QUORUM

CONFIDENTIAL — INSTITUTIONAL DISTRIBUTION ONLY

Document Reference: QRM-BENV-001
Revision: 2.0.0
Issued: May 2026

METHODOLOGY · RESEARCH

Behavioral Signal Synthesis & Multi-Tier Intelligence Architecture

How QUORUM collects, analyzes, and adjudicates behavioral fraud signals through a three-tier LLM consensus engine

CLASSIFICATION	DOCUMENT ID	VERSION
Confidential	QRM-BENV-001	2.0.0
ISSUED	**PAGES**	**DISTRIBUTION**
May 2026	34	Institutional

KB ANALYTICAL SOLUTIONS INC.

SAVE AS PDF

TABLE OF CONTENTS

SECTION 01

The Behavioral Risk Intelligence Problem

The central problem in transaction fraud detection is not data scarcity — financial institutions generate enormous volumes of event data with every interaction. The problem is signal quality and reasoning depth. Transaction metadata alone — amount, merchant category, channel — describes what happened but reveals almost nothing about *who* performed the action, whether they were acting freely, or whether the request was generated by a human or a machine. Behavioral signals answer the question that transaction data cannot: is the entity behind this session a legitimate account holder behaving normally, or is something else in control?

Rule-based fraud systems, which dominated the industry through the 2010s, operate on a fundamentally deficient model. They encode known fraud patterns into explicit logical conditions — "flag if amount exceeds X and merchant category is Y and IP is in country Z." This works precisely once per fraud pattern, against adversaries who have not yet observed the rule. Modern threat actors adapt faster than rule sets can be maintained. The half-life of a new fraud rule against a sophisticated adversarial actor is measured in hours.

QUORUM's approach inverts the problem. Rather than defining fraud by its observable characteristics, the system defines **authentic human behavior** by its measurable properties — typing rhythm, interaction entropy, physical plausibility, and transactional consistency — and treats deviations from that envelope as the primary risk signal. Critically, QUORUM does not rely on a single model or algorithm to make this determination. Three independent intelligence tiers — each specialized for a different fraud surface — evaluate every request simultaneously. Their conclusions are compared, debated, and adjudicated before a final verdict is issued.

> CORE DESIGN PRINCIPLE
>
> QUORUM does not detect fraud by recognizing known fraud patterns. It detects fraud by measuring the degree to which a session fails to exhibit the characteristics of authentic human behavior, then subjects that measurement to structured adversarial scrutiny across three independent intelligence tiers before any action is taken.

QUORUM

Institutional access, complete benchmark data, and full technical implementation details are available to qualified partners and prospects.

REQUEST ACCESS →