

With you every step of the way.

Latent-Sense Technologies

ENTERPRISE AI DEFENSIBILITY

📍 112 -970 Burrard Street,
Unit 1330, Vancouver,
BC V6Z 2R4
Canada.

📍 Suite 111,
1201 North Market Street,
Wilmington, DE 19801,
United States.

🌐 www.latentsense.ai

✉ inquiries@latentsense.ai

© 2026 Latent-Sense Technologies Inc.

ENTERPRISE AI DEFENSIBILITY

LST provides defensibility, and operational control over AI reasoning in your organization.


40-80%
AI incident &
remediation costs



60-90%
audit prep & regulatory
response time



De-risking
Enterprise AI



AI Defensibility Solutions

AI Risks Analysis Fast-Track

A 2-4 week rapid-response engagement designed for executives who need an immediate, credible view of their AI risk exposure.

Reasoning Stress Test

A forensic evaluation of your AI workflows that deliberately probes for hallucinations, hidden contradictions, broken causal chains, and audit-breaking failures under adversarial and edge-case conditions.

Evidence-First AI Retrofit

A surgical retrofit that injects evidence capture, causal tracing, and contradiction detection into legacy AI systems without requiring a full rebuild.

Agentic Workflow Forensics

A forensic close examination of multi-agent or autonomous workflows to reveal hidden coordination failures, compounding hallucinations, and delegation errors across agent swarms.

AI Governance Readiness Report

A practical readiness assessment that stress-tests your current policies, controls, and decision rights against emerging AI regulations and internal accountability standards.

Cognitive Supply-Chain Risk Audit

An audit of all third-party models, APIs, datasets, and AI services embedded in your workflows to surface hidden dependencies and systemic risk.

AI Reasoning Benchmarking

An objective, model-agnostic scoring framework that measures reasoning quality across LLMs, agents, and composite AI systems.

Audit-Ready AI Pack

A rapid transformation kit that makes your existing AI systems audit-survivable by adding evidence capture, decision provenance, contradiction detection, and accountability mappings.

Hyperscaler Gap Analysis

A deep comparative analysis of what AWS / Azure / GCP “native AI stacks” provide versus what high-stakes, regulated AI actually requires for auditability, determinism, and accountability.

AI Procurement Certification

A formal certification process that evaluates vendors’ AI products against LST’s reasoning, auditability, and governance benchmarks. This gives enterprises a defensible, third-party basis for procurement decisions.

Cognitive Infrastructure Assessment

A board-grade diagnostic that evaluates whether your organization’s technical stack, governance model, data flows, and decision processes are structurally capable of supporting high-stakes AI.

M&A AI Due-Diligence Pack

A pre-acquisition audit of a target company’s AI systems to uncover hidden technical debt, regulatory risk, and reasoning fragility that traditional due diligence misses.

112 -970 Burrard Street,
Unit 1330, Vancouver,
BC V6Z 2R4
Canada.

Suite 111,
1201 North Market Street,
Wilmington, DE 19801,
United States.


COGNITIVE AI
Latent-Sense Technologies

WHAT SETS LATENT-SENSE APART

Every LST AI Defensibility solution is part of our proprietary Cognitive AI infrastructure, a market defining category. Currently, no other vendor can attain the same level of regulatory survivability, defensibility, and operational control over AI reasoning across multi-model, multi-agent, and multi-vendor stacks.

Evidence-First Architecture

Traditional firms bolt logging onto opaque systems and call it governance. LST flips the paradigm. We make evidence generation a first-class computational primitive with forensic-grade AI evidence.

- ✓ regulator-grade audit artifacts
- ✓ rebuild causal chains across AI workflows
- ✓ detect hallucinations as logical failures, not text errors
- ✓ bind decisions to accountable actors
- ✓ reasoning stress tests
- ✓ regulatory simulation labs
- ✓ agentic workflow forensics
- ✓ AI liability mapping
- ✓ cognitive supply chain audits.

Where We Fit in the AI Stack

Foundation Model Providers (OpenAI, Anthropic, Google, etc.)

LST's rxMaps and cognitive infrastructure integrate seamlessly with leading LLMs, enhancing traceability, structured memory, and decision transparency across workflows. Our technology is model-agnostic, designed to orchestrate, ground, and strengthen LLM outputs within enterprise environments, without replacing or competing with model providers.

Cloud & Infrastructure Providers (AWS, Azure, GCP)

LST operates as an interoperable cognitive layer within cloud ecosystems, enriching enterprise data with structured semantic graphs (rxMaps) enabling reasoning-aware workflows on top of existing cloud primitives. We are cloud-agnostic by design, deployable within any major cloud environment to strengthen governance, traceability, and reasoning observability.