes - rapid single-record reads, many doctor identities in one session, back-to-back endpoint bursts - reached 99-100% classifier confidence by their 5th event, live.

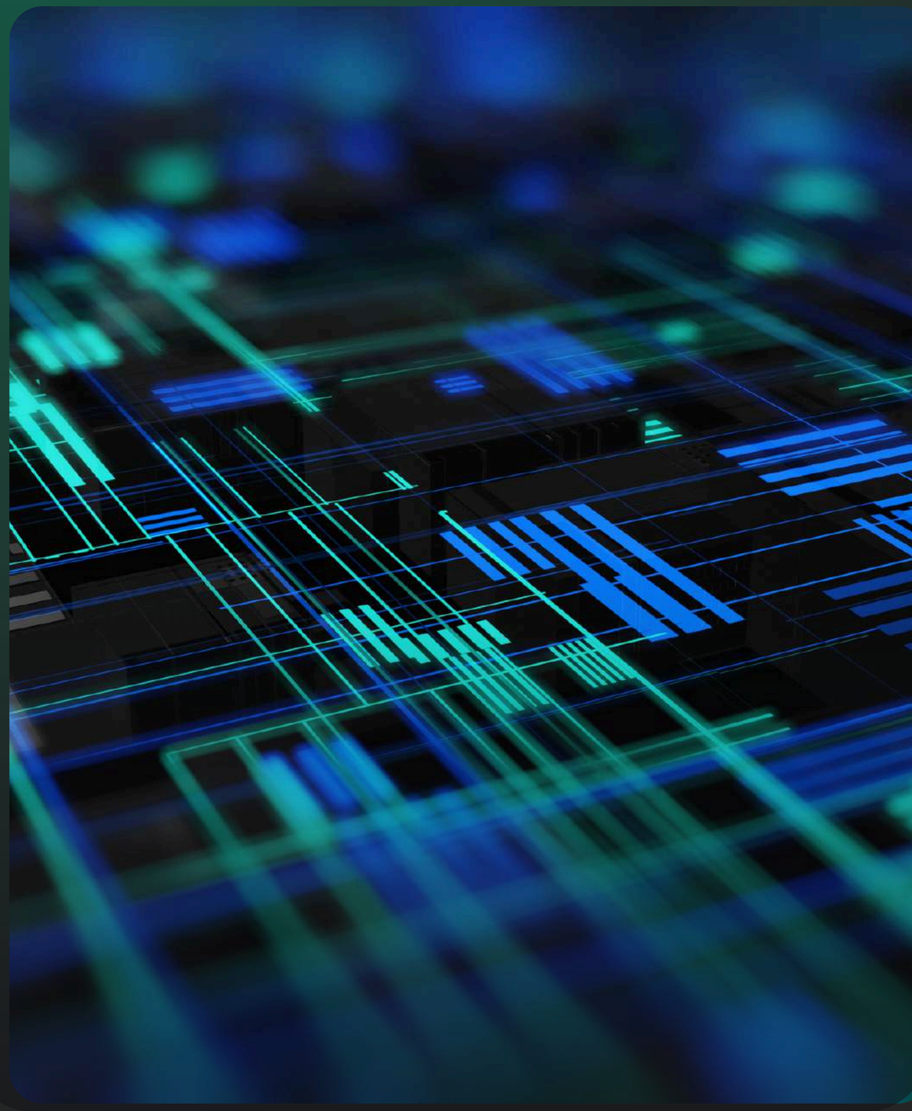
The explanation came out correct on the first try: "A single doctor account (id 1) read the same Prescription record... five times in the last 30 seconds... Review this session's access log to verify the doctor's authorization" - generated the moment it was flagged, nothing edited afterward.

Wayfare

A Pattern Nobody Trained It On

A write-heavy session (write ratio climbing to ~94%, a new record touched almost every time) matched none of the three trained attack shapes - classifier confidence stayed at 17-24%, well under the 70% alert threshold

The novelty layer flagged it anyway, on its own, purely for being statistically unlike the 238-session baseline - no label, no retraining, no prior example



Wayfare

A Known Rough Edge

At a very high write ratio, the local 3B model once produced an internally inconsistent sentence - said "read" while describing actions that were 83% writes.

The underlying data was correct (write_ratio was reported accurately); the wording slip was the model's, not the pipeline's.

A known trade-off of running a free, fully local 3B model instead of a larger one.



Wayfare

What It Costs, Per Event

Both trained models together: under 1.5MB on disk

Measured on this machine, 2,000 iterations each:

~12.4ms of total detection overhead per event (feature extraction + both models)

COMFORTABLY INSIDE THE 2-SECOND POLL CYCLE

The novelty model is the heavier of the two

~6.8ms VS ~5.5ms for the classifier

FEATURE EXTRACTION ITSELF IS UNDER A TENTH OF A MILLISECOND

Wayfare

The One Slow Step Is Rare by Design

The local LLM explanation call took ~7.1 seconds on this machine - real, but it runs once per session, only for the first anomaly, never per event and never blocking detection itself

Wayfare


Wayfare
ASK ME
A QUESTION

Our PUG Challenge sessions

PROVIDENCE



CLAUDIU

Understanding database performance

Monday, October 5, 5:00 PM

AMSTERDAM



ALEXANDRU

Exploring Local AI: Capabilities and Limits of Self-Hosted Models for Development and Documentation

Thursday, November 5, 3:15 PM



CARLA

AI Security Auditor: Real-Time Threat Detection for OpenEdge Applications

Friday, November 6, 10:15 AM



CLAUDIU

Understanding database performance

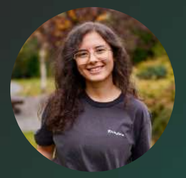
Friday, November 6, 1:45 PM



DIANA

One BA, two teams: How I standardized workflows without slowing anyone down

Friday, November 6, 9:00 AM



DIANA

RAG 2.0 - The Next Generation of RAG Systems

Wednesday, November 4, 4:15 PM



OANA

Guideline to an Accessible Web Application

Thursday, November 5, 9:00 AM

Wayfare

Wayfare



thank you

wayfare.ro

wayfare.community

wayfare.academy

business@wayfare.ro