



Verification Architecture for Autonomous AI



Mar 18, 2026



Straive is a Data and AI Operationalization Specialist

Size and Scale


20K+
Employees

4K+
Masters + PhDs


350+
Clients

30+
Markets

Capabilities



Differentiation



Straive Team



**Ankor
Rai**

**President and
CEO**



**Anand
Subramanian**

Head of R&D
(and LLM Psychologist!)



**Manish
Suthar**

**SVP and
Client Partner**




Key Takeaways from Today's Session

- ☑ **AI without robust verification cannot be operationalized.**

Unchecked outputs create operational errors, regulatory exposure, and reputational risk

- ☑ **Because LLMs are probabilistic, enterprises need a layered verification architecture to operationalize autonomous AI. Verification is not optional; it is a critical layer of the Agentic AI stack.**



- ☑ **Discuss 8 relevant use cases & their verification architecture; demo 3 use cases.**

What Happens When AI Solutions Do Not Have Robust Verification



Healthcare / AI Safety

Generated cancer treatment recommendations that were described as 'unsafe and incorrect' by oncologists



Real Estate / Model Drift

ML model systematically overvalued homes, could not adapt to market reversal; \$304M write-down, division shutdown



Insurance / Claims

nH Predict algorithm denied post-acute care to elderly patients, overriding physicians; 90% of denials reversed on appeal



Operational / Hallucination

AI drive-thru ordering abandoned across 100+ locations after persistent order-processing failures and customer complaints



Regulatory / Compliance

Municipal AI chatbot advised business owners to violate city health codes, labor laws, and tenant protections




Audit / Professional Services

A welfare compliance IT system review report contained AI-fabricated court quotes and non-existent academic citations

And... It Can Get Worse !

Global Generative AI Market (\$Bn)



While 78%

enterprises use AI in at least one function...

Only 14%

have enterprise-level governance & verification...

And 74%

In-production AI agents still rely upon manual verification

Source: Bloomberg Intelligence, GenAI Market Forecast (42% CAGR, 2022–2032); Statista (2024 confirmed)

State of AI Deployment and Verification Today

How Enterprises Verify Agents In Production Today


Industry-wide Observations

	AI Use Case	Verification Strategy
Morgan Stanley	AI @ Morgan Stanley Assistant	Retrieve Only from Curated Library of ~100,000 Internal Research Documents
J.P.Morgan	AML Surveillance, Fraud Detection, Coach AI	Cross-validation and Retrieval from Internal Research
BlackRock	Analyze Portfolio Risk & Market Exposure	Multi-model Validation and Stress Testing
HSBC	AML Transaction Monitoring	Anomaly Detection and Human Investigator
UBS	Internal Research Knowledge Access	Retrieval From Repository and Human Expert


AI End-user



Allowed Autonomous Steps




Benefits from current Agentic AIs



Source: ‘Measuring Agents in Production’ (Dec 2025) - Stanford/UC Berkley/UIUC

Strategies to Verify Agentic AI outputs



Source: ‘Measuring Agents in Production’ (Dec 2025) - Stanford/UC Berkley/UIUC

The Future of Verifiable Autonomy

**Moving Beyond
“Testing
Correctness” to
Standard
Software CI/CD
Testing**

*Reduction in 74%
human-in-loop for
verification*

**Better
Understanding
and Definition of
“Agent Failure
Modes”**


**Production and
Verification
Integrated Into
Same Agent
Sub-tasks**

**Agents Become
Software-facing
(vs.
Human-facing)**

*Agents interfacing with
machines increases
from current 7%*

**Verification
Becomes a
Standard Layer
in the Agentic AI
Stack**

Enterprise AI Stack with Verifiable Autonomy



Verification in Action

Use Case List
Not Exhaustive

★

Claims Adjudication

Triage, Assessment, and Decisioning

AI evaluates claims against policy terms, flags anomalies, and recommends approval or escalation

Domain Constraints

Automated Validation

Runtime Oversight

★

Sales Intelligence

Prospect Research & Advisor Enablement

AI synthesizes market data, CRM history, and public filings into advisor-ready briefings

Knowledge Grounding

Automated Validation

Human Oversight

★

Private Credit

Deal Due Diligence & Monitoring

AI analyzes borrower financials, covenant compliance, and risk signals across the credit lifecycle

Knowledge Grounding

Domain Constraints

Automated Validation

Regulatory Intelligence

Regulatory Tracking & Impact Assessment

AI monitors regulatory updates, maps them to affected policies, and drafts compliance impact summaries

Knowledge Grounding

Domain Constraints

Automated Validation

Digital Actuary

Actuarial Modeling & Reserve Analysis

AI accelerates mortality/morbidity modeling, scenario analysis, and reserve adequacy testing

Domain Constraints

Simulation & Testing

Runtime Oversight

Exceptions Management

Trade & Settlement Reconciliations

AI reconciles data across internal and external sources, identifies root causes of breaks, and verifies audit completeness

Knowledge Grounding

Domain Constraints

Human Oversight

Contract Analysis

Insurance Policy & Reinsurance Review

AI extracts key clauses, checks against standard terms, and flags deviations or missing provisions

Domain Constraints

Automated Validation

Simulation & Testing

Investment Research

PGIM Equity & Fixed Income Analysis

AI processes earnings calls, filings, and market data to generate research drafts with sourced citations

Human Oversight

Simulation & Testing

Domain Constraints

★ Today's Demos