ReadingMinds.AI

The Decision Evidence Layer

Privacy-safe, verifiable signals that gate AI agent actions.

AI agents are advancing quickly. Their failure rate is not driven by lack of capability. It is driven by poor inputs at the moment of decision. This whitepaper defines the Decision Evidence Layer (DEL): a privacy-first system that converts raw voice interactions into verifiable, auditable signals, cryptographically signed, policy-scoped, and queryable by agents before any action is taken. The result is a new control point in the AI stack. Agents no longer act on assumptions. They act on evidence.

> *Agents do not fail because they lack intelligence. They fail because they lack evidence at the moment of action.*

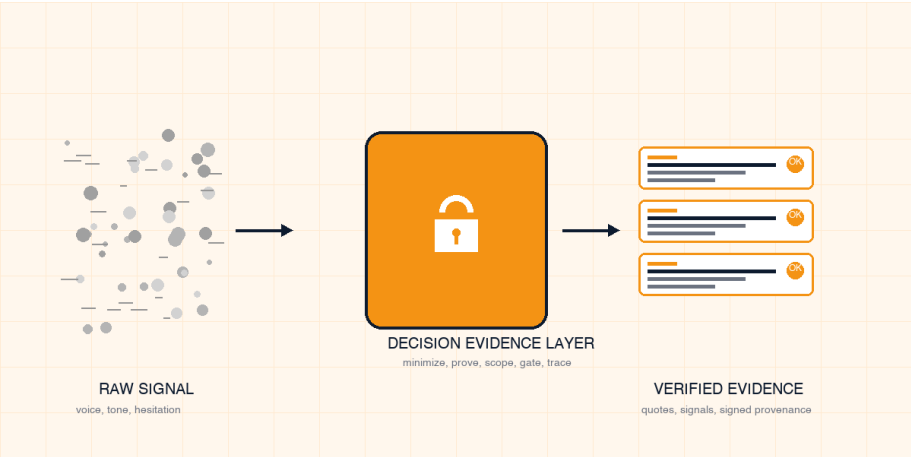

Stu Sjouwerman
Founder & CEO, ReadingMinds.AI
Published April 28, 2026

Table of Contents

> *Real governance is not monitoring agents after the fact. It is requiring evidence before the action.*

#1: Agents Fail With Confidence

Early AI agent deployments reveal a consistent pattern. Agents execute correctly. Workflows trigger as designed. Outcomes are still wrong. The root cause is not execution. It is judgment quality at the point of action.

When an agent makes a decision based on weak inputs, it does not hesitate. It executes. That is where the risk shows up: not in capability, but in the evidence underneath the decision.

Three real-world failure shapes

Marketing agent. Launches a campaign based on survey data that lacks emotional context. Conversion craters and the team cannot explain why.

Customer success agent. Misclassifies churn risk because tone, hesitation, and dissatisfaction were never captured as structured signal. Saves the wrong accounts.

Product agent. Prioritizes features based on volume of mentions instead of intensity of user frustration. Ships polish while the real pain compounds.

The root cause

Most agents act on platform context: CRM fields, email opens, page views, ticket counts. These signals show what people did. They do not show what people felt. Without grounded evidence, agents act in the dark, just faster.

Agents do not fail quietly. They fail with confidence.

#2: Privacy, Procurement, Auditability

For enterprise adoption, three realities must be addressed before any agent is allowed to act on customer data.

Privacy requirements. Regulators and internal policy require strict data minimization and clear purpose limitation. Storing raw recordings expands the threat surface and the legal surface at the same time.

Procurement scrutiny. Security, legal, and compliance teams need defensible artifacts during review, not black-box outputs. They want to see the policy, the retention, the keys, and the audit trail before they sign.

Auditability. Every automated decision must be explainable after the fact. Not just ‘the model said so’ but the specific signals the agent saw, the thresholds applied, and the action taken.

Why traditional analytics and voice tools fail here

They either store raw recordings (which fails minimization and inflates risk) or generate non-verifiable summaries (which fail auditability). Neither posture clears procurement at a regulated enterprise. The Decision Evidence Layer is designed for both at the same time.

Procurement does not block insight. It blocks insight that cannot be defended.

#3: Principles of the Decision Evidence Layer

The DEL is built on five principles that compose into a single discipline. Drop any one of them and the layer stops being defensible.

1. Minimize. No recordings are stored by default. Only derived signals persist. The audio leaves no permanent footprint.

2. Prove. Every insight is backed by an exact quote with start and end timestamps. No claim without a citation.

3. Scope. Access to signals is governed by policy, purpose, and retention limits. The right team sees the right signal for the right window.

4. Gate. Agents must query evidence before taking action. The query is not optional and not skippable in an automated path.

5. Trace. Every decision can be reconstructed end-to-end: the evidence retrieved, the thresholds applied, the resulting action.

#4: How the Decision Evidence Layer Works

The DEL is not another analytics tool. It is an infrastructure layer that sits between raw interaction data and agent execution. Five stages, in order.

1. Capture (transient, no permanent records maintained). Voice interactions are processed in memory. Speech is converted to text. Paralinguistic features (tone, pace, pausing) are extracted. No permanent recordings stored.

2. Derive (structured evidence). Each interaction is converted into a compact, standardized schema, designed to be queried by downstream systems instead of summarized.

3. Harden (integrity and security). Each evidence bundle is secured by hash chains linking quotes to themes, cryptographic signatures, and policy metadata describing access and retention. Tamper-evident, audit-ready.

4. Expose (controlled agent interface). Agents access the layer through a single function such as **getEvidence(query, context)** that returns only the minimum required signals, never raw media, with explicit provenance and confidence.

5. Audit (end-to-end traceability). Every decision an agent makes is logged with the evidence it used, the confidence thresholds applied, and the resulting action. Defensible to security, compliance, and regulators.

The evidence schema

Every interaction collapses into the same compact, queryable structure. This is what makes the layer composable across CRM, marketing, customer success, and product workflows.

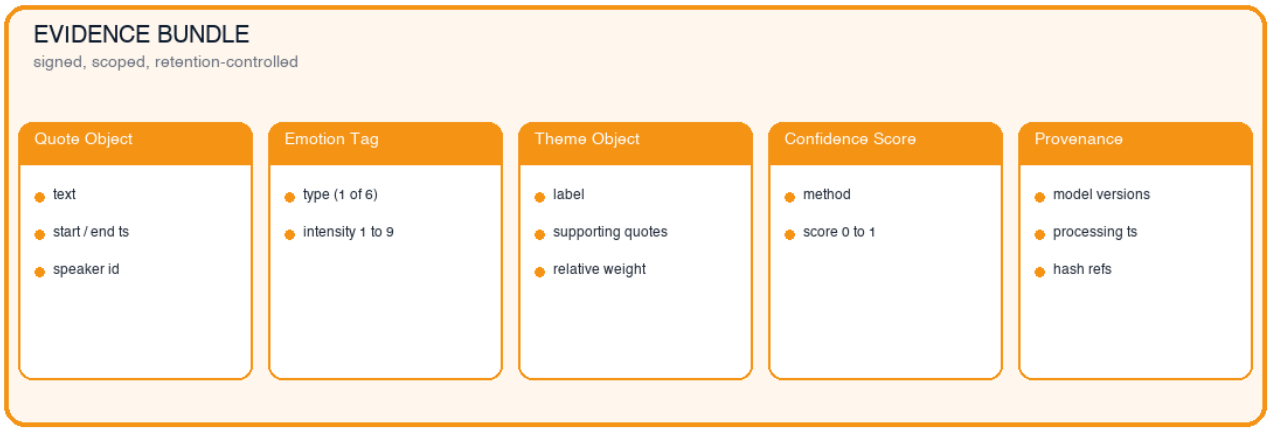

What an evidence bundle looks like in practice

Each agent query returns a compact bundle with quotes, emotion tags, themes, confidence, and provenance. Procurement teams receive the same shape on day one as part of the Evidence Pack, so security review can begin immediately.

EXAMPLE EVIDENCE BUNDLE (REDACTED)


```
{
 "bundle_id": "evb_8c1f...a902",
 "policy_scope": "marketing.campaign_launch",
 "retention_ttl_days": 30,
 "quotes": [
  { "text": "...", "start": 12.4, "end": 18.7, "speaker": "r1" }
 ],
 "emotion": { "type": "confrontational", "intensity": 7 },
 "theme": { "label": "pricing friction", "weight": 0.62 },
 "confidence": { "method": "ensemble", "score": 0.83 },
 "provenance": { "models": ["asr.v3", "emo.v2"], "hash": "sha256:..." },
 "signature": "ed25519:..."
}
```


Precise, verifiable signals beat vague summaries. Every time.

#5: Privacy by Design and Security Posture

The Decision Evidence Layer is built to clear privacy and security review on day one. Every design choice is driven by what regulated enterprise teams actually require, not by what is convenient for the vendor.

Aligned to core privacy principles

Data minimization. Only derived signals persist. No permanent recordings stored, by default and by design.

Purpose limitation. Access is scoped by policy and use case, enforced at query time, not at review time.

Retention controls. Time-to-live policies are attached to each evidence bundle and honored automatically.

Consent traceability. Consent state is part of the evidence metadata, not a separate spreadsheet.

Security posture for enterprise review

Encryption. At rest and in transit, with documented algorithms and key sizes.

Access control. Role-based, least privilege, with policy enforced at the evidence query.

Key management. Custody and rotation processes that survive a key compromise without losing audit trail.

Tamper-evident logs. Hash-chained decision logs that detect retroactive edits and preserve forensic value.

> **Outcome**
>
> The artifacts produced by the DEL are designed to drop directly into a security review. No translation layer between the platform and the people who have to sign.

How DEL maps to the questions security teams actually ask

“Where is the raw audio?” There is none to find. Audio is processed transiently in memory and never stored as a persistent artifact.

“Who can see what, and for how long?” Policy scope and TTL are attached to every evidence bundle and enforced at the query, not by review committee.

“Can you prove this decision?” Every automated action carries a tamper-evident trail back to the exact quotes, signals, and thresholds that justified it.

“What happens if a key is compromised?” Key rotation and custody procedures preserve the audit trail without requiring a full re-derivation of historical evidence.

> *Privacy and security are not after-the-fact paperwork. They are the design constraints every signal in the layer is built around.*

#6: A Procurement-Ready Artifact Bundle

To accelerate procurement, every deployment produces a standard Evidence Pack. These are the artifacts security, legal, and RevOps teams ask for in week one. Shipping them by default compresses the review cycle from quarters to weeks.

Sample evidence bundle (JSON). A real, redacted bundle showing the schema, a quote, an emotion tag, and the confidence score, end to end.

Policy scope and consent metadata. Who can query, for what purpose, against which respondents, with consent state attached.

Retention configuration. Per-bundle TTL, deletion guarantees, and the audit trail proving deletion.

Model and processing manifest. Versions of every model in the pipeline, with hash references for reproducibility.

Example audit trace. A real decision walked back from action to confidence threshold to source quotes.

Why this matters

Most enterprise reviews stall not because the technology is risky, but because the vendor cannot produce defensible artifacts on demand. The Evidence Pack is the answer to that stall, packaged on day one.

If the artifact does not exist by default, it does not exist when procurement asks.

#7: Outcomes and KPIs

Organizations that adopt a Decision Evidence Layer should expect measurable change inside 30 days. Track these metrics from day one so the value is visible to Finance, not just to the team that implemented it.

What changes in 30 days

Faster insight cycles. Customer-conversation-to-decision time compresses from weeks to days.

Reduced incorrect automated actions. Agents stop firing on assumptions and start gating on evidence.

Higher confidence in campaign and product decisions. Internal debate moves from opinion to citation.

Earlier churn and dissatisfaction detection. Tone, hesitation, and intensity surface before activity drops.

Key Performance Indicators

KPI	What It Measures	Why It Matters
Time to actionable insight	Hours from interview to a cited, decision-ready signal.	Compressing this metric is the single best leading indicator of agent ROI.
Percent of agent actions gated by evidence	Share of automated decisions that called **getEvidence()** first.	Should approach 100 percent for any action that materially affects a customer.
Audit pass rate	Share of automated decisions that can be reconstructed end-to-end on demand.	The metric security and compliance will ask for first. Make it boring on purpose.
Campaign and retention lift	Conversion, retention, and pipeline change vs. pre-DEL baseline.	Where evidence-gated decisions earn their seat at the budget table.

> *If you cannot show the percentage of agent actions gated by evidence, you do not yet have a Decision Evidence Layer. You have a label.*

#8: The DEL Adoption Checklist

Run every line before you authorize an agent to act on customer signal. If you cannot check all twelve, you are deploying agents on assumptions.

- [ ] No permanent recordings stored, by default and by design.
- [ ] Every insight is backed by an exact quote with start and end timestamps.
- [ ] Each evidence bundle carries policy scope and a documented retention TTL.
- [ ] Evidence is hash-chained and signed, with tamper-evident decision logs.
- [ ] Agents call a **getEvidence()** interface before any customer-affecting action.
- [ ] The interface returns the minimum required signals, never raw media.
- [ ] Confidence thresholds for each gated action are documented and version-controlled.
- [ ] Every automated decision is reconstructable end-to-end (evidence, threshold, action).
- [ ] Consent state is part of the evidence metadata, not a separate system.
- [ ] An Evidence Pack (sample bundle, policy, retention, manifest, audit trace) ships by default.
- [ ] The six emotions (sad, angry, confrontational, neutral, cheerful, enthusiastic) are scored 1 to 9 per turn.
- [ ] KPIs are published: time-to-insight, percent gated, audit pass rate, lift vs. baseline.

#9: The Upshot

AI agents will not fail because they lack intelligence. They will fail because they lack evidence at the moment of action. The Decision Evidence Layer introduces a new operating standard for the agent era.

The new standard

No decisions without proof. No automation without traceability. No insights without provenance. Every agent that acts on a customer first proves it is allowed to.

This is not another analytics layer. It is a control layer for the agent era. The organizations that install it now will compound an audit-defensible, decision-quality advantage that compounds for years. The ones that wait will spend the same years explaining incorrect automated actions to customers, regulators, and their own boards.

No decisions without proof. No automation without traceability. No insights without provenance.

Start With A Pilot

Conduct 10 to 20 customer interviews. Generate a full Decision Evidence bundle. Trace a real decision from action back to source quote. In less than 30 days you will see the difference between insight-driven automation and evidence-driven execution.

#10: Resources & About ReadingMinds

Two ways to put this whitepaper to work this week.

Live Test Drive
Talk to Emma for 3 minutes about your biggest customer feedback challenge. See the ReadingMinds Expression Fingerprint™, themes, and cited quotes she delivers from a real conversation. No permanent recordings stored.
readingminds.ai/live-test-drive

Procurement & Trust
Review the trust architecture, data retention policy, DPA, incident response plan, and subprocessor list. Built for the security review on day one.
readingminds.ai/trust-compliance

ABOUT READINGMINDS

ReadingMinds is the AI-native customer insight platform built for teams that need to understand what customers actually express, not just what they say. Emma, our AI voice interviewer, runs natural short conversations and captures the ReadingMinds Expression Fingerprint™ of every respondent across six emotions and intensity 1 to 9. Decision-ready customer truth in hours, not weeks. No permanent recordings stored.

Source notes and related reading

- ReadingMinds emotion taxonomy: six emotions (sad, angry, confrontational, neutral, cheerful, enthusiastic) scored 1 to 9 in intensity per conversational turn.
- Companion blog post: **Agents Don't Fail Quietly. They Fail Confidently.** at readingminds.ai/blog.
- Trust architecture, data retention, DPA, incident response, and subprocessor list at readingminds.ai/trust-compliance.