WHITE PAPER

JIL SOVEREIGN TECHNOLOGIES

BI-DIRECTIONAL PAYMENT INTEGRITY NETWORK

Adversarial, Validation, and Adjudication Across Verticals

How JIL's AVA layer validates fraud across healthcare, government, and commerce using 110+ global data sources, 234 active attestation checks expanding to 456, Apache Kafka streaming, and proprietary agentic AI. 83 patents filed. CourtChain®. CREB®. Measurable ROI.

DOCUMENT Comprehensive White Paper	PUBLISHED May 2026
AUDIENCE C-Suite, Technical, Institutional	CLASSIFICATION Institutional

CONTENTS

01 The Adversarial Approach

*AVA stands for **Adversarial, Validation, and Adjudication**. The mindset is simple: assume every transaction is fraudulent until proven legitimate. Validate adversarially. Adjudicate decisively. The default is block. Proof is required to pass.*

This is the opposite of reactive fraud detection, where the default is allow and you catch the bad ones after. Adversarial validation reverses the burden. Every claim, every payment, every benefit transaction is questioned. Every signal is examined. The decision is rendered, and it is defended with evidence.

Why adversarial matters

Reactive detection is always too late. A fraudster hits you once, you see it, you block it. By then the damage is done. Adversarial validation stops the fraud before it happens. Every transaction is suspect until the signals prove otherwise. This is why healthcare fraud, benefits fraud, and payment fraud persist at scale: institutions are reactive, not adversarial.

The cost of staying reactive is measurable. Healthcare fraud costs 3-10% of healthcare spending annually. Benefits fraud spikes 500% in crisis periods. Payment fraud absorbs 1-3% of merchant revenue. Government processes billions in transactions with no real-time integrity layer. These gaps persist because the burden of proof sits in the wrong place.

AVA is built on adversarial logic. Every attestation is a question: Is this transaction legitimate? The answer requires evidence. Device reputation. User history. Network behavior. Sanctions status. The decision is not a score the institution can quietly ignore. It is an adjudication that requires action and is defensible in court.

> **ADJUDICATION, NOT SCORING**
>
> Fraud detection returns a score. A reviewer sees a 78 and decides. AVA returns an adjudication: PASS, REVIEW, BLOCK, or ESCALATE. That is the decision. The institution can override it, but the adjudication remains the evidence-backed determination. It is defensible. It is admissible. It is the institution's shield in court.

02 The Wedge: Why JIL Is Different

The fraud market is crowded at the scoring layer and almost empty at the evidence layer. That is JIL's wedge. Legacy vendors tell an institution that something looks risky. JIL tells the institution what to do, preserves why the decision was made, seals the proof, and turns the event into evidence that can be defended, recovered, audited, and prosecuted.

Most fraud platforms live inside operations. They help a reviewer approve, deny, or manually investigate. JIL lives at the control point between operations, legal, compliance, audit, recovery, and law enforcement. That is the strategic difference. JIL is not another model score. It is the evidentiary operating system behind the decision.

> **THE WEDGE IN ONE LINE**
>
> JIL converts fraud detection from a disposable operational signal into a signed, self-authenticating, CourtChain®-anchored evidentiary asset that can stop loss before settlement and support recovery after the fact.

The market gap: scores do not survive challenge

The fraud-detection industry is built around probability. A model looks at a transaction and returns a number. A reviewer interprets the number, makes a judgment, and moves on. That may be useful inside a queue, but it breaks down when the decision is challenged by a customer, provider, merchant, regulator, auditor, opposing counsel, or court.

At that point the institution needs more than a risk score. It needs to prove what was known, when it was known, which signals were used, which model version made the determination, who changed anything, whether the record was altered, and whether the evidence can be authenticated without dragging engineers and data scientists into testimony.

JIL inverts the category. Every adjudication is packaged as a **Court Ready Evidentiary Bundle (CREB®)**: the signed decision, every signal consumed, the model version and weights, the outcome labels, the entity graph, the cryptographic signatures, and a timestamp/content proof anchored to CourtChain®. The CREB® is designed to support admissibility under FRE 902(13) and FRE 902(14), which means the output is not merely useful for operations. It is useful for recovery, litigation, regulatory defense, audit response, and prosecution.

> **902**
>
> FRE RULES TARGETED BY CREB® ADMISSIBILITY: 902(13) COMPUTER-GENERATED RECORDS AND 902(14) BLOCKCHAIN-ANCHORED EVIDENCE

The real wedge: JIL owns the proof layer

The winner in fraud integrity is not the company with the prettiest score. Scores commoditize. Models can be copied, rented, tuned, or replaced. The durable control point is the proof layer: the standard by which institutions memorialize decisions, verify chain of custody, share bad-actor intelligence, defend adverse actions, and recover losses.

That is where JIL sits. JIL does not ask the buyer to trust a black box. It produces a sealed record showing the decision path. That record is usable by operations on day one, compliance during review, finance during recovery, outside counsel during litigation, and investigators during prosecution. The same artifact follows the transaction through its entire life cycle.

BUYER QUESTION	LEGACY ANSWER	JIL ANSWER
Should we stop it?	Here is a score.	Here is an adjudication: PASS, REVIEW, BLOCK, or ESCALATE.

BUYER QUESTION	LEGACY ANSWER	JIL ANSWER
Why did we stop it?	The model flagged it.	Here are the exact signals, weights, thresholds, data sources, and entity relationships.
Can we defend the decision?	Maybe, with expert explanation.	Yes, with a sealed CREB® designed for self-authentication.
Can we recover money?	Usually no. The evidence is fragmented.	Yes. The evidence package is already assembled and portable.
Can we warn the network?	No, each customer remains siloed.	Yes. Bad Actor Repository signals can inform future adjudications.

Why buyers care: the economics move from avoidance to recovery

Traditional fraud tools are budgeted as loss-avoidance software. They reduce exposure but rarely create a recovery asset. JIL changes the budget conversation. A prevented claim, denied benefit, blocked merchant transaction, or flagged grant payment becomes a documented evidentiary record. That enables recoveries, recoupments, chargeback defenses, civil actions, False Claims Act support, audit defense, and cross-merchant intelligence.

This is why the wedge matters commercially. JIL can be sold to the fraud team, but justified by legal, compliance, audit, finance, and recovery. The buyer is not only paying for detection. The buyer is paying for an evidentiary infrastructure layer that turns institutional knowledge into admissible, reusable proof.

Anatomy of a CREB

This is the part competitors cannot match, so it is worth being concrete. JIL does not return a score and stop. A score is one field. The CREB is the entire case file, assembled automatically for every adjudication and sealed so it cannot be altered after the fact.

CREB COURT READY EVIDENCE BUNDLE — one sealed package per adjudication

1. **The adjudication**
   Verdict: PASS / REVIEW / BLOCK / ESCALATE, with the risk basis
2. **Every signal consumed**
   Each check fired and data source queried, with its value and weight
3. **The model**
   Version, dimension weights, and thresholds applied for that vertical
4. **The outcomes**
   Confirmed labels: chargeback, fraud confirmed, recovery, clean
5. **The entities**
   Graph of device, IP, payee, card, claim ID, and benefit account
6. **The proof**
   Ed25519 + Dilithium-III signature and SHA-256 hash of the bundle
7. **The anchor**
   CourtChain block height, timestamp, and Merkle inclusion proof

SEALED & SELF-AUTHENTICATING
Admissible under FRE 902(13) and 902(14). No expert testimony required.

Every adjudication produces one CREB: the decision, the full evidence behind it, and cryptographic proof that both are authentic and unchanged.

In practice those components arrive as a single sealed bundle. The specimen below is illustrative, with sample values drawn from healthcare payment integrity. It is one example: JIL produces a CREB for every adjudication across every line of business, from digital commerce to government benefits to federal grants.

COURT READY EVIDENCE BUNDLE

CREB-HC-2026-0530-7F3A9C21 · SPECIMEN · ILLUSTRATIVE

ADMISSIBLE EVIDENCE

BLOCK **Risk 0.94** · Confidence High · 41 of 234 checks fired · adjudicated in 38 ms

VERTICAL
Healthcare Payment Integrity

MODEL
AVA Healthcare v4.2 · profile HC-2026.05

SUBJECT CLAIM
8842166-03

PROVIDER NPI
1043••••27 (masked)

PAYEE
Meridian Wellness Group LLC

DATA SOURCES QUERIED
63 of 110+

EVIDENCE MANIFEST · TOP WEIGHTED SIGNALS

Beneficiary deceased per SSA DMF; date of service after date of death	+0.21
Phantom provider: license inactive, OIG LEIE exclusion match	+0.19
Duplicate encounter billed across two payers, same date of service	+0.17
Upcoding: E/M level unsupported by documented complexity	+0.14
Velocity: 312 claims from a single NPI within 24 hours	+0.13
Service address classified vacant / non-clinical storefront	+0.10
+ 35 additional checks fired (full manifest contained in bundle)	···

CRYPTOGRAPHIC SEAL

```
bundle.sha256    0x9f3ac8b1d4e7…6f21e
signature        Ed25519 + Dilithium-III (ML-DSA-65)
anchor           CourtChain block #1,482,309 · 2026-05-30T14:22:07Z
merkle.root      0x7c2e44a9b0…03b8
quorum           14 of 20 validators · BFT finality
```

Self-authenticating. Admissible under FRE 902(13) for computer-generated records and FRE 902(14) for blockchain-anchored timestamps. No expert testimony required to enter into evidence.

Specimen shown for the healthcare line of business. The same bundle structure and admissibility apply to every vertical JIL adjudicates.

Because the bundle is signed with post-quantum cryptography and anchored to CourtChain, a CREB pulled six months or six years later can be proven to be the exact decision made on the exact date, on the exact signals, under the exact model. Nothing added. Nothing removed. That is why it is admissible without a JIL engineer taking the stand, and why it travels: the same bundle supports a chargeback dispute, a regulatory response, a False Claims Act filing, or a cross-merchant warning.

Legacy detection versus JIL

DIMENSION	LEGACY FRAUD DETECTION	JIL AVA
Output	A risk score (e.g., 78 of 100)	An adjudication (PASS, REVIEW, BLOCK, ESCALATE) plus a CREB
Posture	Reactive, catches fraud after it happens	Adversarial, blocks before settlement
Defensibility	Internal logic, needs expert testimony	Admissible under FRE 902(13) and 902(14), no expert needed
What you can do with it	Block the transaction, absorb the loss	Recover losses, defend adverse action, build litigation and qui tam cases
Institutional value	Operational decision support	Enterprise proof layer for fraud, legal, compliance, audit, and recovery
Intelligence	Single-tenant, siloed per customer	Cross-merchant Bad Actor Repository
Coverage	One model rebuilt per problem	One hypercube weighted across every vertical
Evidence ledger	Proprietary logs, mutable	CourtChain L1, post-quantum, immutable, evidence-only

Why this is hard to copy

It requires a different architecture. Incumbents are built around scoring engines, case queues, and dashboards. Producing admissible evidence is not a feature that can be bolted onto that stack later. It requires deterministic evidence capture, signed event provenance, model/version lineage, signal manifests, entity graph attribution, post-quantum signatures, immutable timestamping, and a chain-of-custody model from the first millisecond of the transaction.

It requires a different legal posture. Most vendors avoid being part of the courtroom record because it creates scrutiny. JIL is designed for that scrutiny. The CREB® exists so the institution can defend the decision without relying on vague vendor assertions, mutable logs, or after-the-fact reconstruction.

It requires a different network. Single-tenant fraud tools see what one institution sees. JIL's Bad Actor Repository compounds across participants while preserving privacy through hashed identifiers and cryptographic tokenization. The more institutions participate, the stronger the signal becomes. That creates a data advantage that is difficult to recreate from a standing start.

It is fenced by IP and implementation depth. 83 filed patents cover the AVA hypercube, CourtChain consensus and protocol, CREB® generation, post-quantum signing, and cross-merchant federation. The implementation is not a slideware concept: the platform spans 300 production services and approximately 1.5 million lines of code. Competitors would need to rebuild the scoring layer, evidence layer, cryptographic layer, legal-authentication layer, and network layer at the same time.

It aligns incentives differently. Legacy vendors are paid to score transactions. JIL is positioned around prevention, proof, and recovery. That makes the commercial case stronger because the product is tied to dollars prevented, dollars recovered, and institutional risk reduced.

> **THE ONE-SENTENCE WEDGE**
>
> Everyone else tells you something looks like fraud. JIL hands you a signed, timestamped, CourtChain®-anchored CREB® that proves it, defends it, shares the intelligence safely, and turns fraud detection into recoverable, court-ready institutional evidence.

03 AVA: The Architecture

JIL's AVA layer is multi-dimensional. Think of it as a hypercube: a single core platform where every vertical (healthcare claims, digital payments, government benefits, federal grants) shares the same reusable signal dimensions but weights them differently based on its own fraud risk profile. One architecture. Many markets. No rebuild per vertical.

Multi-dimensional design

Rather than building six different fraud platforms, JIL built one platform with independent, reusable dimensions: user risk, device risk, network risk, velocity risk, sanctions risk, payment method risk, and creator risk. Each dimension is computed once and reused everywhere. The same device fingerprint that flags fraud in healthcare flags it in commerce. Each vertical assigns its own weights to the same dimensions based on the way fraud actually behaves in that market.

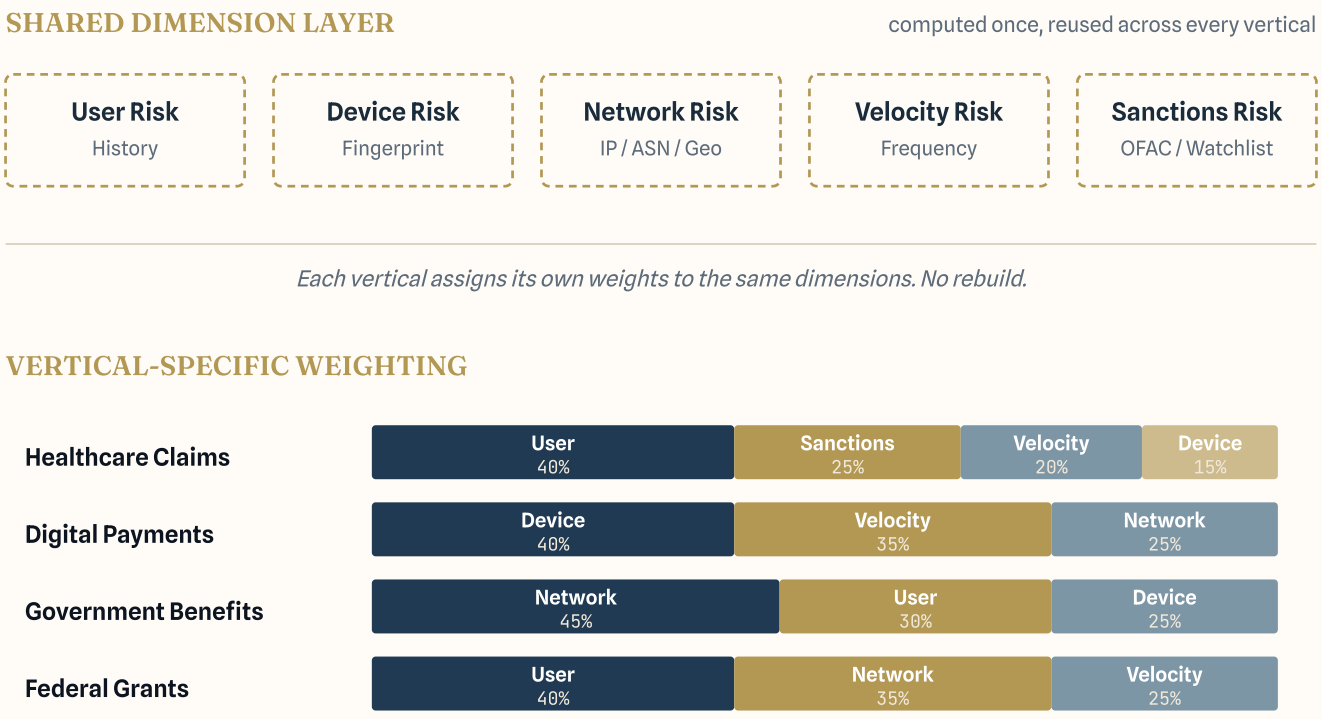

The AVA hypercube: shared dimensions computed once, weighted per vertical. Adding a market is configuration, not a six-month rebuild.

The same six core shared dimensions produce a distinct risk fingerprint for each vertical. Healthcare leans on user history and sanctions. Digital payments lean on device and velocity. Government benefits lean on network and sanctions. One platform, read several different ways.

DIMENSION PROFILE BY VERTICAL

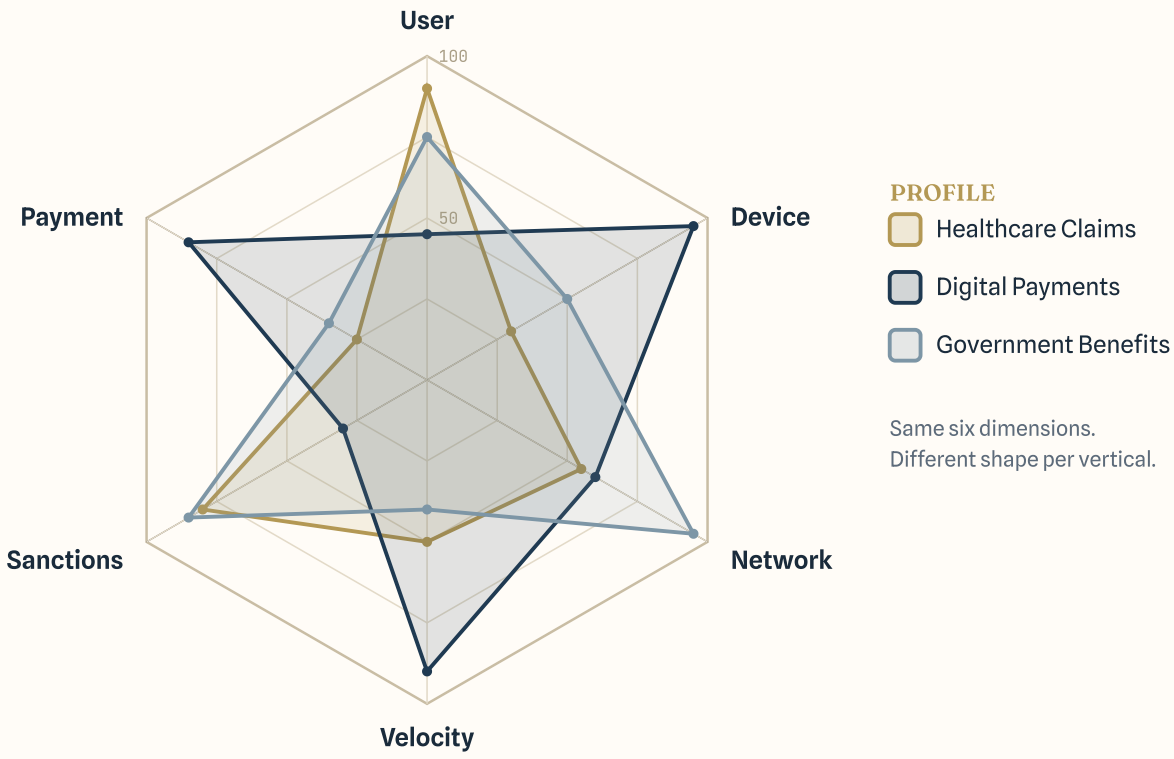

The six core shared dimensions, weighted into a different shape for each line of business. Adding a vertical means drawing a new profile on the same axes, not building new ones.

The four-layer stack

Dimension layer. Independent, reusable signal sources. User, device, network, velocity, sanctions, payment method, and creator risk. Each can be computed once and used across all verticals. This is the foundation everything else is built on.

Weighting layer. Each vertical assigns different weights to the dimensions. Healthcare claims prioritize behavioral and user history because claims fraud follows patterns. Digital payments prioritize device and velocity because card testing and account takeover are the threats. Government benefits prioritize network and sanctions because synthetic identities and mules cross jurisdictions. Same platform, different weights per vertical.

Orchestration layer. Rules, ML models, thresholds. This is where vertical customization lives, but it is built on the shared dimension layer. Adding a new vertical is not six months of feature engineering. It is configuring weights and thresholds on dimensions that already exist.

Evidence layer. Every attestation is signed, timestamped, and anchored to an immutable record. The evidence layer is identical across all verticals because the requirement is identical: prove the decision later, in court if necessary.

WHY THIS SCALES

The hypercube lets JIL add a new vertical, say disaster assistance fraud or pension death fraud, without a complete rebuild. New vertical owners configure weights on existing dimensions, set thresholds, and tune the model. The core

platform does not change. The infrastructure does not change. Only the tuning changes. That is the difference between a product and a platform.

04 Scale and Operational Infrastructure

JIL processes attestations at scale using 110+ integrated data sources across 50+ countries, 234 active checks expanding to 456, Apache Kafka for real-time streaming, and proprietary agentic AI automating operational workflows. Built for millions of events per day.

Data sources: 110+ global integration

JIL ingests signals from 110+ domestic and international sources spanning:

- **Government sanctions and regulatory.** OFAC, EU sanctions, UK OFSI, UN Security Council, world bank debarment, GSA excluded parties lists.
- **Network intelligence.** IP reputation databases, ASN data, Tor exit nodes, VPN/proxy detection, geolocation covering 50+ countries.
- **Device fingerprinting.** Emulation and VM detection, device reputation, browser fingerprinting, biometric signals.
- **Payment and velocity.** Card BIN reputation, chargeback histories, velocity counters, card testing patterns, consortium alerts.
- **Behavioral and transactional.** Claims pattern analysis, billing code validation, eligibility verification, duplicate detection.
- **Cross-merchant intelligence.** JIL proprietary Bad Actor Repository, on-chain transaction attribution for crypto fraud.

> **110+**
> DATA SOURCES INTEGRATED GLOBALLY

Attestation checks: 234 today, 456 roadmap

JIL executes 234 distinct checks across 15+ categories:

- **Healthcare claims (85 checks):** encounter patterns, provider license verification, bundling rules, diagnosis-procedure validity, duplicate detection, phantom provider matching, billing compliance, modality verification.
- **Digital commerce (62 checks):** device fingerprint consistency, card velocity testing patterns, account takeover indicators, refund abuse sequences, promo exploitation, chargeback likelihood, wallet risk assessment, geolocation anomalies.
- **Government benefits (51 checks):** synthetic identity detection, duplicate enrollments across jurisdictions, income verification, employment validation, categorical eligibility rules, mule account patterns, benefits stacking, residency confirmation.
- **Federal grants (36 checks):** contractor sanctions matching, beneficial ownership verification, lending abuse patterns, grant diversion signatures, false certification detection, cost substantiation, timeline inconsistencies.

Roadmap expands to 456 checks across five new verticals: pension death fraud, workers comp, telehealth/DME, pharmacy/340B, digital assets.

> **234**

ACTIVE CHECKS TODAY (456 ROADMAP)

Kafka streaming: real-time architecture

Apache Kafka is the backbone for real-time data ingestion. Raw transaction streams arrive in Kafka topics, are validated and enriched with JIL dimensions in real-time, triggering immediate attestation decisions. Ensures fault tolerance, message ordering, and replay capability. Enables both hot-path decision making (sub-100ms) and cold-path evidence generation.

Agentic AI: autonomous operational efficiency

JIL's proprietary agentic AI layer autonomously manages five critical functions:

- **Outcome analysis and labeling.** When chargebacks, fraud confirmations, or outcomes arrive, the agent automatically extracts, labels, and queues for feedback loop. Processing 1000s per hour automatically.
- **Model retraining orchestration.** Agent triggers retraining when outcomes accumulate or data drift detected. Coordinates feature engineering, training runs, validation. Humans review before production deployment.
- **Anomaly detection and signal optimization.** Agent watches dimension performance across verticals. Flags performance degradation and proposes new data sources. Humans decide based on evidence.
- **Feedback loop orchestration.** Routes confirmed outcomes to Bad Actor Repository, updates reputation scores, triggers cross-merchant intelligence updates. Fully automated.
- **Operational monitoring and intelligent alerting.** Tracks latency, error rates, false positive/negative rates, decision distribution drift, data quality issues. Alerts ops team with root cause analysis.

5

AGENTIC AI AUTOMATION FUNCTIONS

05 Lines of Business: Current and Pipeline

Four lines of business are active today. Five more in advanced roadmap. Each uses the same AVA platform with vertical-specific check weights.

Currently active

Healthcare Payment Integrity (live pilots). MCOs, health plans, hospital networks. Centene partnership active. 10B+ recoverable fraud/year in U.S. HIPAA-compliant native deployment available. Revenue: tiered platform fees (1-7.5% of fraud recovered) plus annual subscription ($500K-$5M).

Digital Commerce Payment Integrity (shadow mode, Phase 1 Q2 2026). Merchants, payment platforms. Epic Games anchor account. Sub-50ms latency. Bi-directional enforcement. Revenue: per-decision pricing ($0.002-0.01) plus platform subscription.

Government Benefits Integrity (active pilots, 2 states). State Medicaid, unemployment, SNAP, disaster assistance. Billions of transactions/month. Revenue: per-state annual contract ($100K-$2M) plus success-based recovery sharing (10-25%).

Federal Grants Integrity (live partnerships). SBA PPP/EIDL recovery, federal contractor fraud, disaster assistance. Retrospective and predictive. CREBs as deliverable. Revenue: success-based (15-25% of recovery) plus qui tam attorney partnerships.

LINE OF BUSINESS	STATUS	PRIMARY CUSTOMERS	ESTIMATED TAM
Healthcare Claims	Live pilots	MCOs, health plans	10B+/year
Digital Commerce	Shadow mode, Phase 1 Q2	Merchants, platforms	2B+/year
Government Benefits	Active pilots (2 states)	State agencies	5B+/year
Federal Grants	Live partnerships	Federal agencies, qui tam	200M+/year
Pension Death Fraud	Roadmap Q3 2026	Pension boards	500M-1B/year
Workers Comp / P&C	Roadmap Q4 2026	Insurers, state WC boards	30B+/year
Telehealth / DME	Roadmap Q1 2027	PBMs, providers	5B-10B/year
Pharmacy / 340B	Roadmap Q2 2027	Pharmacies, PBMs	2B-5B/year
Digital Assets	Roadmap Q3 2027	Exchanges, SROs	1B+/year

Roadmap verticals

The pipeline is not speculative. It is the same hypercube applied to adjacent fraud markets where the dimensions already exist. **Public pension and annuity death fraud** applies death-record temporal anchoring and identity continuity checks to catch continued payments to deceased beneficiaries. **Workers comp and P&C** apply claims pattern and provider-network analysis to insurer and state board fraud. **Telehealth, DME, hospice, and genetic testing** extend the healthcare check library into the highest-abuse billing categories. **Pharmacy and 340B diversion** attack drug-pricing and discount-program abuse. Each launch is configuration on existing dimensions, not a new platform, which is why the roadmap moves in quarters rather than years.

06

Intellectual Property and Differentiation

JIL has filed 83 patents covering AVA architecture, algorithms, and infrastructure. Combined with proprietary blockchain (CourtChain) and admissible evidence systems (CREB), JIL's IP moat is defensible and unique.

Patent portfolio: 83 filed as of May 2026

Patents cover AVA hypercube architecture, multi-dimensional risk scoring with dynamic weighting, agentic AI orchestration, post-quantum cryptographic signing (Dilithium-III, Kyber), CourtChain consensus and protocol, CREB generation and FRE admissibility, cross-merchant intelligence federation, Kafka-based real-time pipelines, outcome feedback loop design, and data residency frameworks. Both defensive and offensive patents spanning all core innovations.

> **83**
> PATENTS FILED

CourtChain®: proprietary Layer 1 blockchain

JIL's Layer 1 blockchain, CourtChain, is built for evidence, not transactions. BFT consensus with 14-of-20 validator quorum. Post-quantum cryptography (Dilithium-III, Kyber KEM, Ed25519) ensures evidence remains valid post-quantum. Hourly block time. Sub-100 byte per attestation footprint. Zero transaction fees (attestations are batched).

Not a payment layer. A proof layer. No bridges, tokens, or DeFi. Just immutable evidence with cryptographic proof of timestamp and content.

> **14/20**
> COURTCHAIN BFT QUORUM

CREB®: Court Ready Evidence Bundle

JIL's CREB format is a forensic package containing: the signed attestation, all signals consumed, the risk model version and weights used, outcome labels, blockchain proof of timestamp and immutability, and chain of custody. Admissible under FRE 902(13) (computer-generated records) and FRE 902(14) (blockchain timestamps). No expert testimony required.

No other fraud vendor produces admissible evidence. JIL does. This matters in litigation, regulatory review, and qui tam cases.

Bad Actor Repository and graph attribution

JIL's cross-merchant Bad Actor Repository is built on proprietary graph store using Neo4j backend tracking entity relationships (devices, IPs, users, wallets, payees, cards, claim IDs, benefit accounts). Relationships weighted by signal strength and corroboration. Observations labeled as observations. Resolutions anchored to CourtChain with timestamps. Graph grows stronger with each merchant. Only JIL has federation, trust, and governance.

07 Deployment: You Choose Your Environment

JIL is infrastructure-agnostic. Same AVA code runs natively in your cloud, on JIL's Databricks cloud, or on JIL's Snowflake cloud. You choose based on compliance, data residency, volume, and existing infrastructure. JIL inherits your system.

Three paths

Native in your cloud. JIL code runs in your AWS/Azure/GCP account. Data stays in your VPC. You own and manage infrastructure. Best for HIPAA/PHI compliance, banks, strict data residency. You handle patching and scaling. JIL manages AVA logic. Pricing: JIL platform fee plus your cloud costs. $50K-$500K annually.

JIL Sovereign Databricks Cloud. Data sent to JIL's managed Databricks workspace. JIL handles ML, streaming, real-time serving. PII hashed at rest. Best for digital commerce, payments, high-volume (1M to 100M+ events/day). Sub-50ms p99 serving. Pricing: per-DBU plus cloud VM. $200K-$5M+ annually at scale.

JIL Sovereign Snowflake Cloud. Data sent to JIL's managed Snowflake account. PII tokenized and encrypted. JIL manages infra. Best for analytics-heavy, SQL-first (100K to 5M events/day). Per-second billing with auto-suspend. Pricing: per-credit. $100K-$2M annually at scale.

> **JIL INHERITS YOUR INFRASTRUCTURE**
>
> You run Databricks? JIL deploys to Databricks. You run Snowflake? JIL uses Snowflake. You run both? We switch based on workload. JIL's architecture is cloud-agnostic because the goal is fitting into your ecosystem, not forcing you into ours.

08 Databricks vs Snowflake Cloud: The Trade-offs

Both are excellent. The choice depends on your workload. JIL uses whichever fits your problem best.

Databricks: 20-40% cheaper than Snowflake for ML, streaming, real-time (can be 2-9x better on large-scale scans). Per DBU ($0.15-0.55) plus cloud VM. Sub-50ms serving at 300K+ QPS. Best for 1M to 100M+ events/day. Requires discipline in tuning. Cost explodes if misconfigured.

Snowflake: Comparable cost for SQL analytics and BI. Includes infrastructure (per credit, $2-4). Per-second auto-suspend saves money on bursty workloads. Not the right tool for real-time serving or intensive ML. Best for 100K to 5M events/day.

DIMENSION	DATABRICKS CLOUD	SNOWFLAKE CLOUD	NATIVE (YOUR CLOUD)
Workload	ML, streaming, real-time	SQL, analytics, BI	Any, you control
Volume	1M to 100M+ events/day	100K to 5M events/day	Any scale
Real-time serving	Yes, sub-50ms	No	Your choice
Cost predictability	Medium (tuning required)	High (auto-suspend)	Your responsibility
Network effects	Yes, full Bad Actor Repository	Yes, full Bad Actor Repository	No, single-tenant only

09 ROI and Business Impact

JIL's ROI is strongest when measured against enterprise-scale payment-integrity exposure, validated recoverable loss, and evidentiary leverage. The platform is priced as institutional infrastructure, not as a narrow fraud-point tool. Shadow mode establishes the validated addressable pool before production scale.

National MCO payment integrity: enterprise-scale ROI

For national Managed Care Organizations, JIL uses an annual flat-fee enterprise model rather than a narrow per-fraud-event pricing model. The target buyer is a national-scale MCO or healthcare payor platform with estimated annual claims exposure that may exceed **$130B**, including organizations such as UnitedHealthcare / UHG, Humana, Centene, and similar national payors.

JIL's annual enterprise platform fee is typically structured in the **$3M-$5M annual range**, depending on claims volume, deployment scope, integrations, and operational footprint. This fee covers the AVA platform, claims-integrity workflows, anomaly detection, validation logic, audit trails, provider-risk monitoring, dashboards, escalation workflows, and enterprise deployment across payment-integrity and SIU functions.

The ROI case is strongest when measured in basis points. In a **$130B annual claims environment**, a $3M-$5M platform fee represents only approximately **0.23-0.38 basis points** of annual claims exposure. At that scale, JIL does not need to recover billions to justify the platform. It only needs to influence a small fraction of claims leakage through prevention, validation, dispute support, recovery support, audit readiness, or defensible escalation.

ANNUAL CLAIMS EXPOSURE	7% GROSS FWA / LEAKAGE EXPOSURE	12% GROSS FWA / LEAKAGE EXPOSURE
$130B	$9.1B	$15.6B
$150B	$10.5B	$18.0B
$200B	$14.0B	$24.0B

MCOs commonly face significant fraud, waste, abuse, and claims-leakage exposure across medical claims, pharmacy claims, provider billing, eligibility, duplicate claims, upcoding, phantom billing, medically unnecessary services, and network-level abuse. Using a 7-12% gross exposure range, a $130B annual claims environment represents approximately **$9.1B-$15.6B** in potential annual fraud, waste, abuse, and payment-integrity exposure. JIL does not assume this full amount is recoverable. Instead, shadow mode establishes the validated addressable pool and shows where the platform can prevent, validate, escalate, support recovery, or strengthen evidentiary posture.

METRIC	ILLUSTRATIVE NATIONAL MCO EXAMPLE
Estimated annual claims exposure	**$130B**
JIL annual enterprise platform fee	**$3M-$5M**
Platform fee as percentage of claims exposure	0.0023%-0.0038%
Platform fee in basis points	**0.23-0.38 bps**
Gross 7-12% FWA / leakage exposure	**$9.1B-$15.6B**
Required influenced value for 3x ROI on $5M fee	$15M
$15M as percentage of $130B claims exposure	0.0115%

METRIC	ILLUSTRATIVE NATIONAL MCO EXAMPLE
$15M in basis points	**1.15 bps**

NATIONAL MCO ROI LOGIC

At $130B+ in annual claims exposure, a $3M-$5M platform fee is justified if JIL influences only a few hundredths of one percent of claims value through prevention, validation, dispute support, recovery support, audit readiness, or defensible escalation.

Escalated matters and CREB® creation

When matters escalate beyond routine platform review, JIL charges separately for **CREB® creation**. A CREB® - Court Ready Evidentiary Bundle - packages validated evidence, claim history, anomaly rationale, supporting records, audit trail, attestation data, and chain-of-custody artifacts needed for dispute resolution, settlement, enforcement referral, regulatory review, or litigation support.

For escalated matters, JIL charges a small negotiated upfront or preparation-based percentage and a small negotiated portion of settlement, recovery, or financial resolution. This allows the MCO to use JIL broadly as payment-integrity infrastructure while reserving additional economics for high-value matters where JIL creates formal evidentiary and recovery leverage.

The result is a disciplined commercial structure: the MCO pays a predictable enterprise platform fee for broad operational use, while JIL participates separately in high-value escalations only when CREB® creation contributes to recovery, settlement, or financial resolution.

Digital commerce: ROI through chargeback prevention

Digital commerce ROI is driven by chargebacks, refund abuse, account takeover, card testing, promo abuse, seller fraud, creator-payout fraud, and cross-merchant bad-actor detection. For this line of business, JIL should be priced through a mix of platform fee, volume-based decisioning, and validated savings. The economics are strongest when the merchant or platform has large transaction volume, high dispute burden, or network-level fraud exposure.

JIL's value is not limited to a model score. The system produces adjudications, audit trails, Bad Actor Repository signals, chargeback evidence, and CREB® packages where disputes escalate into recovery, arbitration, or litigation support. This improves loss prevention while strengthening the customer's ability to defend adverse decisions.

Government benefits: ROI at program scale

Government benefits ROI is driven by synthetic identity, duplicate enrollment, deceased-beneficiary payments, mule networks, eligibility manipulation, cross-program abuse, and improper payments. At state or federal scale, even a modest reduction in improper payments can justify a substantial deployment.

JIL's benefit to public-sector programs is strongest when it validates eligibility, identifies duplicate or synthetic identities before disbursement, preserves audit trails, and creates defensible evidence for recovery, referral, or enforcement. Pricing should reflect program size, procurement rules, deployment scope, and validated preventable or recoverable loss.

Federal grants: ROI through recovery and litigation support

Federal grants should be evaluated as a recovery and evidence product, not merely as SaaS. JIL identifies high-confidence grant-fraud patterns and generates CREB® packages that support investigative triage, civil recovery, agency enforcement, DOJ referral, or litigation support.

For retrospective grant programs, the relevant economic pool is the portfolio of recoverable matters, not only annual transaction volume. A single high-value recovery can justify meaningful JIL economics when the fee is tied to actual recovery, settlement, or formal financial resolution.

THE COMMON PATTERN

Across all verticals, JIL proves the loss first, validates the addressable pool, and prices against institutional value. National MCOs pay a predictable enterprise platform fee for payment-integrity infrastructure, with CREB® economics reserved for escalated matters. Commerce, benefits, and grants use the same principle: platform value plus evidence-backed recovery where JIL materially contributes to prevention, settlement, or financial resolution.

---

10 Blockchain for Evidence

JIL uses blockchain for one purpose: immutable evidence. Not transaction processing. Not real-time. Not in the hot path. Just immutable record keeping that holds up in court.

CourtChain: the evidence ledger

Every AVA attestation is hashed and batched to CourtChain once per hour. Post-quantum cryptography (Dilithium-III, Kyber, Ed25519) ensures evidence remains valid post-quantum. A fraud attestation from six months ago can be retrieved, proven valid, and presented in court. The chain proves the decision was made on a specific date with specific signals. No tampering. No rewriting history. A proof system.

Why immutable records matter

Litigation. JIL's signed attestation anchored to blockchain is admissible under FRE 902(13) and FRE 902(14).

Regulatory review. When OIG audits a hospital, the attestation links to full evidence bundle. Decisions defensible with on-chain proof.

Whistleblower cases. Qui tam attorneys use JIL's attestations to build False Claims Act cases. Evidence is portable and trustworthy.

Cross-merchant network. When device is confirmed fraudulent across three merchants, JIL anchors to chain. Next merchant verifies reputation is real because it is on-chain.

NOT A PAYMENT PROCESSOR

Critical: JIL's blockchain does not move money, settle payments, or tokenize assets. It is read-only. Merchants query it, retrieve evidence, present in court. The chain is never in hot path of any real-time system. Blockchain is good at one thing: immutable record keeping. JIL uses it for exactly that.

11 Cross-Merchant Network Effects

A stolen card hits Hospital A and is flagged. Two weeks later, the same card hits Digital Commerce Platform B independently. A month later, Government Agency C is still processing it. The card is worth 50K to the attacker because each merchant has blind spots.

JIL's Bad Actor Repository changes the math. When Hospital A flags the device, Platform B is warned in real time. The attacker's window collapses from weeks to minutes. Network effect only works if merchants share the same database. That is why cloud deployments enable network effects, and native deployments do not.

Card testers hit a merchant hundreds of times looking for valid cards. One merchant sees velocity and flags it. But if the card is hitting five merchants with low velocity per merchant, none flag it individually. JIL computes velocity across merchant boundaries. Same card being tested at three merchants reveals a pattern none see alone.

Repository uses hashed identifiers and cryptographic tokenization so merchants cannot reverse-engineer each other's customer data. Merchant knows "this device is bad" but does not know who used it at competitor. Privacy by design. Network effect works because signal is strong enough without identity.

12 The Attestation Pipeline

End-to-end attestation takes tens of milliseconds to decide and hours to become evidence. Hot path is computational. Cold path builds proof.

Hot path: tens of milliseconds

Event arrives via Kafka. Gateway validates and checks idempotency. Decision engine pulls features from in-memory store. Applies weighted dimensions. Returns score and adjudication (PASS, REVIEW, BLOCK, ESCALATE). No blockchain. No external calls. Just computation.

Cold path: hours

In parallel, event is logged with all signals. Bad Actor Repository updated. Decision signed and queued for anchoring. At end of block (hourly), batch of attestations is hashed, merkle-treed, and submitted to CourtChain. Later, full evidence trail retrieved and CREB generated on demand.

Feedback loop and continuous learning

When outcome is known (chargeback, fraud confirmed, successful payment), agentic AI routes through feedback loop. Updates Bad Actor Repository. Retrains ML model. Closes loop. System learns and gets sharper with every confirmed outcome.

13 Why This Matters

Today, when fraud is suspected, merchants absorb losses or fight chargebacks with weak evidence. AVA changes that. Every attestation is a signed, timestamped, cryptographically-verifiable proof. It holds up in arbitration. It holds up in court. It is defensible in regulatory review.

Chargeback losses decrease because evidence is strong. Fraud recoveries increase because case is provable. Regulatory settlements are based on proven fraud, not negotiated value. Insurance premiums go down because risk is proven lower. Whistleblowers can build qui tam cases on solid evidence. For institutions, AVA is not a feature. It is a shield.

14 Roadmap and Ecosystem

JIL's vision: fraud attestation as institutional infrastructure. Evidence that is portable. Decisions that are provable. A network that is stronger because every merchant contributes. 234 checks expanding to 456. 4 verticals expanding to 9.

2026 roadmap

- **Q2 2026.** Epic Games payment integrity go-live (Phase 1 creator payouts). Healthcare claims live with Centene. Government benefits live with two state MFCUs. Expand checks from 234 to 280.
- **Q3 2026.** Federal grants vertical go-live. Expand merchant network to 10+ institutional partners. Launch consortium data integrations at scale. Pension death fraud roadmap finalized.
- **Q4 2026.** Launch pension death fraud vertical. Expand checks to 350+. Reach 100M+ attestations per month across all verticals.

2027 and beyond

Workers comp and P&C fraud vertical launch. Telehealth and DME fraud launch. Pharmacy and 340B diversion launch. Digital assets and crypto fraud launch. Total of 456 checks across 25+ categories spanning all nine verticals. Cross-vertical Bad Actor Repository at massive scale. Strategic partnerships with payment networks, cloud providers, compliance platforms. Regulatory recognition from SEC, FINRA, OCC for JIL attestations in regulatory filings and enforcement.

> **THE NORTH STAR**
>
> Five years from now, "JIL attestation" is a meaningful phrase in fraud and compliance conversations. Institutions trust it. Regulators cite it. Courts accept it without challenge. Whistleblowers use it to build cases. Merchants bid for customers knowing JIL is behind fraud controls. That is the outcome JIL is building toward.

Adversarial. Validation. Adjudication.

JIL's AVA layer assumes fraud until proven legitimate. Validates with 234 checks across 110+ global data sources via Kafka. Adjudicates with evidence anchored to CourtChain and delivered as admissible CREB. Agentic AI automates operational workflows. Saves institutions tens of millions annually.

Contact JIL Sovereign to discuss how AVA can strengthen your fraud controls and your evidence in the courts.

Contact us

Email contact@jilsovereign.com
or visit jilsovereign.com/connect