HCLTech | Supercharging Progress™

AI-Augmented Evidence

Empowering Analysts in Next-Gen Fraud Management



HCLTech Fraud Risk Management

Protecting revenue, ensuring compliance
and shaping the AI era

Speaker: Ivano Cassis

Revenue Intelligence User Forum – 2026 – Milan, Italy

The Expanding Blast Radius of Telecommunications Fraud

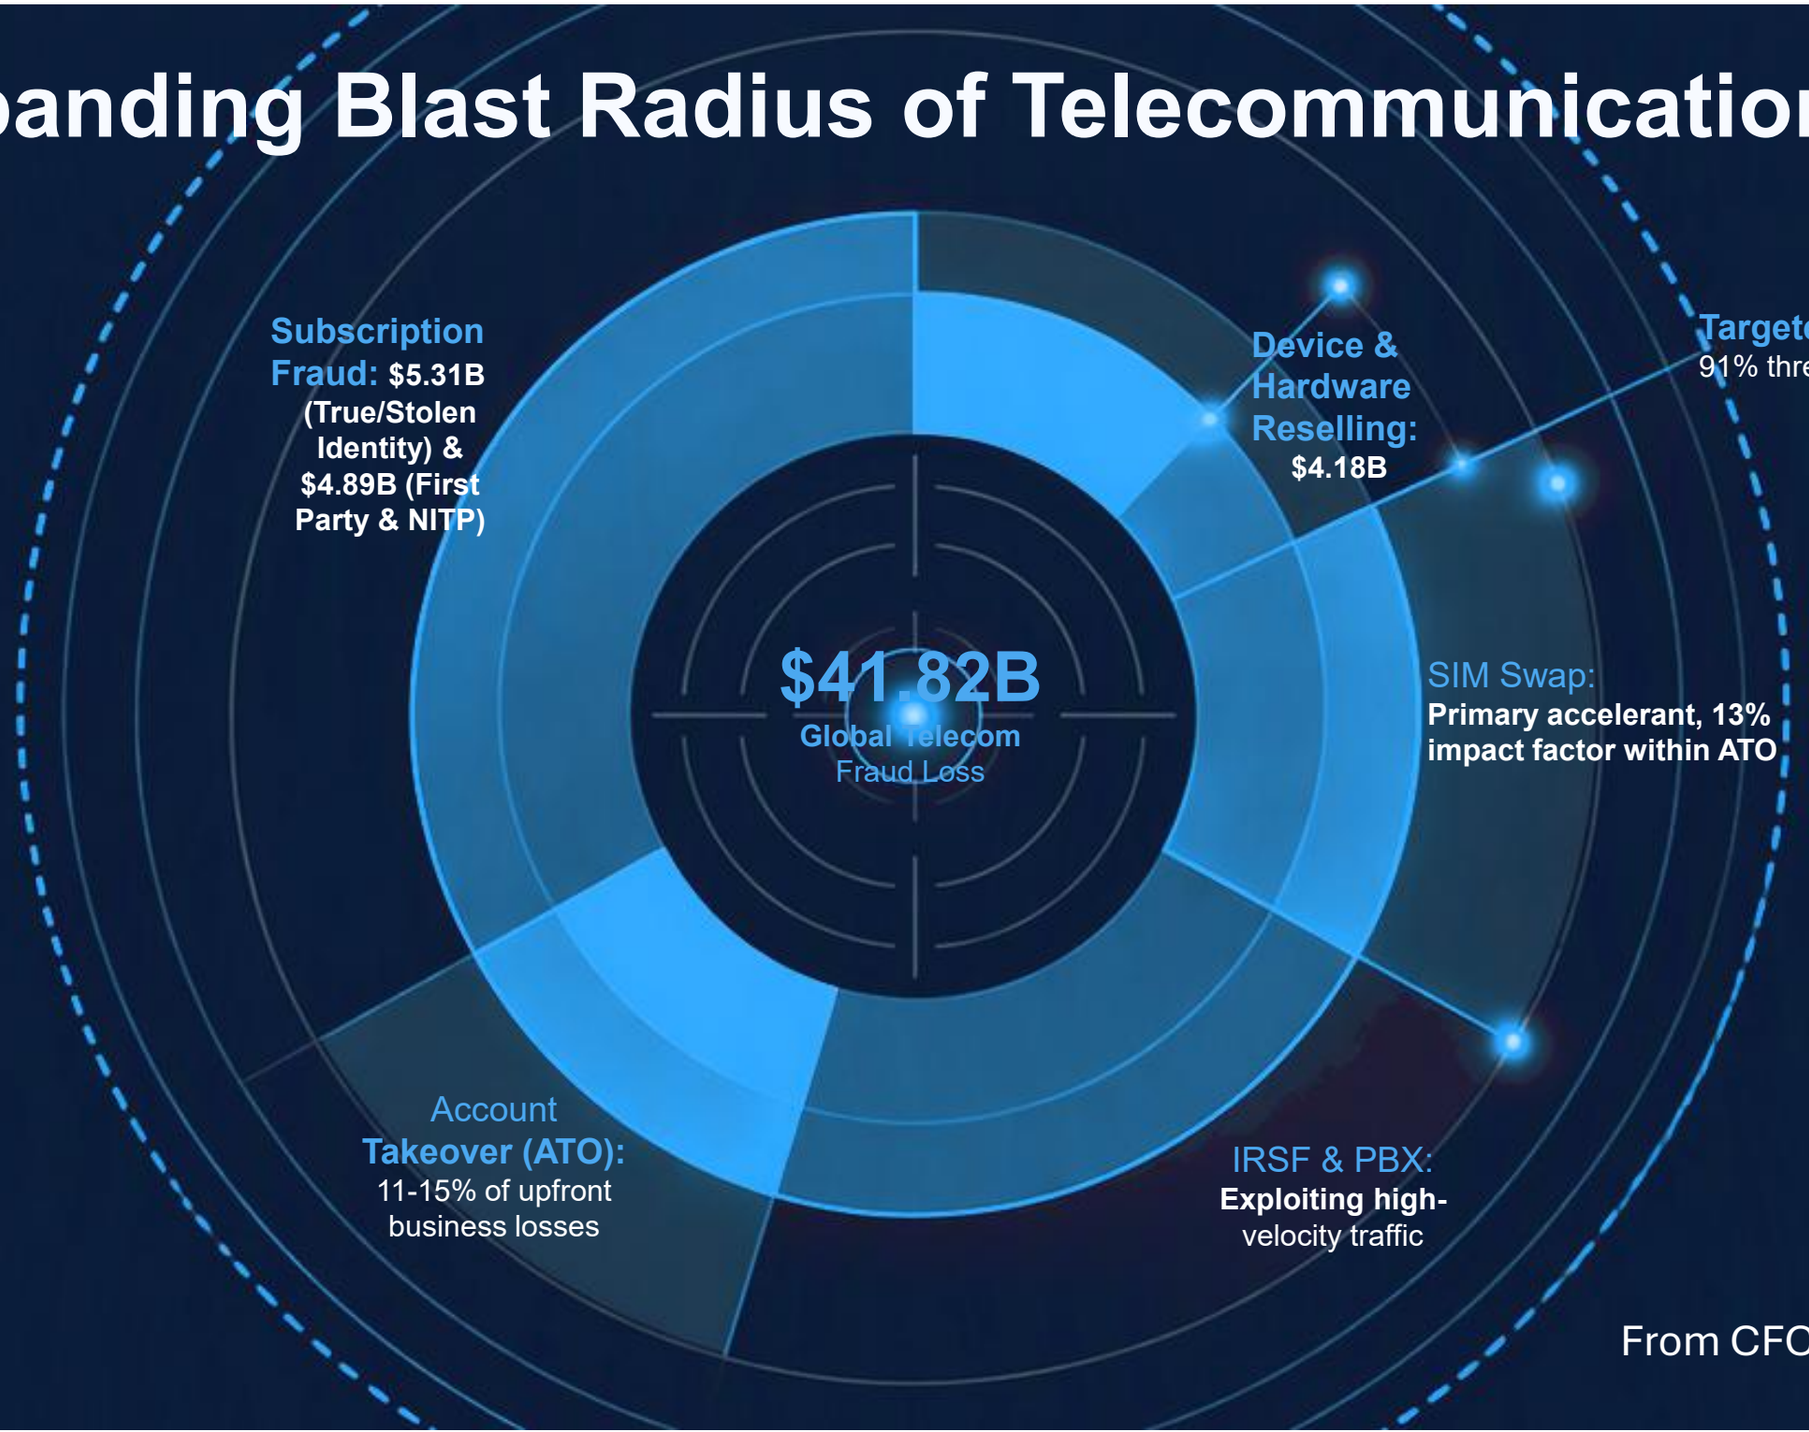

From CFCA 2025 Report

31+ Years of Revenue Intelligence

1994-2004
Detect Usage
Mobile, wireline, and convergent services fraud detection.

2007-2010
Subscription
Dealer activity and subscriber fraud prevention.

2014-2018
Analytics
Real-time scoring and early AI/ML workflow design.

2023-2024
Technology
Real-time cloud operations, behaviour propensity, and anomaly detection.

2025-2026
AI Agentic
Chat, Agentic Crew Orchestration, Domain knowledge, AI findings and Hunting the Unknown

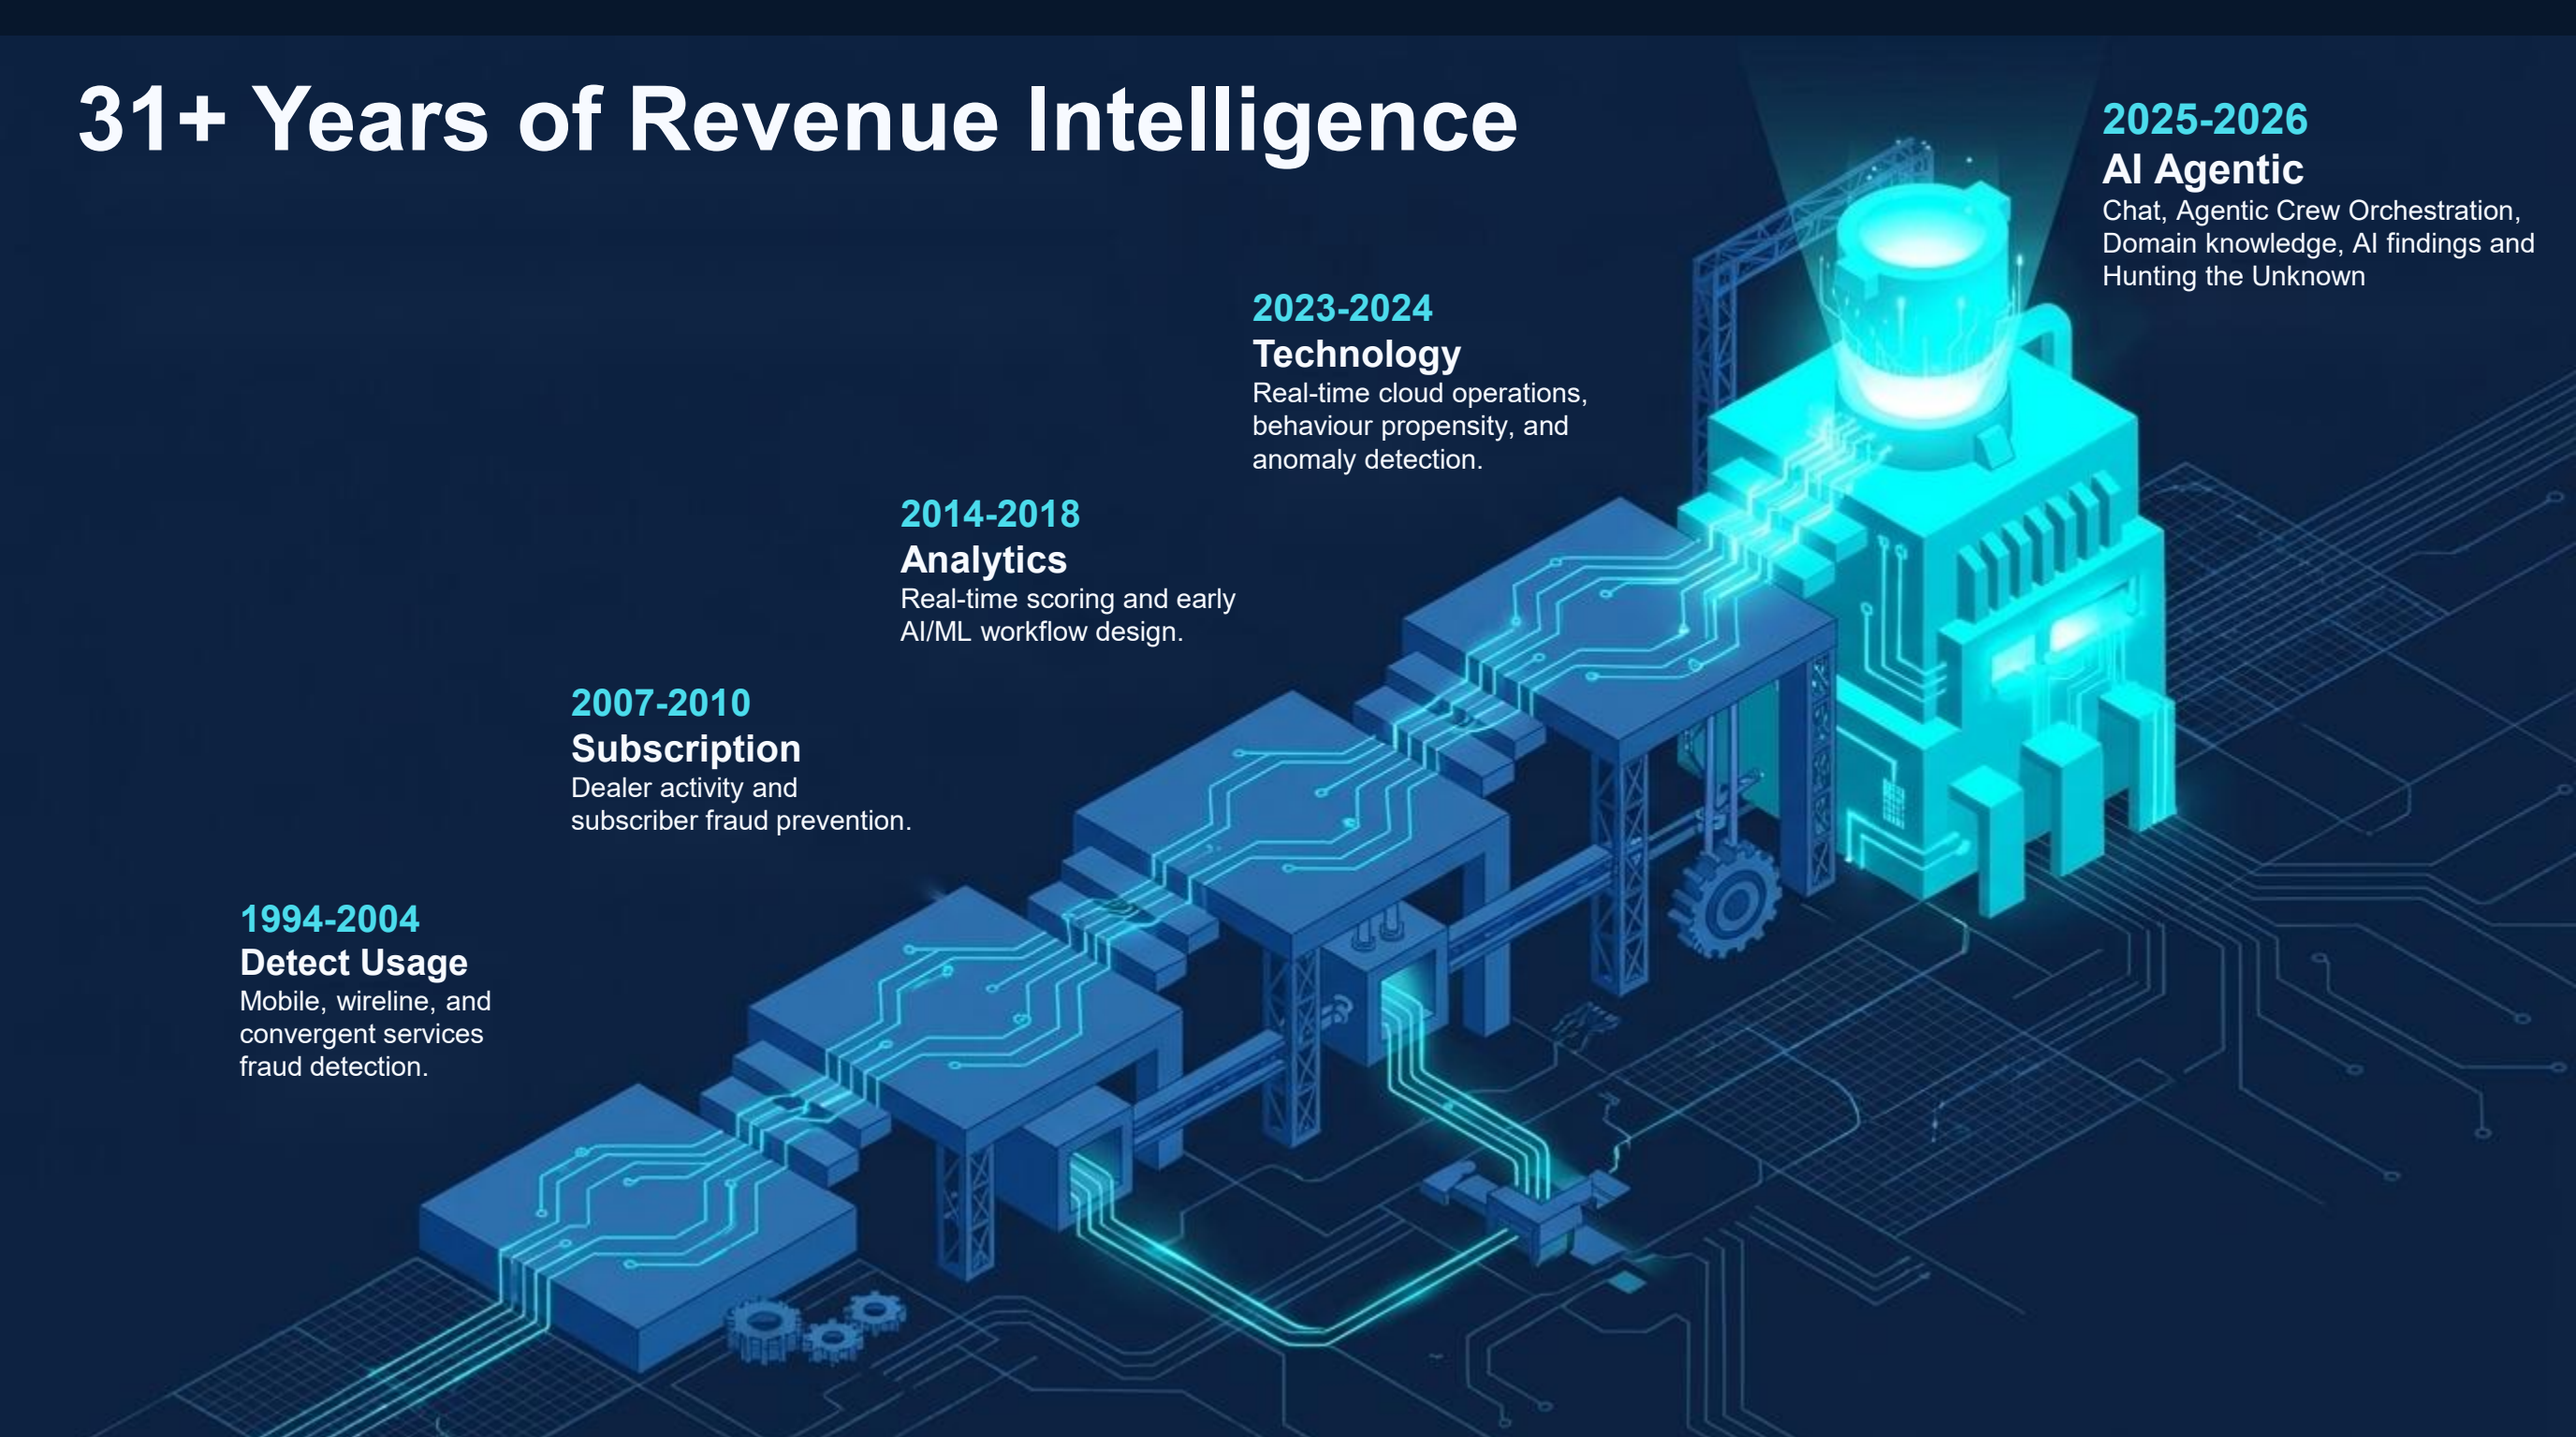

Encoding Expertise Without Losing Control

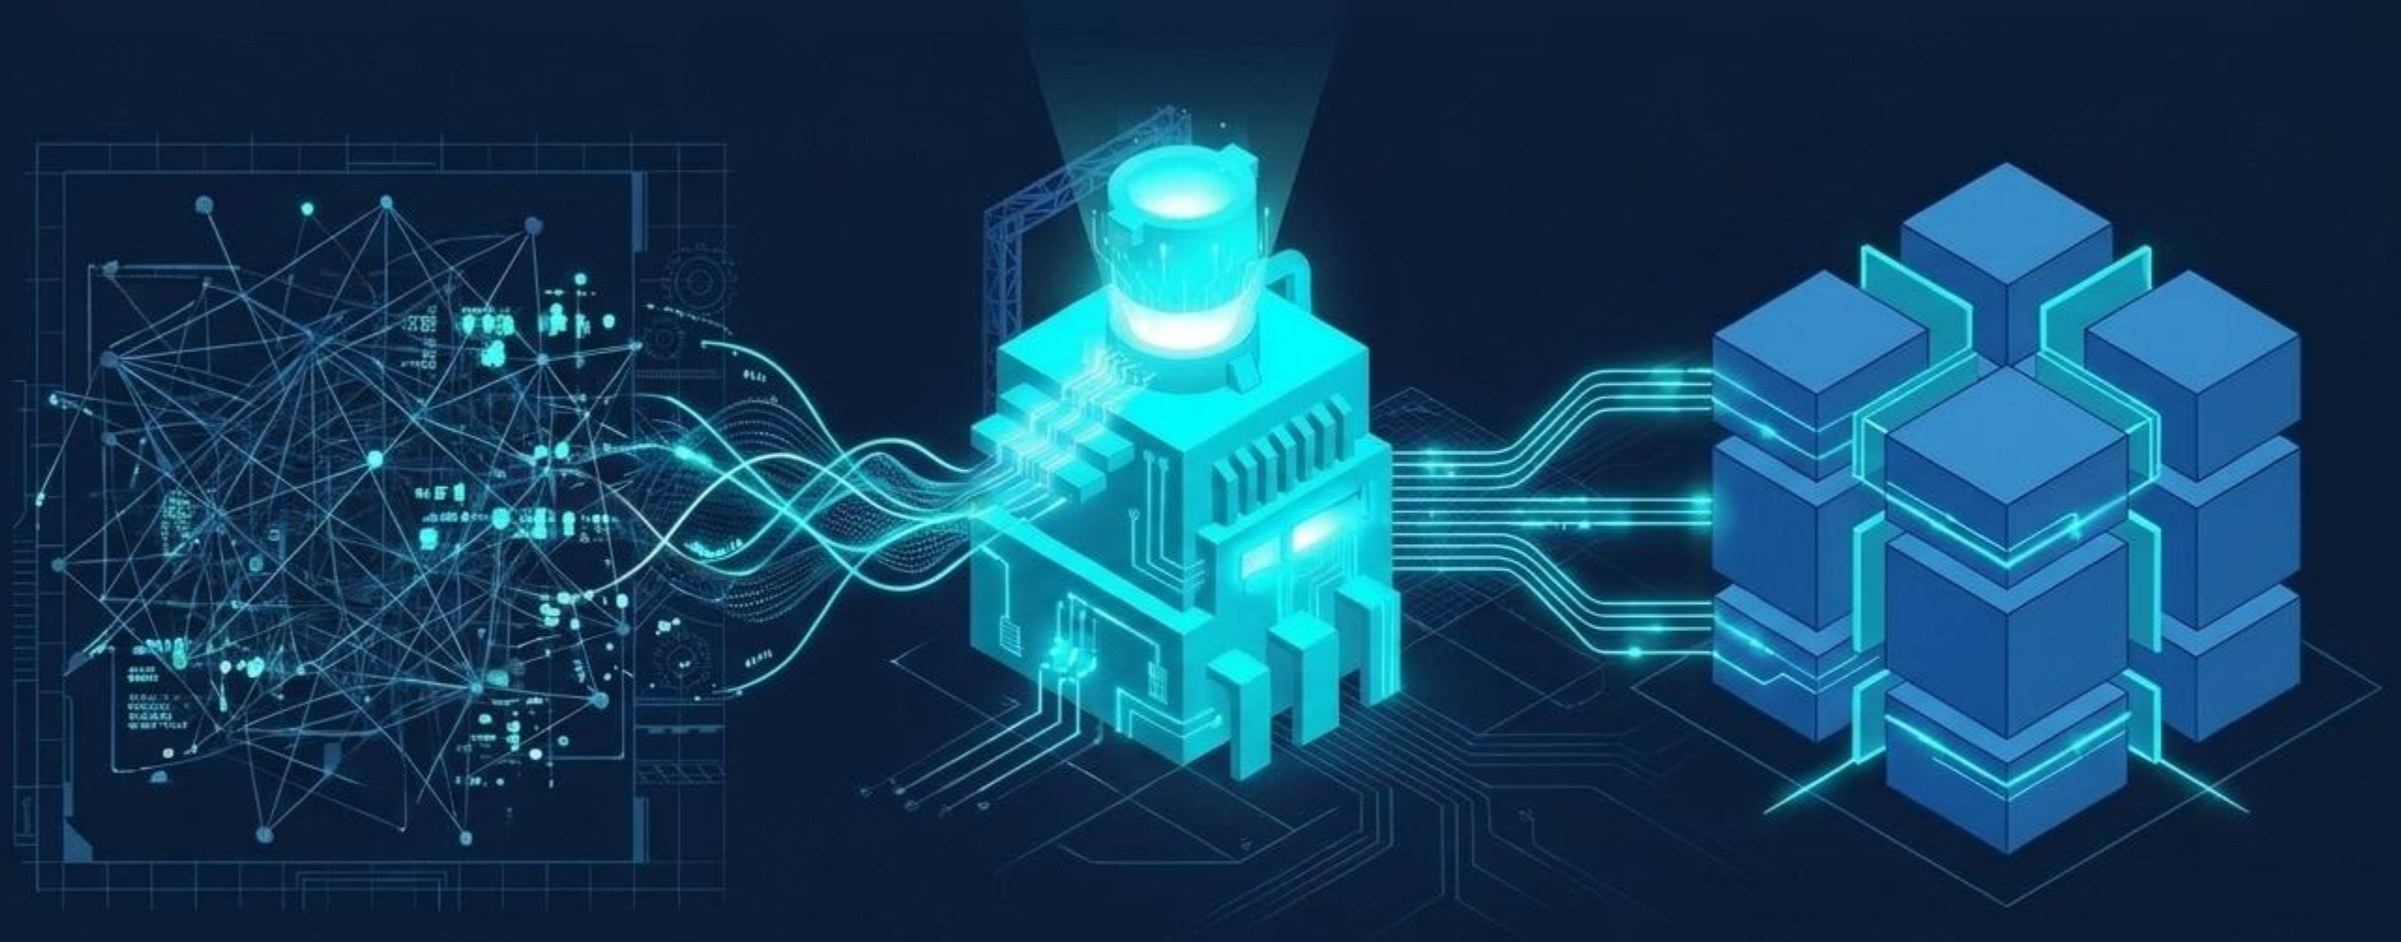

The Dilemma
Manual rule engineering cannot keep pace with evolving fraud patterns while AI black-box decisions are difficult to explain and govern.

The Transformation
Domain-Tuned Small Language Models (SLMs) translate fraud scenarios, expert knowledge and known typologies into structured detection logic.

The Deterministic Advantage
The resulting validated logic execution remains deterministic, secure, explainable, and audit-proof

The Domain Knowledge

Scenario (Fraud Typology & Model)

Structured Attack Pattern: Standardized specification of a fraud scenario
Detection Logic: Defines the threat typology, case evaluation window, and investigative rationale.
Portability: Decouples fraud rules from installation-specific network tokens for cross-platform standardization.

Evidence & Counter-Evidence (Confidence & Validation)

Characteristic Evidence: Mandatory structural indicators required to confirm a fraud pattern
Corroborating Evidence: Secondary signals that increase case confidence
Counter-Evidence: Legitimate behavior signals

Observable (Network Metrics & Indicators)

Indicators: Atomic, measurable data points extracted from event detail records.
Rules & Thresholds: Quantitative parameters (min/max bounds, ratios, counts) evaluated over time windows



Domain Knowledge Base
Centralized repository storing standardized scenarios, observable catalogues, and action definitions.

Signal Indicators
Real-time monitoring rules tracking cumulative observable thresholds across defined time windows

Verified case outcomes
Result of investigation used for tuning and refinement

AI-Assisted Case Handling: The Analyst's Exoskeleton

Multi-Agent Crew

Aggregate system findings, evidence, score and explain the case



Case Dashboard



AI Finding

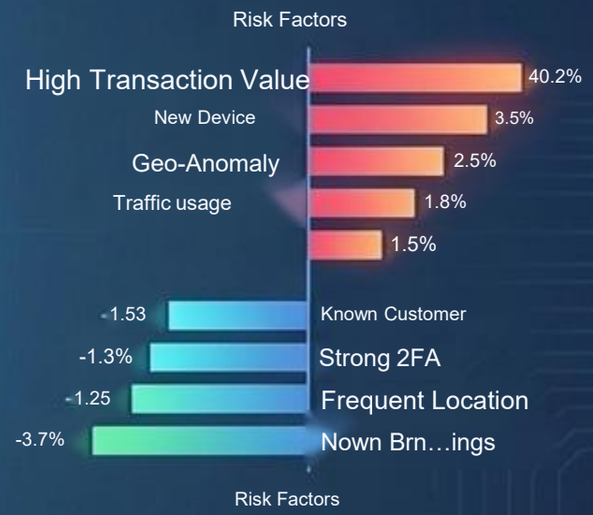

Breaks down the exact features that triggered the alert using SHapley Additive exPlanations.

Disposition Action

Suggests concrete remediation steps:

- Suspend Service
- Contact Customer
- Flag Device

Hunting the Unknown

Discovering emerging patterns

AI Profiling

Ai crew combine indicators, ML models analyse subscriber behaviour to identify emerging patterns, unusual usage clusters and deviations from established norms.

AI Assisted Explanation

When an unexplained anomaly is surfaced, it is routed into the AI-Assisted Handling pipeline for analyst review

Human-in-the-Loop Validation

Analysts review the anomaly together with supporting evidence,historical context and explainability artefacts before validating the discovery.



The HCLTech FRM Execution Engine

Carrier-grade detection at deterministic scale

Execution Velocity

Processes hundreds of millions of EDRs and network events per day with predictable performance and deterministic outcomes.



Multi Dimensional Correlation

Evaluates behaviour across time, destinations, usage patterns, velocity indicators and historical context simultaneously.



Decision

Applies fraud patterns, rules and other detection logic to generate traceable and reproducible alerts.

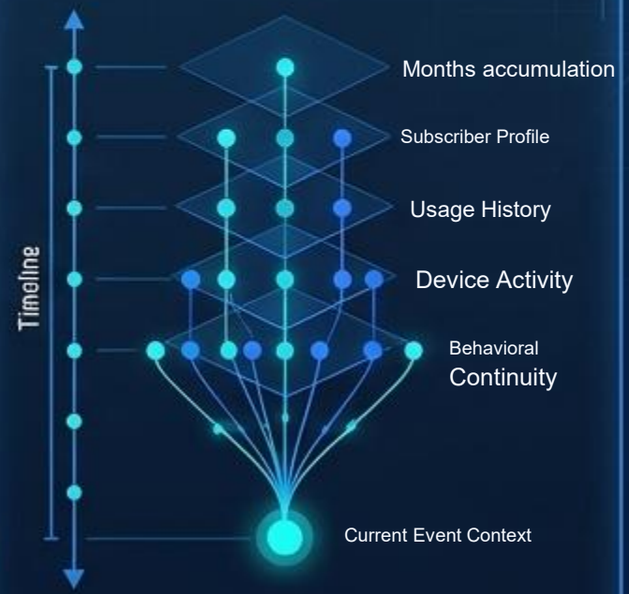

The journey

1994-2004
Detect Usage
Mobile, wireline, and convergent services fraud detection.

2007-2010
Subscription
Dealer activity and subscriber fraud prevention.

2014-2018
Analytics
Real-time scoring and early AI/ML workflow design.

2023-2024
Technology
Real-time cloud operations, behaviour propensity, and anomaly detection.

2025-2026
AI Agentic
Chat, Agentic Crew Orchestration, Domain knowledge, AI findings and Hunting the Unknown

2027- The Closed-Loop Milestone

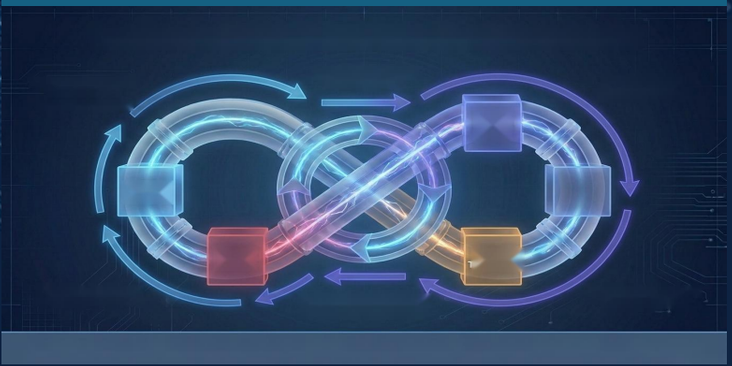

The Closed-Loop Milestone

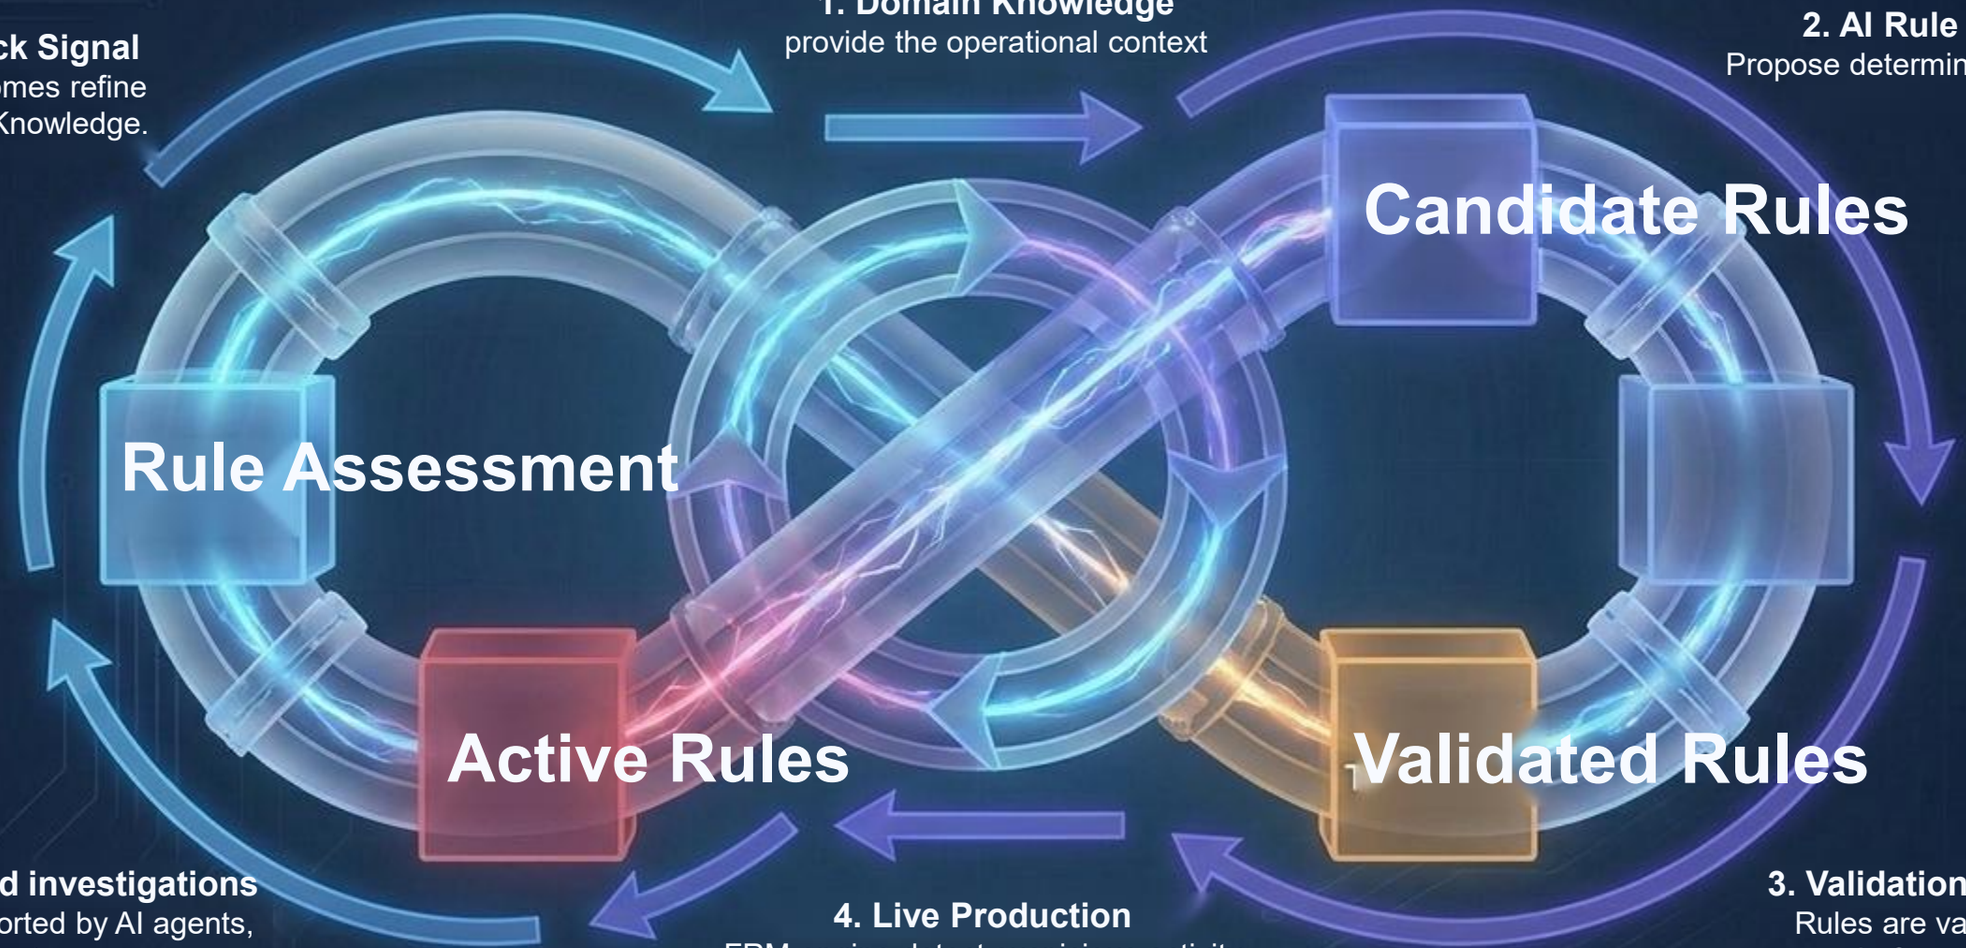

A governed lifecycle that transforms fraud expertise into validated detection logic, executes it deterministically at carrier-grade scale, and improves through verified operational feedback.



Ivano Cassis
Revenue Intelligence R&D Manager | HCLTech

Mobile: +39 335 7390943
Email: ivano.cassis@hcltech.com

DEMOs

AI-Assisted Case Handling: The Analyst's Exoskeleton
Hunting the Unknown: discovery emerging pattern
Realtime discovery graph:
Dashboard: Dealer

Classification: Internal

FRM | LLM integration Approach | step 1: BYOM

BYOM enables customer-owned model selection while FRM provides a consistent integration layer.

Integration approach

FRM connects to the model selected and operated by the customer through an API-based adapter layer.

External LLM services VALIDATED

Product integration validated with ChatGPT and Claude.

Local / private LLMs VALIDATED

Local model scenarios validated with OpenAI GPT-OSS and Qwen3.

Customer-selected LLM EXTENSIBLE

Additional models can be integrated efficiently when they expose the OpenAI-compatible API standard.



Built-in AI governance controls

Cost & SLA controls
Usage, cost and service-level governance

Privacy enforcement
Policy-based privacy safeguards

Auditing module
Traceability and governance evidence

HCLTech | Internal
© Internal © 2026 HCL Tech | Confidential

HCLTech

Continuous Rule Optimization

AI Rule Auditor

The Rule Bloat challenge

Over time, rulebases accumulate ineffective, redundant and obsolete detection logic, increasing operational complexity and false-positive noise.

Governed Optimization

AI-generated optimisation recommendations are reviewed before updates are applied to the production rulebase.

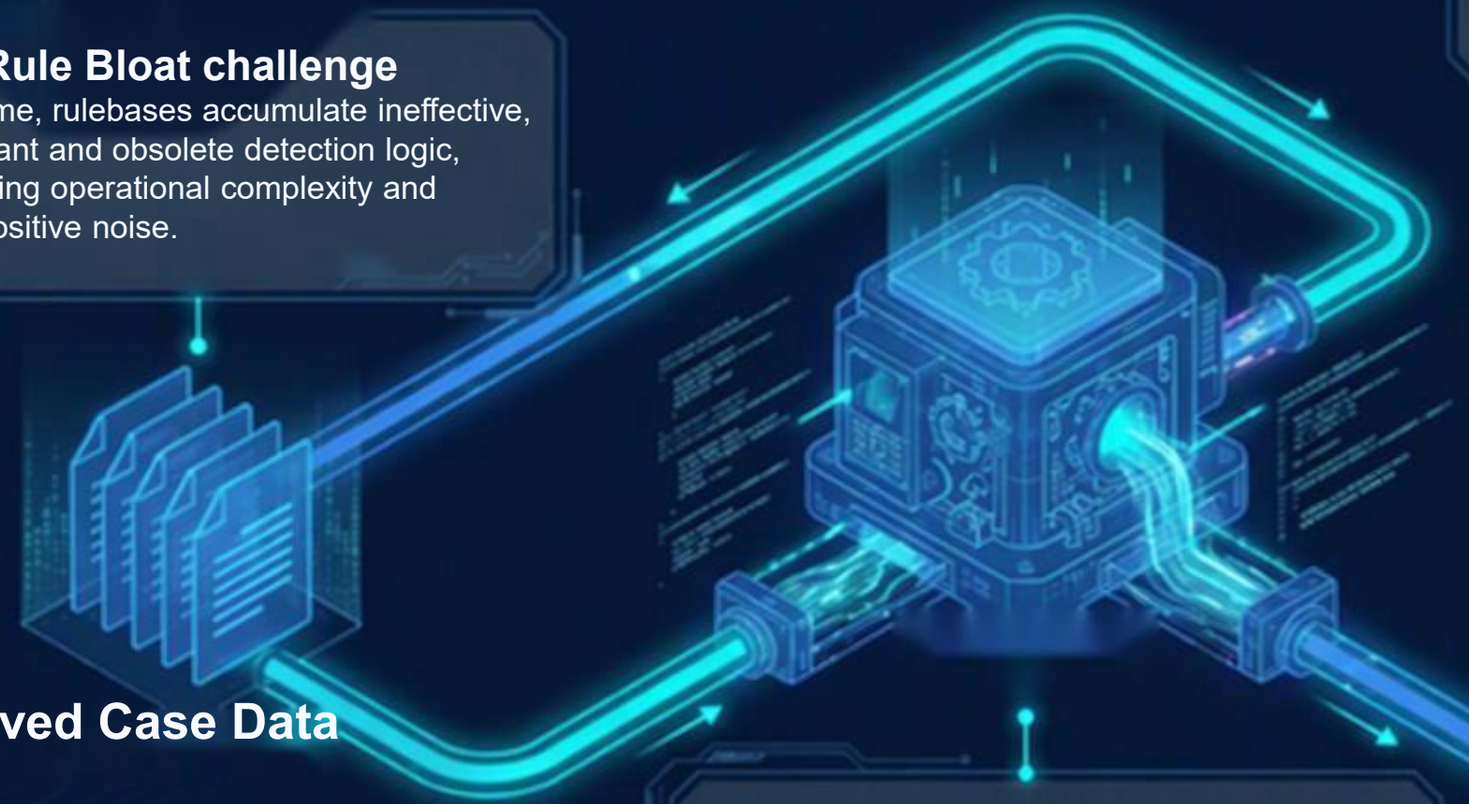

Archived Case Data

Inactive Rules

Redundand Logic

Recommendations

Optimization Candidates

Rule Tree

Agentic Auditing

AI agents analyse historical alerts and compare rule triggers with verified case dispositions to evaluate rule effectiveness and precision.

Frictionless Transition from Sandbox to Production

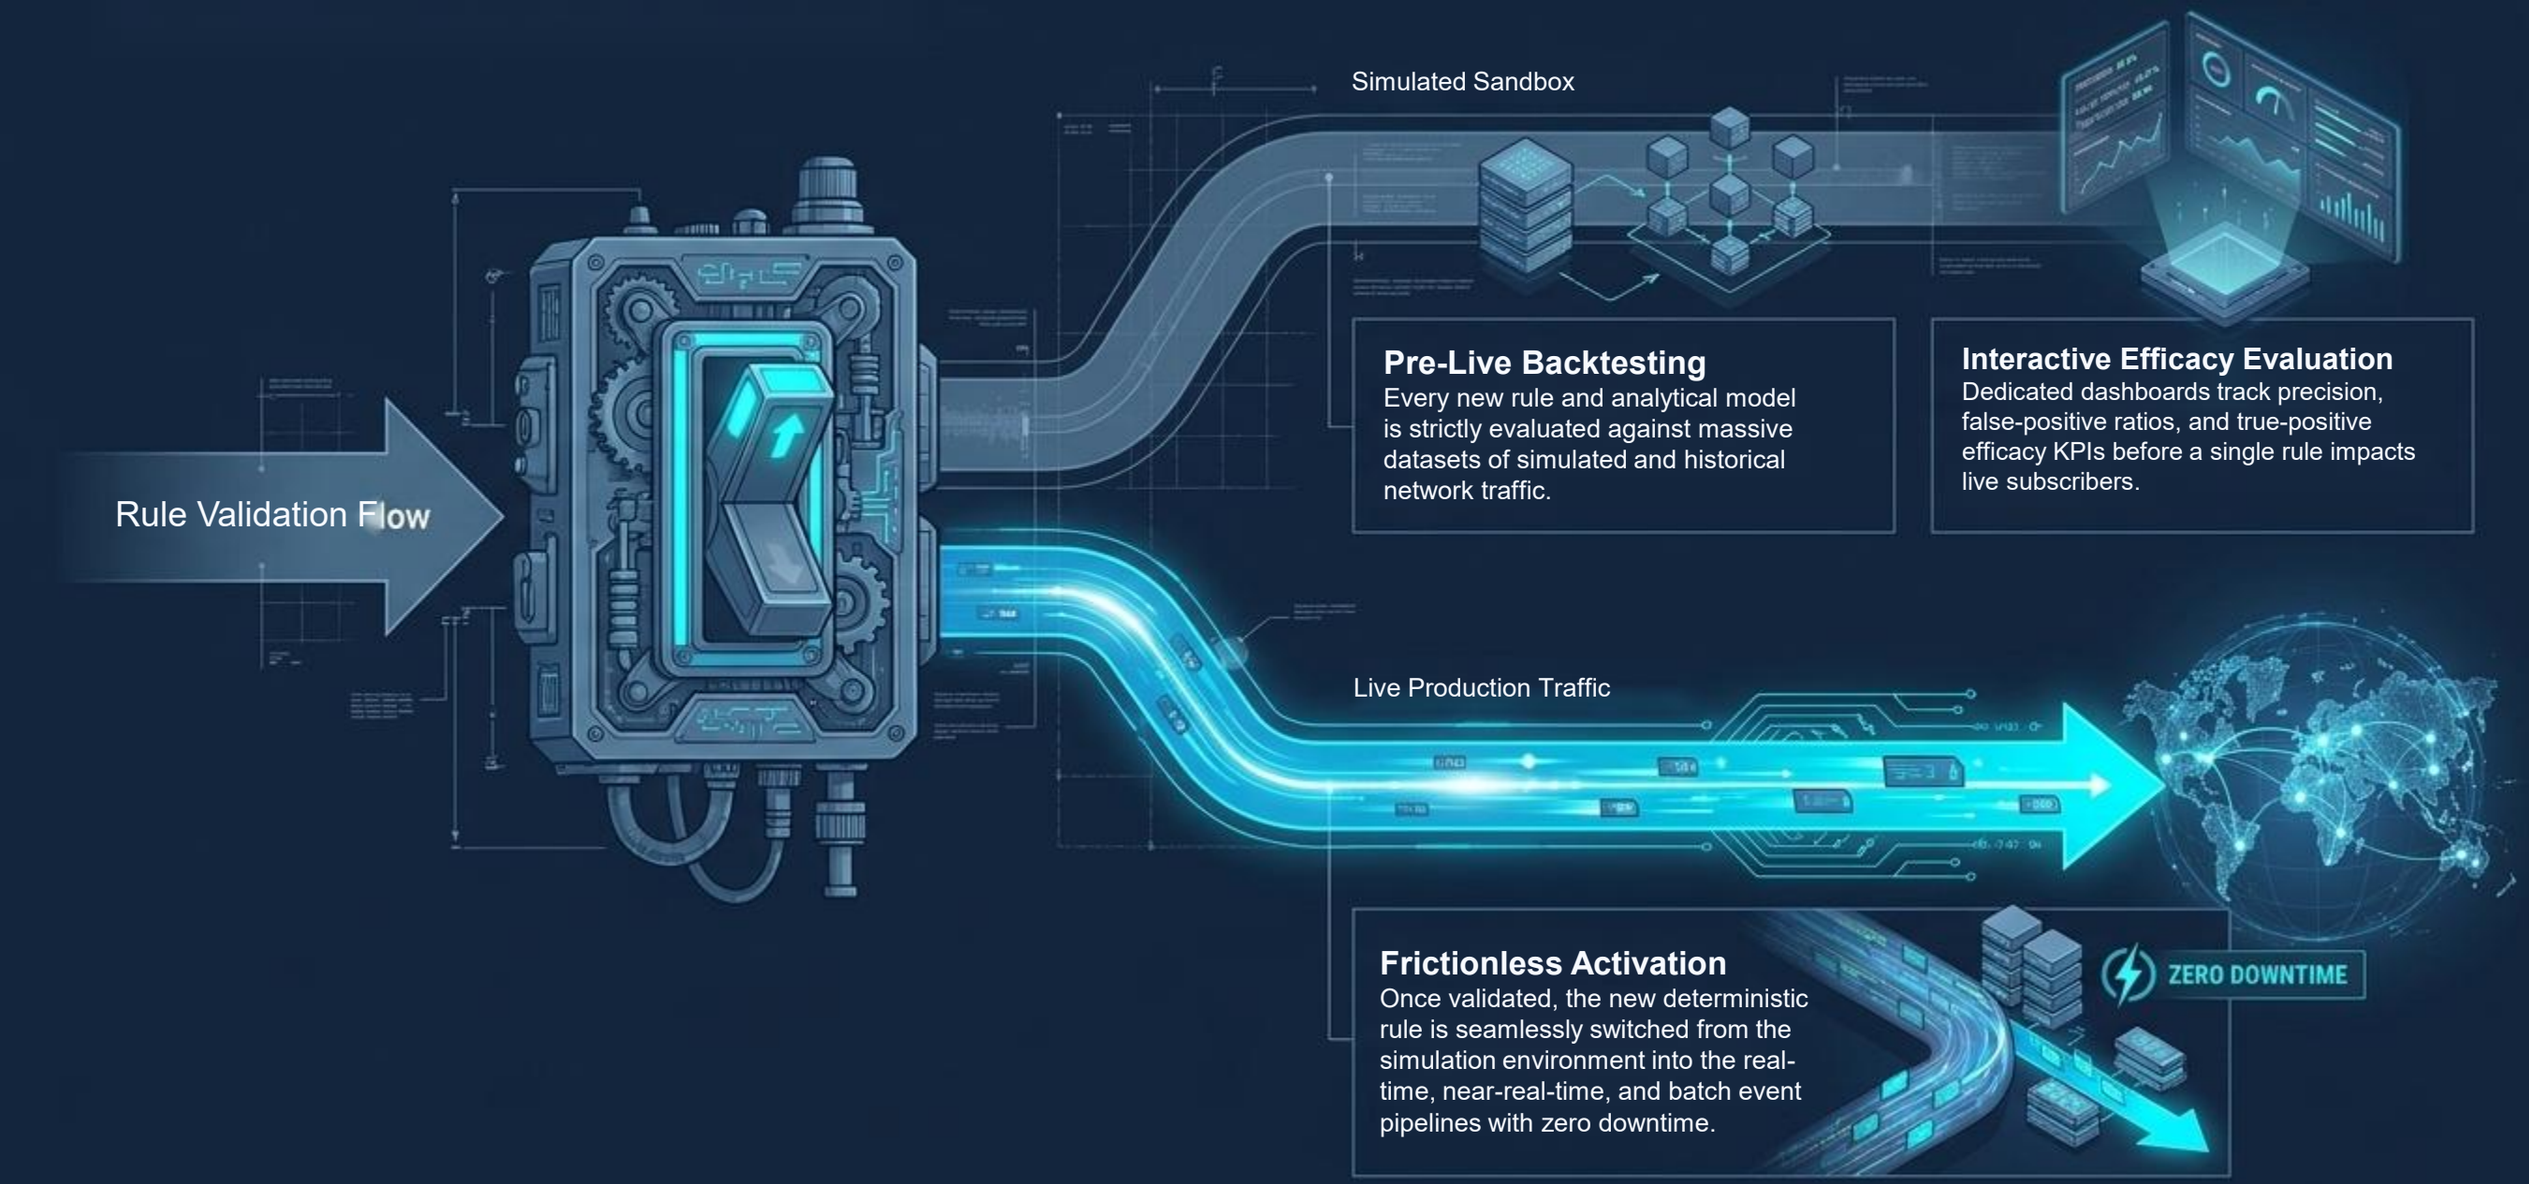

Pre-Live Backtesting
Every new rule and analytical model is strictly evaluated against massive datasets of simulated and historical network traffic.

Interactive Efficacy Evaluation
Dedicated dashboards track precision, false-positive ratios, and true-positive efficacy KPIs before a single rule impacts live subscribers.

Frictionless Activation
Once validated, the new deterministic rule is seamlessly switched from the simulation environment into the real-time, near-real-time, and batch event pipelines with zero downtime.










Adaptive Domain Knowledge Refinement

Knowledge distribution
Once the knowledge is updated the loop restart.
Approved knowledge updates if world-wide valid can be propagated across customer installations, fortifying the entire ecosystem against identical future attacks.

The Refinement Brain
These signals are reviewed and approved before being incorporated into the Domain Knowledge Base, Adapting the knowledge to the environment.

Signal discovery
Polling the latest victories and mistakes from recent Cases resolution to extract indicators that distinguish fraud vs. legitimate behaviour.

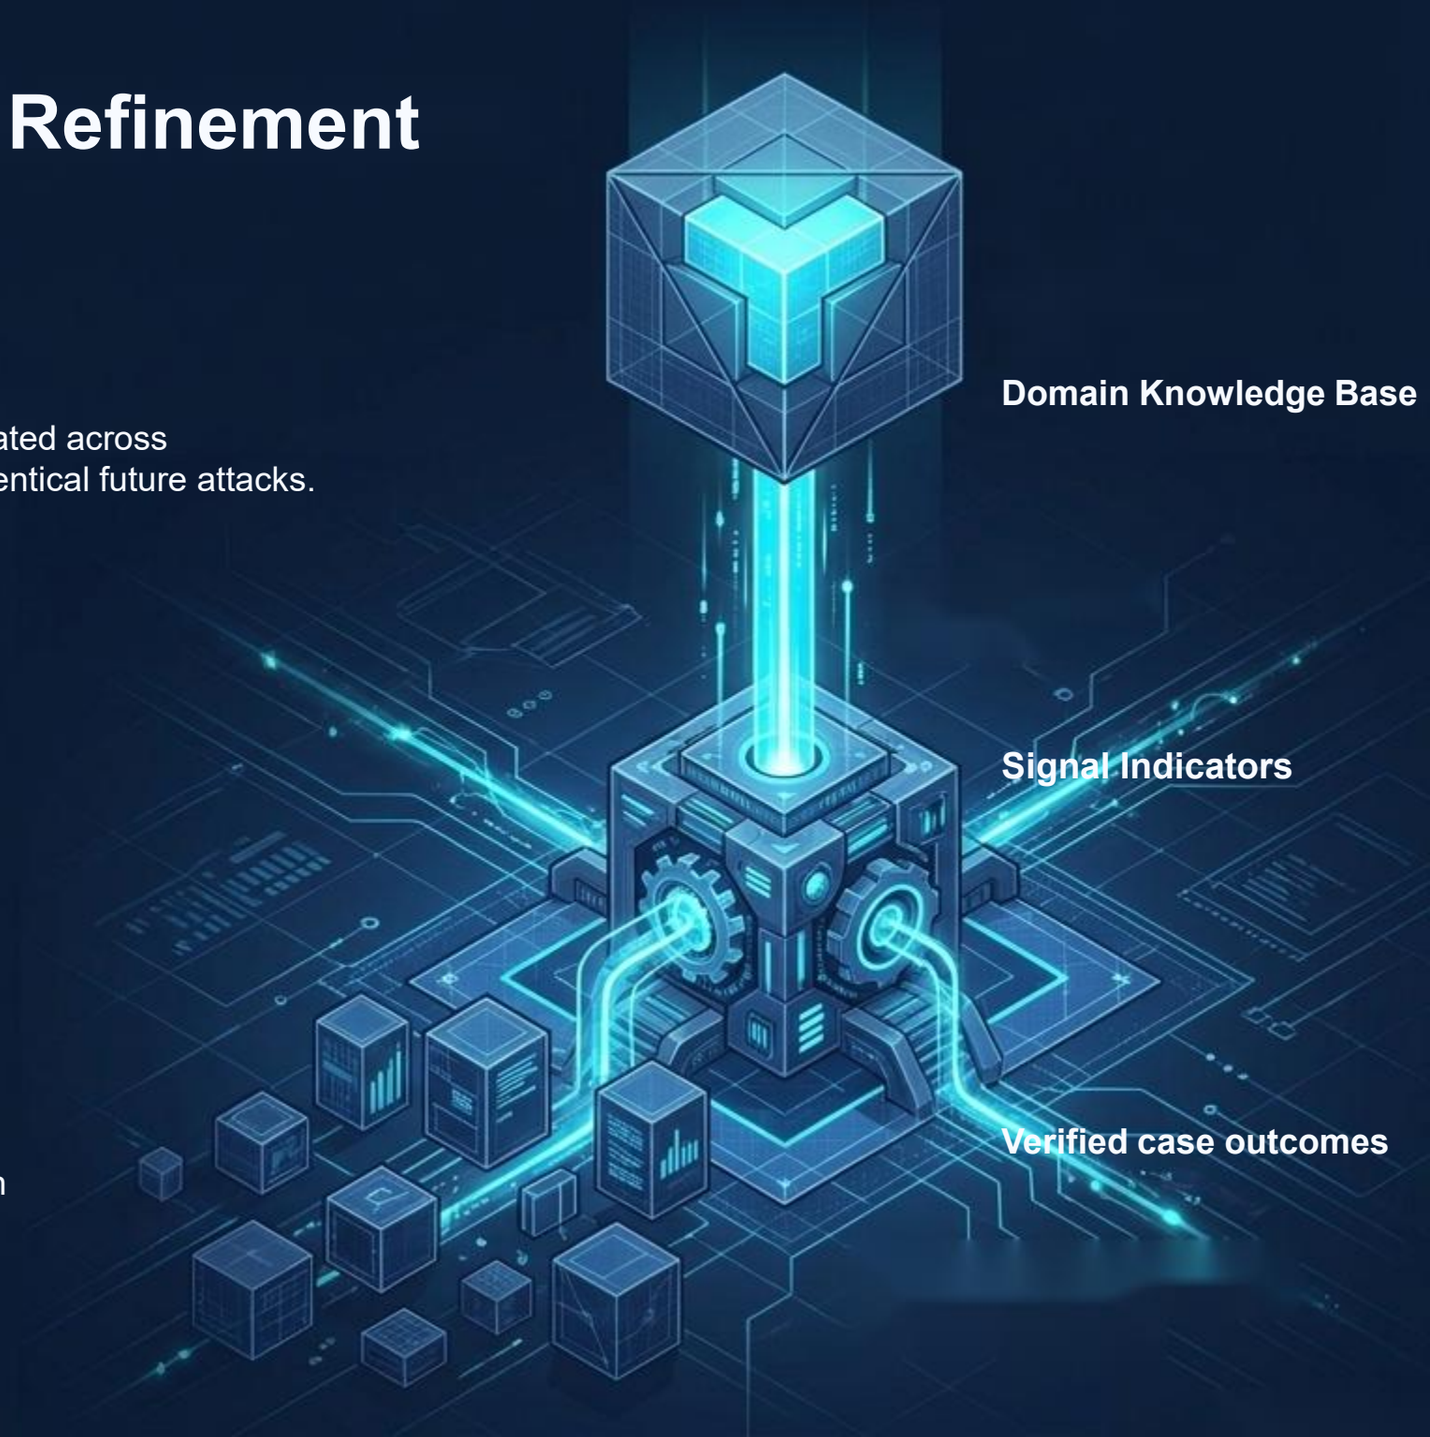

HCLTech's Closed-Loop Fraud Intelligence Cycle

From Where We Are

Manual Outcome Review
Analysts review verified case outcomes and manually refine knowledge, rules and thresholds.

Domain Knowledge Base
Fraud scenarios, typologies, and historical intelligence.

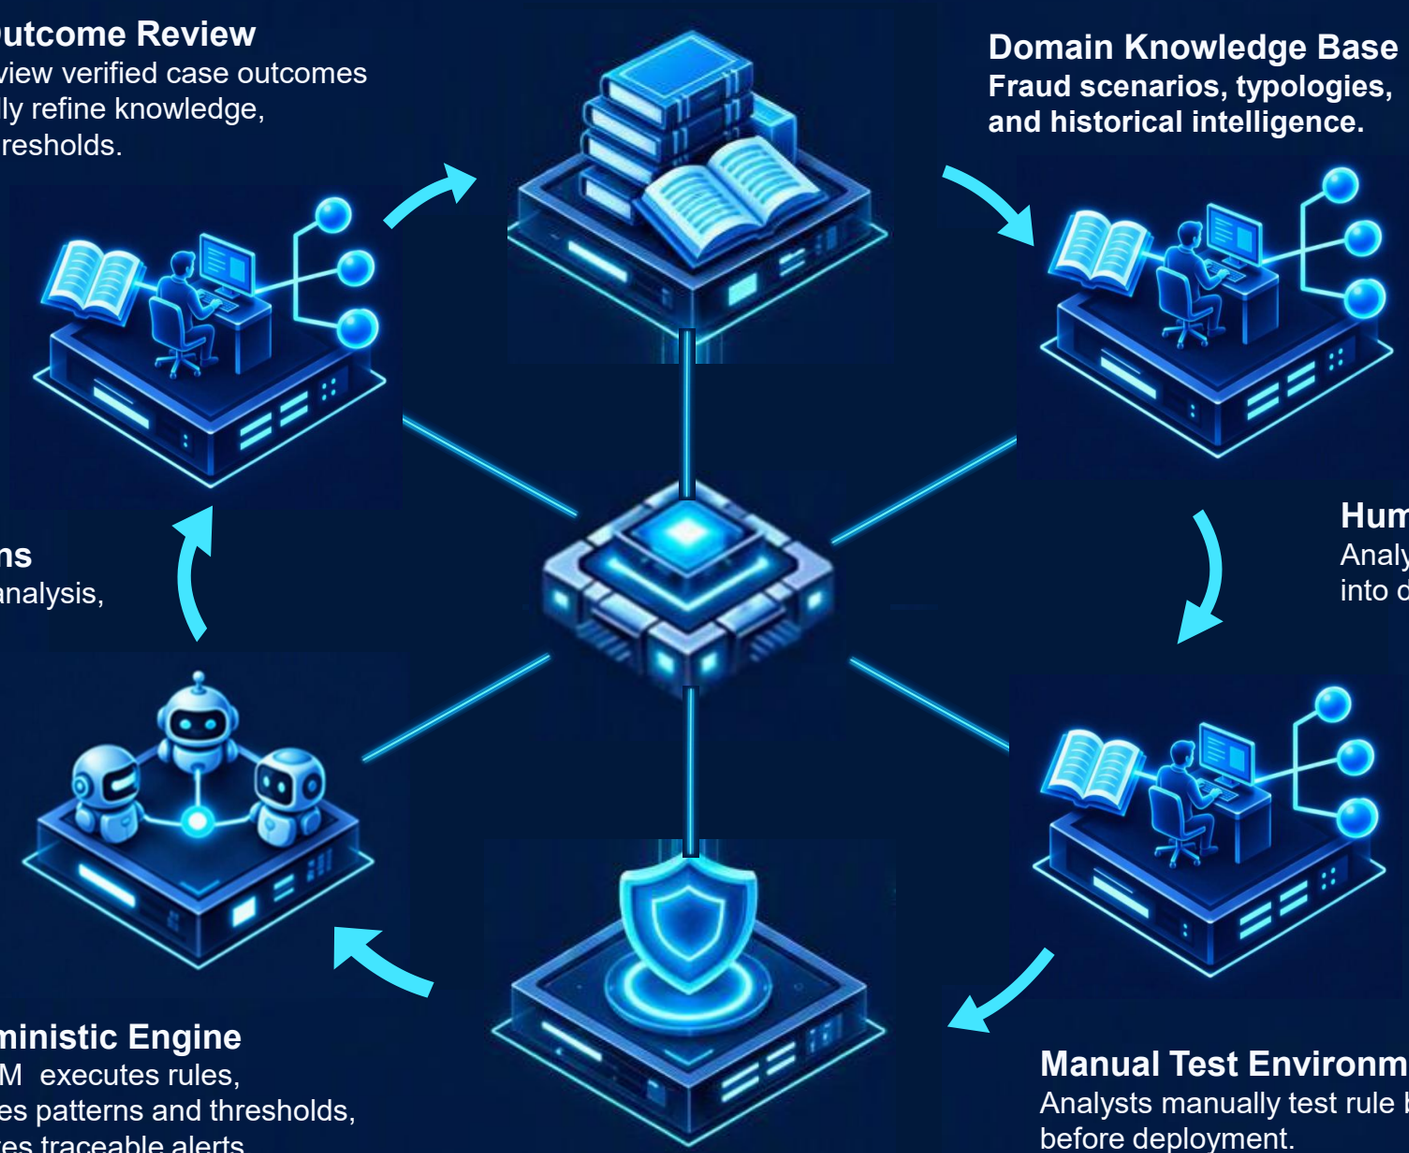

Human Rule Engineering
Analysts translate fraud scenarios into detection rules and patterns.

AI-Assisted Investigations
Multi-agent AI supports case analysis, evidence gathering and analyst decision-making.

Deterministic Engine
Our FRM executes rules, evaluates patterns and thresholds, generates traceable alerts.

Manual Test Environment
Analysts manually test rule behaviour before deployment.

HCLTech's Closed-Loop Fraud Intelligence Cycle

To Where We Are Going

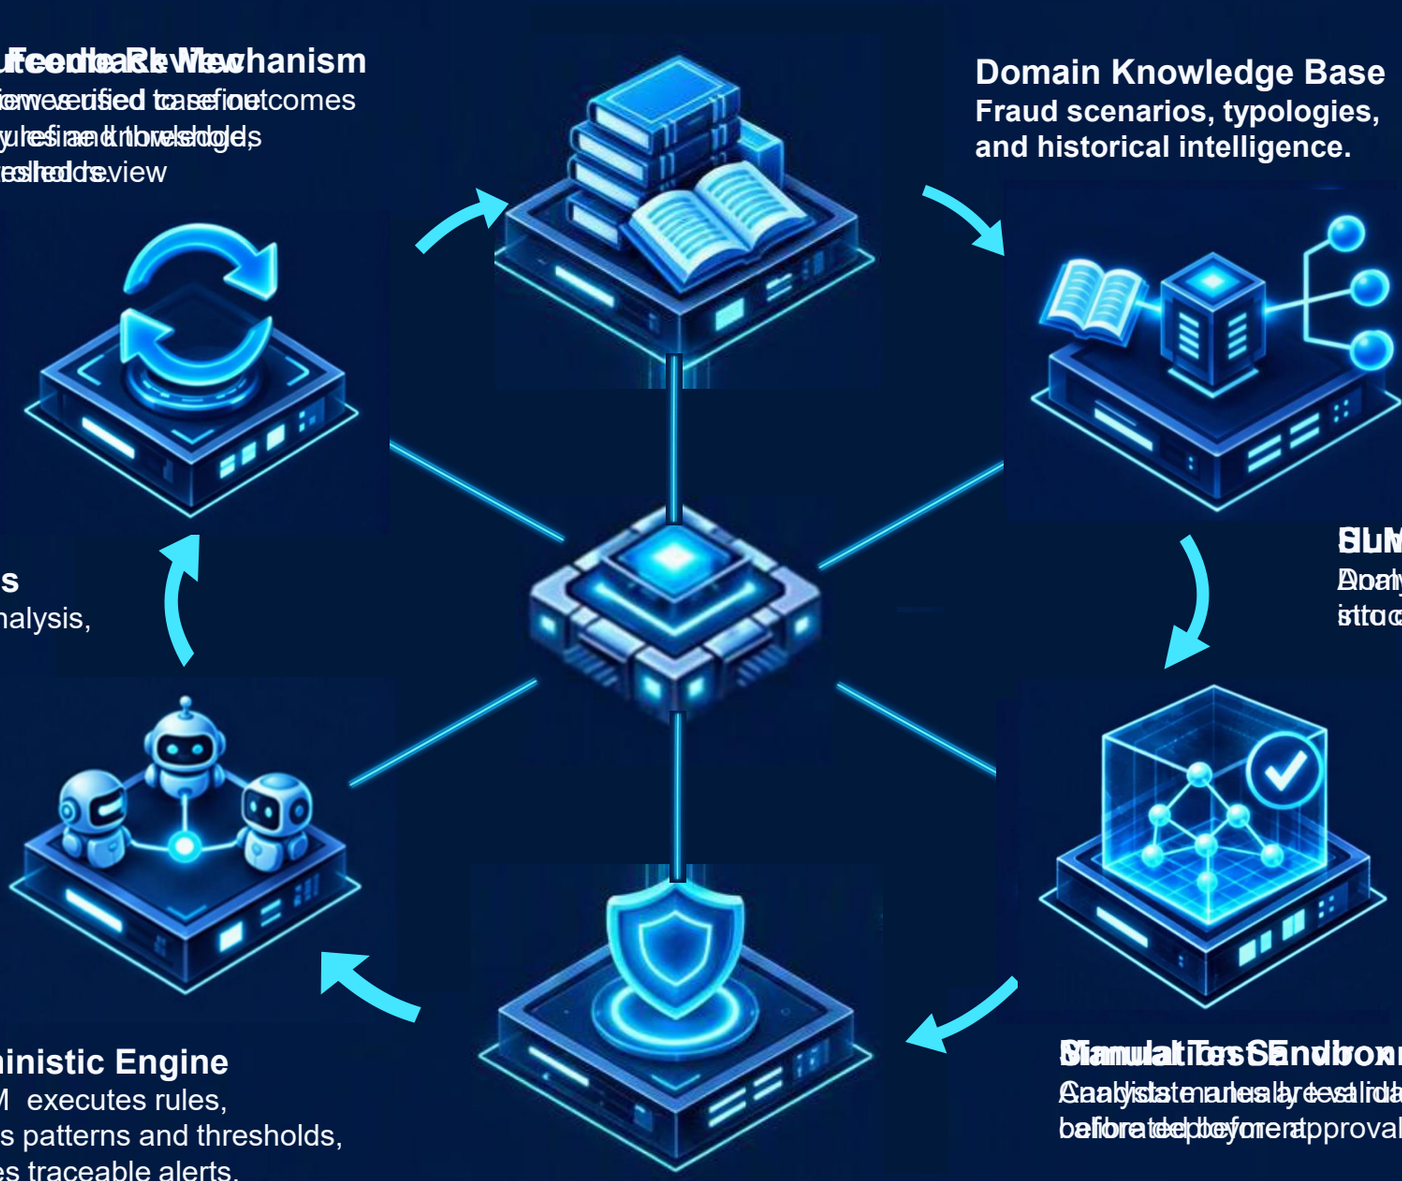

Domain Knowledge Base
Fraud scenarios, typologies, and historical intelligence.

SLM-Assisted Rule Generation
Analysts translate fraud scenarios into detection rules and patterns.

Simulation Sandbox
Analysts manually test rules and before deployment.

Manual Test Environment
Analysts manually test rules and before deployment.

Deterministic Engine
Our FRM executes rules, evaluates patterns and thresholds, generates traceable alerts.

AI-Assisted Investigations
Multi-agent AI supports case analysis, evidence gathering and analyst decision-making.

Manual Outcome Review
Analysts review verified case outcomes and manually refine knowledge, rules and thresholds.

Manual Outcome Review
Analysts review verified case outcomes and manually refine knowledge, rules and thresholds.

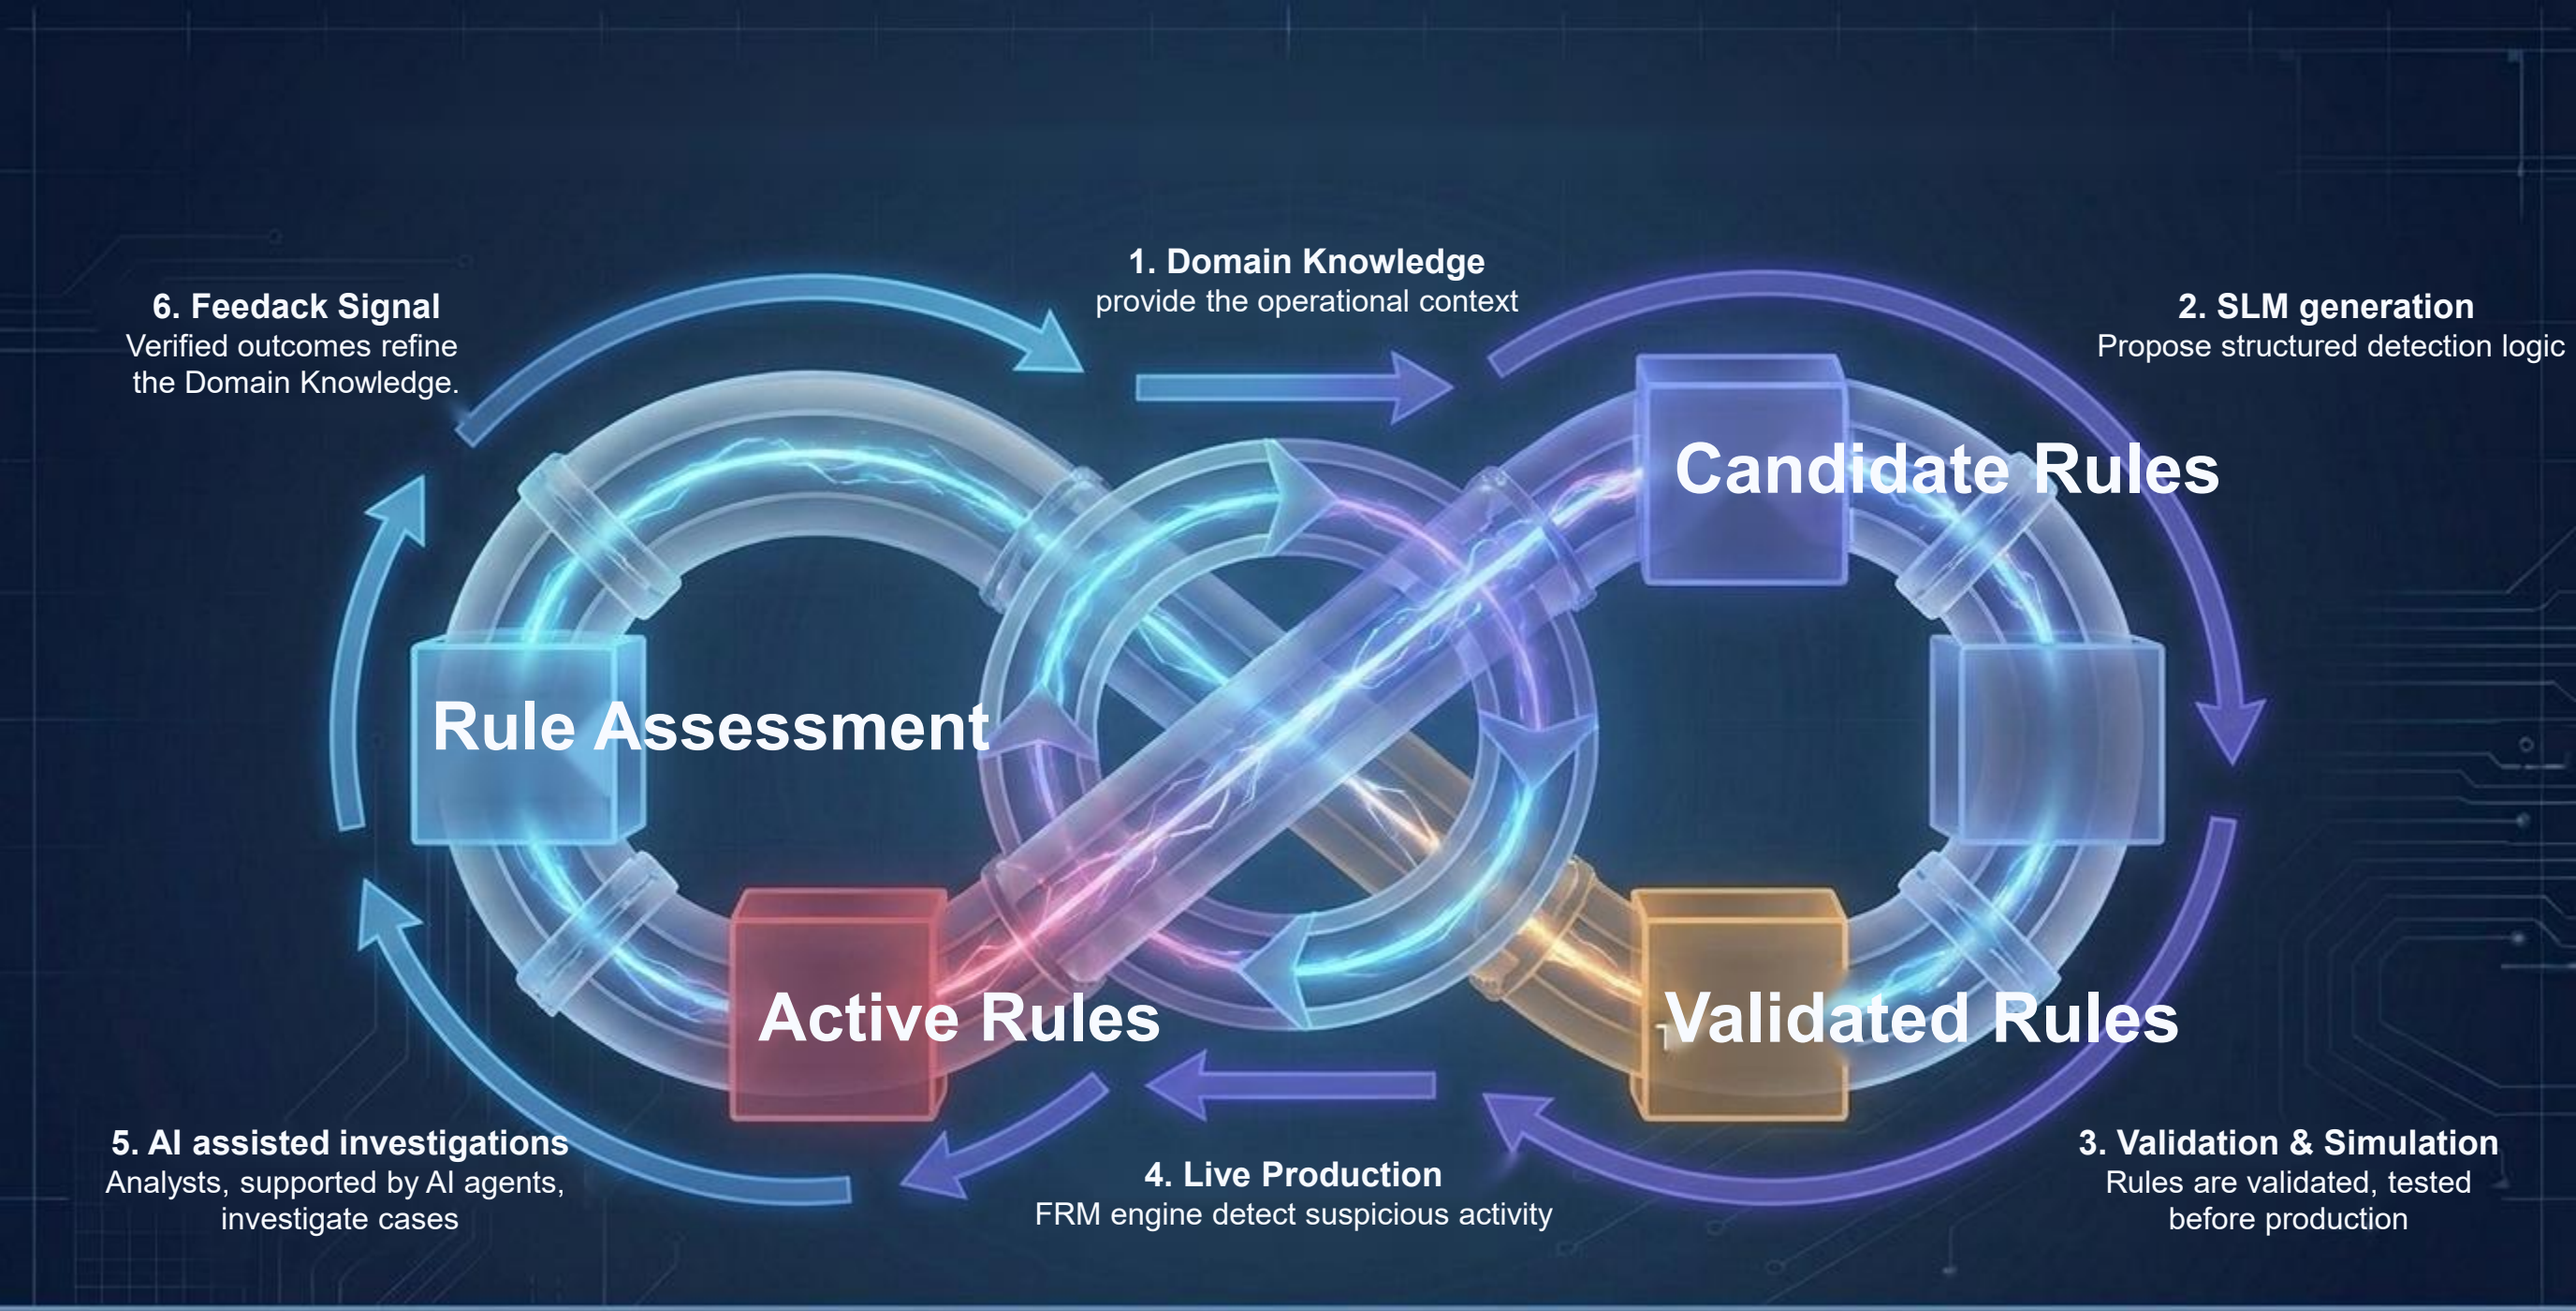

A governed lifecycle that transforms fraud expertise into validated detection logic, executes it deterministically at carrier-grade scale, and improves through verified operational feedback.