TAZI

TECHNICAL BRIEF · 2026-07

TAZI AI Platform

Enterprise-Ready, Grounded AI with Validated Actions

ONE PLATFORM, THREE PATHS TO PRODUCTION

TAZI ships with pre-built use cases; teams also build their own on the same governed foundation. Every path lands on identical serving, monitoring and documentation infrastructure.

1 Bring Your Own Models

Onboard existing in-house models. The platform adds the explanation layer, validation agents, drift monitoring and auto-generated compliance documentation. No rip-and-replace.

2 Deploy Pre-Built Use Cases

Pre-built explainable agents (fraud/ATO & AML, retain & grow, voice of customer) configured on your data and deployed in days.

3 Build New Use Cases

Define a new label and business KPI, build a grounded explainable model from scratch, validate and deploy — for risk types no vendor has templated yet.

★ **Low barrier to entry: start with a joint discovery phase and one use case at a time — data cleaned as you go, a digestible process. Path 3 then becomes the risk team's long-term path.**

HOW A GROUNDED AGENT WORKS

Grounding	Predictive models are trained on your data with explainable ML; continuous learning and GEPA keep models and agents current without full retrains.
Explanation	A companion explanation model produces dynamic, business-friendly reason codes and named patterns for every score.
Validated actions	Agents propose actions from your dynamic population of data; Expert-Panel and Focus-Group agents validate them against your organization, business KPIs and focus population before they reach a queue.
Human in the loop	Risk officers coach models through conversational agents and web interfaces; questions can be tailored to each problem; approval gates can stop any action; every step is logged.
Documentation	Model cards, data dictionaries and compliance documents are auto-generated by default.

ENTERPRISE-READY PLATFORM CAPABILITIES

Agent & Business-User Friendly MCP tools, workflows and ETL for agents; no-code UI and MLOps for humans.	**Continuously Updated** Continuous learning and GEPA keep explanation models and agents current — no full retrain.	**Explainable & Validated** Explainable ML with business-friendly explanations; actions validated by expert-panel and focus-group agents.
BYOK / BYOM Bring your own keys and LLMs; keep control of your models and secrets.	**Cloud · On-Prem · Hybrid** Deploy where your data lives — no vendor lock-in.	**Multitenant** One secure platform across teams and business units, isolated by design.
Real-Time Deployment Single REST API; event-driven or batch; auto-scaling for high throughput.	**Drift Monitoring** Continuous data- and concept-drift detection on every deployed model.	**Secure, Governed & Documented** Auditable activity logs, auto-generated model cards and compliance documentation.

BUILDING A NEW RISK USE CASE (PATH 3)

The workflow the risk team runs when a new risk type emerges — the same governed lifecycle behind every pre-built agent:

1. **Define.** Specify the target label and the business KPI that measures the new risk (e.g., $ exposure caught per inspection); build the data dictionary.
2. **Ground.** Profile the data, train the explainable predictive model and its companion explanation model on your data.
3. **Coach & validate.** SMEs review reason codes and coach the model through conversational agents and the no-code UI; Expert-Panel and Focus-Group agents value-test the proposed actions.
4. **Deploy.** Serve real-time via REST API, event-driven or batch; A/B test against the incumbent process before go-live.
5. **Monitor & adapt.** Continuous learning absorbs change; drift detection and monthly monitoring flag when the risk itself shifts; compliance documentation stays current automatically.

Humans stay in the loop throughout. Risk officers own the label definition, can approve and edit the actions, prompts and data, and can gate any step. The platform automates the documentation, the accelerated development, and the ease of going into production — not the judgment. The day-to-day result: less manual triage, less report-writing, and a market advantage built on TAZI's field-proven explainable-ML and continuous-learning IP — not on a roadmap promise.

WORKED EXAMPLE: ATO FRAUD AGENT

Label: fraud / not_fraud per session and transaction, retrained on investigator feedback.

Business KPI: efficiency = $ fraud caught per inspection; false-negative $ exposure = benefit value × fraud probability.

POV: Prepare (data, label, KPI, data dictionary) → Design & Develop (model, XAI, continuous learning, human-in-the-loop, compare vs incumbent) → Monitor & Maintain (compliance docs, integration, A/B test, go live) — approximately 72 POV hours.

+37% / 53%

Investigation efficiency (test / hold-out); $480K and $2.5M more fraud caught respectively.

50% → ~80%

True positives; false-negative cost $1M → $235K.

VALIDATED-ACTION LAYER: MEASURED IMPACT

What Expert-Panel + Focus-Group validation adds on top of the same predictions:

94% → 42%

Needless manager alerts in Wealth Retain & Grow; 27% of the book cleared as "no action"; 37% of corrupted-data cases caught, zero false alerts.

20 → 0

False SAR exposure on legitimate customers in the top ATO tier: same detector, corrected action layer; 100% auditable rule trail.

98 → 4

Top-tier AML SAR escalations (96% fewer, all corroborated); tipping-off phrasing removed 99% → 0.4%; false-positive overrides added.