RevoAI

Explainable AI for Sovereign Security

PRODUCT BRIEFS

ISNR Abu Dhabi 2026

Hall 8 • Stand 8ST-09

Three Pillars. One Platform.

Purpose-built for UAE government security agencies

See every decision. Trust every model.

01 Counter-Threat Finance & AML Platform

Explainable financial-threat detection with audit-ready evidence

Overview

RevoAI's Counter-Threat Finance & AML Platform helps agencies and regulated institutions detect suspicious transaction networks, financial-crime patterns, sanctions risk, and AML/CTF anomalies. It is designed to make alerts explainable, defensible, and reviewable by analysts, compliance teams, investigators, and authorised oversight bodies.

Target Buyers

Ministry of Interior, Cyber Security Council, SIRA, financial-crime units, AML/CTF teams, regulated financial institutions, and security agencies requiring explainable audit trails.

Core Capabilities

Transaction Monitoring	Screen transaction streams or batch uploads (CSV, JSON, SWIFT-style) with risk scoring and anomaly detection.
Graph Intelligence	Build entity relationship networks to identify suspicious clusters, layered transactions, structuring patterns, and coordinated activity.
Explainable Alerts	Provide SHAP-backed reason codes, feature contributions, and human-readable explanation artifacts for each flagged decision.
Audit Verification	Bind input commitments and explainability artifacts to tamper-evident ledger records to support review and evidentiary integrity.
Resilience Controls	Support state-based risk gating and operational resilience during elevated threat or coordinated fraud scenarios.

Buyer Value & Operational Outcomes

- Reduce black-box alert risk by showing why a transaction or network was flagged.
- Support investigator triage with network visualisation and ranked risk factors.
- Improve defensibility of AI-assisted AML/CTF outputs through verifiable audit trails.
- Provide a controlled 90-day pilot route using synthetic and approved sample data.

Suggested 90-Day Pilot Scope

Scope	Screen a synthetic or approved transaction dataset; demonstrate network detection, explainability, and audit export.
Inputs	Transaction records, entity attributes, sanctions/PEP indicators, risk thresholds, and approved business rules.
Outputs	Alert list, risk category, top SHAP factors, network graph, explainability artifact, and audit commitment.
Success Criteria	Analysts can reproduce the reason for selected alerts and verify that the explanation artifact matches its audit commitment.

Technical Notes

- REST API and optional gRPC interface for high-throughput integration.
- Government deployments can use mTLS and OAuth 2.0 client credentials.
- Designed for air-gapped, on-premise or sovereign cloud deployment.
- Can integrate through batch files, API calls, and streaming architectures such as Kafka.

Illustrative Integration Endpoints

POST /v1/ctf/transactions/screen | GET /v1/ctf/alerts/{id} | GET /v1/ctf/network/{entity_id} | GET /v1/audit/artifact/{id}

Contact: Pedro Nguyen | pedro@revoaix.com | +971 58 579 2788 | Booth 8ST-09

02 AI Governance & Compliance Dashboard

Regulator-ready oversight for high-stakes AI decisions

Overview

RevoAI's AI Governance & Compliance Dashboard supports oversight of high-stakes AI systems by turning model outputs into explainable, auditable, and verifiable decisions. It is designed for regulators, agencies, and regulated organisations that need visibility into model versions, decision trails, compliance status, and explainability artifacts.

Target Buyers

SIRA, Cyber Security Council, regulatory advisory bodies, government AI oversight teams, internal audit teams, and organisations deploying AI in high-trust environments.

Core Capabilities

Model Registry	Track deployed AI models, versions, configurations, rule sets, and engine identifiers.
Decision Audit Trail	Search and review AI decisions by model, subject, outcome, risk tier, epoch, or explainability artifact.
Compliance Scoring	Evaluate decisions against configurable rulesets, including explainability, bias review, human override, and data minimization checks.
Regulator Portal	Provide authorised read-only oversight with verifiable access and audit-ready evidence views.
Tamper-Evident Records	Anchor decision artifacts and model identifiers through cryptographic commitments for independent verification.

Buyer Value & Operational Outcomes

- Give regulators and internal auditors a single place to inspect high-stakes AI decisions.
- Reduce governance risk by pairing model outputs with reason codes and audit packets.
- Support explainability, accountability, and oversight without exposing unnecessary sensitive data.
- Prepare agencies for structured AI assurance workflows and model-risk reviews.

Suggested 90-Day Pilot Scope

Scope	Configure one AI decision use case and load 30 days of synthetic or approved sample decisions.
Inputs	Model identifier, decision record, input summary, outcome, risk tier, explainability artifact, and policy ruleset.
Outputs	Compliance scorecard, flagged observations, model registry view, decision audit trail, and verifiable artifact record.
Success Criteria	Authorised reviewers can search a decision, understand the reason chain, and verify artifact integrity.

Technical Notes

- REST API and optional gRPC interface for technical integration.

- Supports model registry, decision audit, compliance evaluation, and artifact retrieval endpoints.
- Designed around explainability artifacts, engine identifiers, and cryptographic commitments.
- Can support on-premise or sovereign deployment models for sensitive environments.

Illustrative Integration Endpoints

GET /v1/governance/models | GET /v1/governance/decisions/{id} | POST /v1/governance/compliance/evaluate | GET /v1/audit/artifact/{id}

Contact: Pedro Nguyen | pedro@revoaix.com | +971 58 579 2788 | Booth 8ST-09

03 Security Workforce Verification

Privacy-preserving capability verification and readiness analytics

Overview

RevoAI's Security Workforce Verification solution helps security agencies verify credentials, readiness signals, and personnel capability indicators while preserving privacy. It combines explainable capability scoring, telemetry-based verification, SHAP attribution, and zero-knowledge predicate checks for controlled workforce assessment and training decisions.

Target Buyers

Abu Dhabi Police, National Guard, MoI HR divisions, training academies, personnel vetting units, and security organisations requiring trusted qualification verification.

Core Capabilities

Capability Scoring	Create multi-dimensional capability vectors from training, certification, service, and performance indicators.
Privacy-Preserving Verification	Verify identity, liveness, clearance level, and qualification thresholds without unnecessary exposure of raw sensitive records.
Explainable Analytics	Show which verified factors contributed to a capability or readiness score using SHAP-style attribution.
Readiness Dashboard	Provide unit-level views, skills gap indicators, training recommendations, and aggregate readiness signals.
Audit Trail	Generate explainability artifacts and commitment records so workforce decisions can be reviewed and verified.

Buyer Value & Operational Outcomes

- Shorten verification workflows by connecting evidence, scoring, and audit review in one system.
- Reduce privacy risk by verifying predicates rather than exposing unnecessary personal data.
- Support training investment decisions using explainable capability trends and skills gaps.
- Provide an introductory pilot route using mock profiles and approved datasets only.

Suggested 90-Day Pilot Scope

Scope	Load mock personnel profiles and demonstrate capability scoring, radar chart, SHAP factors, and ZK predicate verification.
Inputs	Training records, certifications, fitness indicators, reviews, attestations, timestamps, and signatures.
Outputs	Capability vector, explanation factors, predicate proof results, ledger commitment, and readiness analytics.
Success Criteria	Reviewers can understand why a capability score was produced and verify credentials without exposing raw sensitive details.

Technical Notes

- Supports workforce assessment, profile retrieval, and predicate verification endpoints.
- Can verify identity and clearance predicates while limiting exposure of underlying sensitive attributes.
- Designed for controlled pilot use with synthetic or agency-approved records.
- Architecture can align with sovereign deployment, access controls, and audit requirements.

Illustrative Integration Endpoints

POST /v1/workforce/assess | GET /v1/workforce/profile/{id} | POST /v1/workforce/verify | GET /v1/audit/artifact/{id}

Contact: Pedro Nguyen | pedro@revoaix.com | +971 58 579 2788 | Booth 8ST-09

See every decision. Trust every model.

RevoAI • Explainable • Cryptographic • Privacy-Preserving • Sovereign