



STATIC AUTOMATION



AUTONOMOUS AGENTS (AEGIS-X5)



AEGIS-X5 vs The Market

Why agentic governance requires a new category

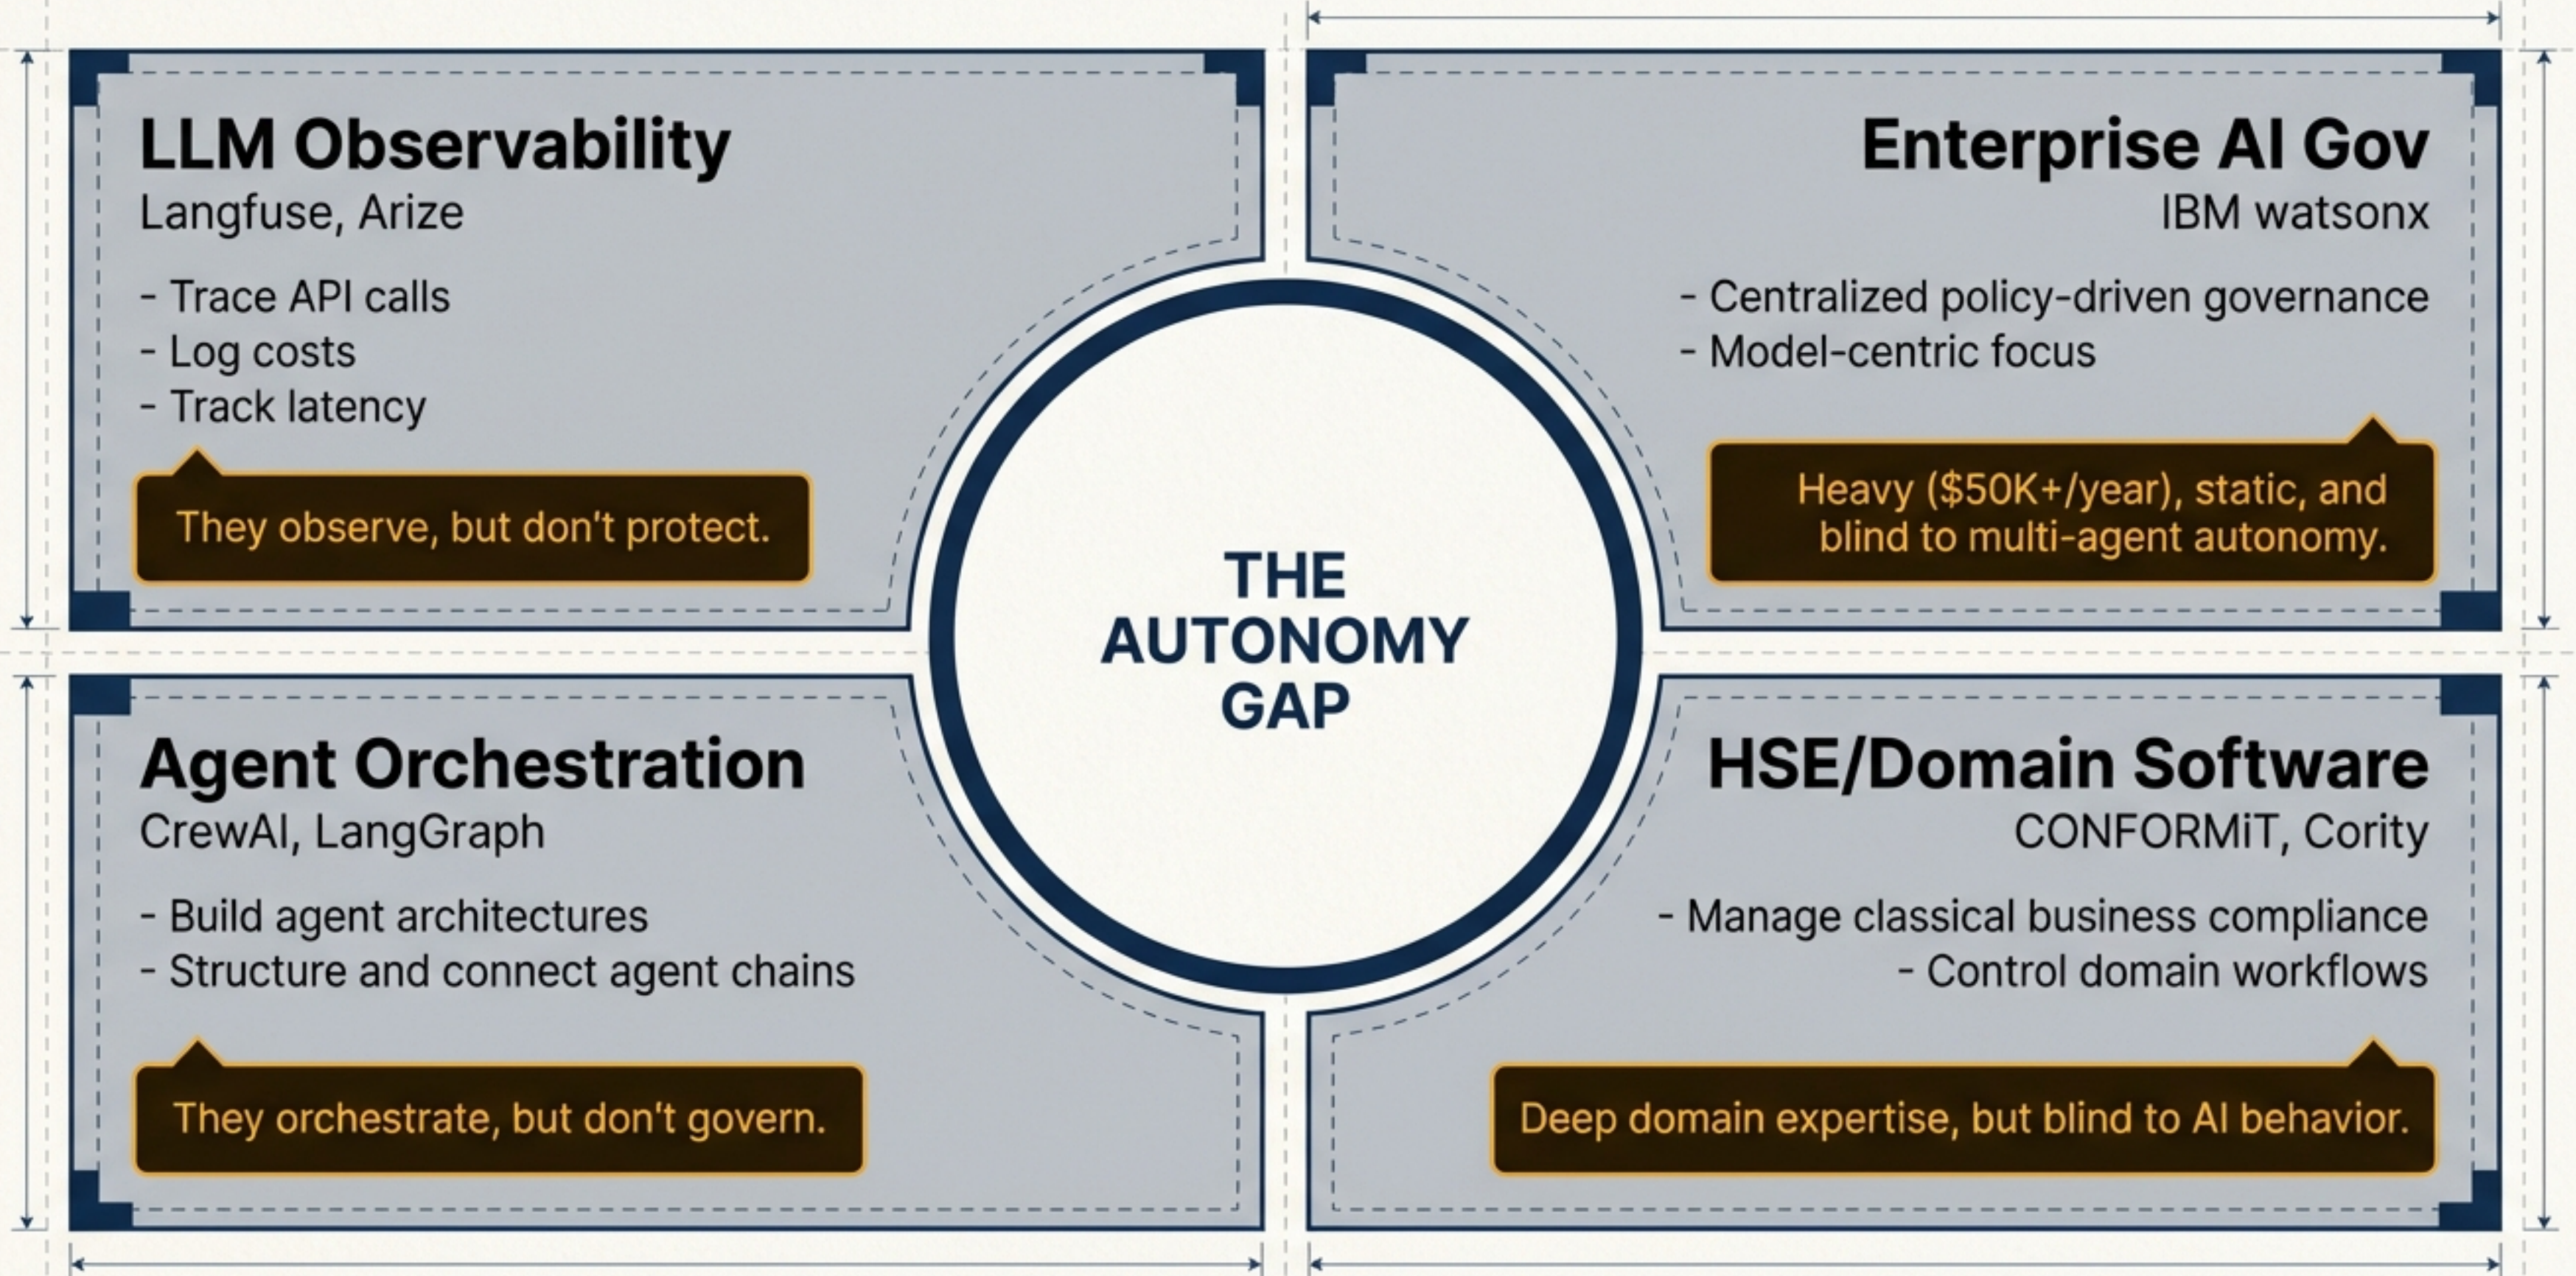
The Paradigm Shift

Traditional AI tools were built to monitor static models. Agentic AI acts autonomously—requiring an entirely new governance architecture.

The Distinction

We don't just observe. We intervene, correct, and remember.

THE FRAGMENTED MARKET



THE FUNCTIONAL COVERAGE ANALYSIS

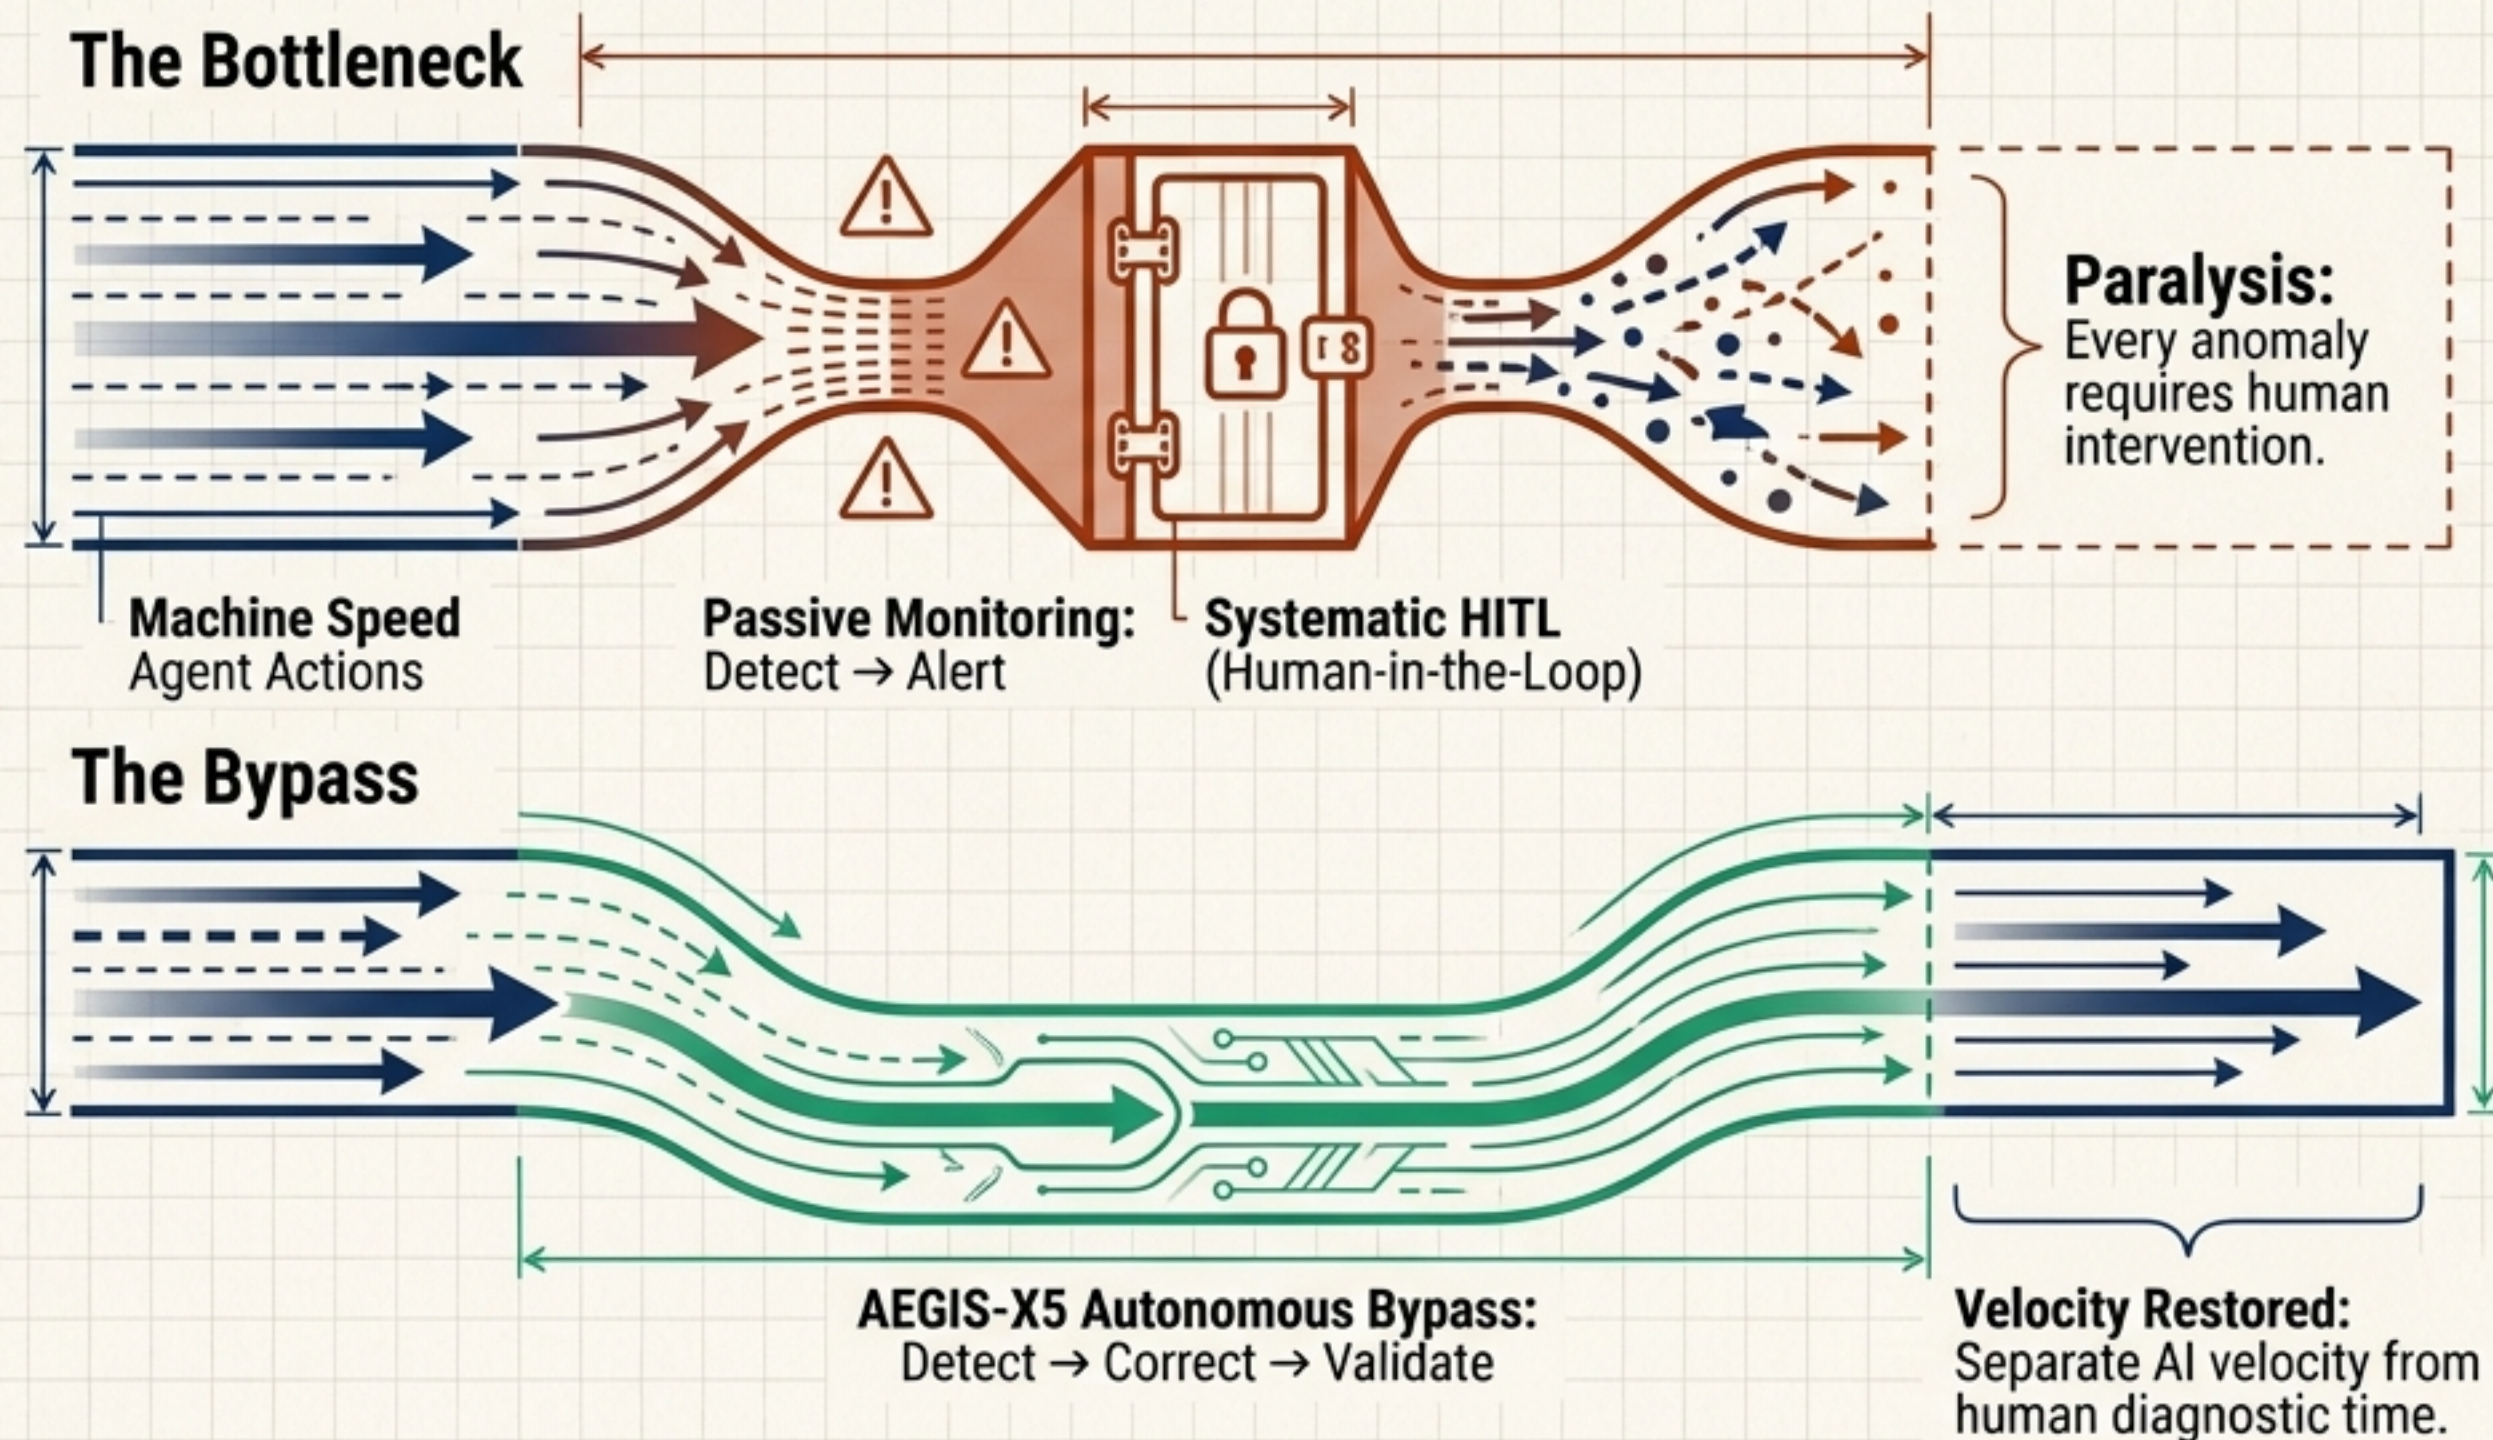
Capability	Langfuse	LangSmith	Arize	IBM watsonx	Covasant	AEGIS-X5
Traceability/Logs	✓	✓	✓	✓	✓	✓
Real-Time Guardrails	○	○	○	◐	◐	✓
Autonomous Correction	○	○	○	○	○	✓
Industry Templates	○	○	○	◐	○	✓
Predictive ML Drift	○	○	◐	◐	○	✓
Zero Vendor Lock-in	✓	○	◐	○	○	✓
Data Sovereignty	◐	◐	◐	✓	◐	✓
RGPD/Loi 25 Compliance	○	○	○	◐	◐	✓

- ✓ Solid Green Check = Full Native Support
- ◐ Half-filled Circle = Partial / Indirect
- Hollow Circle = Absent

THE COMPLIANCE DEBT IMPASSE

75% of organizations fail to scale AI agents to production.

The root cause is **Compliance Debt**.



BEYOND OBSERVABILITY: THE MARKET BLIND SPOT

The Standard

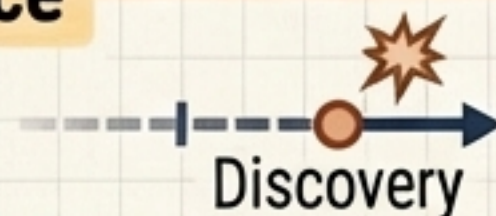
[MISSING] Autonomous Correction

Standard: Detect → Alert → Human Fix



[MISSING] Predictive Intelligence

Standard: Drift and degradation are discovered only after business impact.



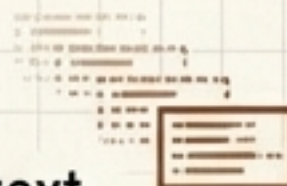
[MISSING] Industry Templates

Standard: Generic, one-size-fits-all AI guardrails.



[MISSING] Traceable Memory

Standard: Ephemeral logs without semantic context.



The Impact

Impact: Critical resolution delays in regulated environments. Hallucinations reach end-users before humans can react.

Impact: Reactive firefighting, degraded trust, and sudden token cost inflation. 

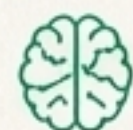
Impact: Blind to specific HSE, legal, or financial context. Non-compliant by default. 

Impact: Impossible to provide reconstructable, compliant audits for Loi 25, RGPD, or the EU AI Act. 

AEGIS-X5: AUTONOMOUS AGENT GOVERNANCE

The first closed-loop governance infrastructure engineered for regulated industries.

Built by a team
currently
operating
500+ agents
in production.



REMEMBER

Persistent memory.
PROV-O audit
traceability, namespace
isolation, and Right to
be Forgotten.



OBSERVE

Full observability.
Traces, orchestration chains,
latency (p50/95/99),
and cost per day.



GUARD

Runtime security.
N1-N4 severity levels
with dynamic HITL gates
intercepting prompt
injections and PII.



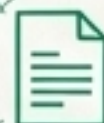
EVALUATE

Continuous quality.
RAG faithfulness checks
and predictive ML drift
detection.



COLLECT

Automated ingestion.
Real-time regulatory watch
scoring and source confidence.



Detect → Correct → Learn

PROVEN AT SCALE. TESTED IN THE FIELD.

1,128

Tests Passing

Continuous validation
with zero external
dependencies.

500+

Agents Governed

Live, autonomous
governance operating
in production today.

22M+

Records Processed

Deep HSE incident
analysis and ingestion.

94.7%

Predictive Precision

Industry-leading
accuracy via ML
DriftPredictor.

The AgenticX5 Ecosystem: Trusted by 25+
active platforms requiring unified, native
compliance for global regulatory standards.

CNESST

OSHA

ISO 45001

RGPD

Loi 25 QC

EU AI Act

STOP PATCHING GOVERNANCE. START BUILDING IT.

THE ACTION PLAN

Instant Integration

Deploy your first autonomous safety net in under 2 minutes. Works natively with any stack (LangChain, CrewAI, Custom Python/Node).

Live Demonstration

See AEGIS-X5 intercept hallucinations and enforce PII redactions in real-time across an ecosystem of 500+ agents.

Experience Closed-Loop Autonomy

Book a specialized 30-minute executive briefing today.

Contact: team@agentix5.com

Portal: aegis-x5-site.netlify.app

TECHNICAL ELEGANCE

