SUTHERLAND

Insurance AI You Can Trust:

A Blueprint for Responsible AI Transformation for Insurance Leaders

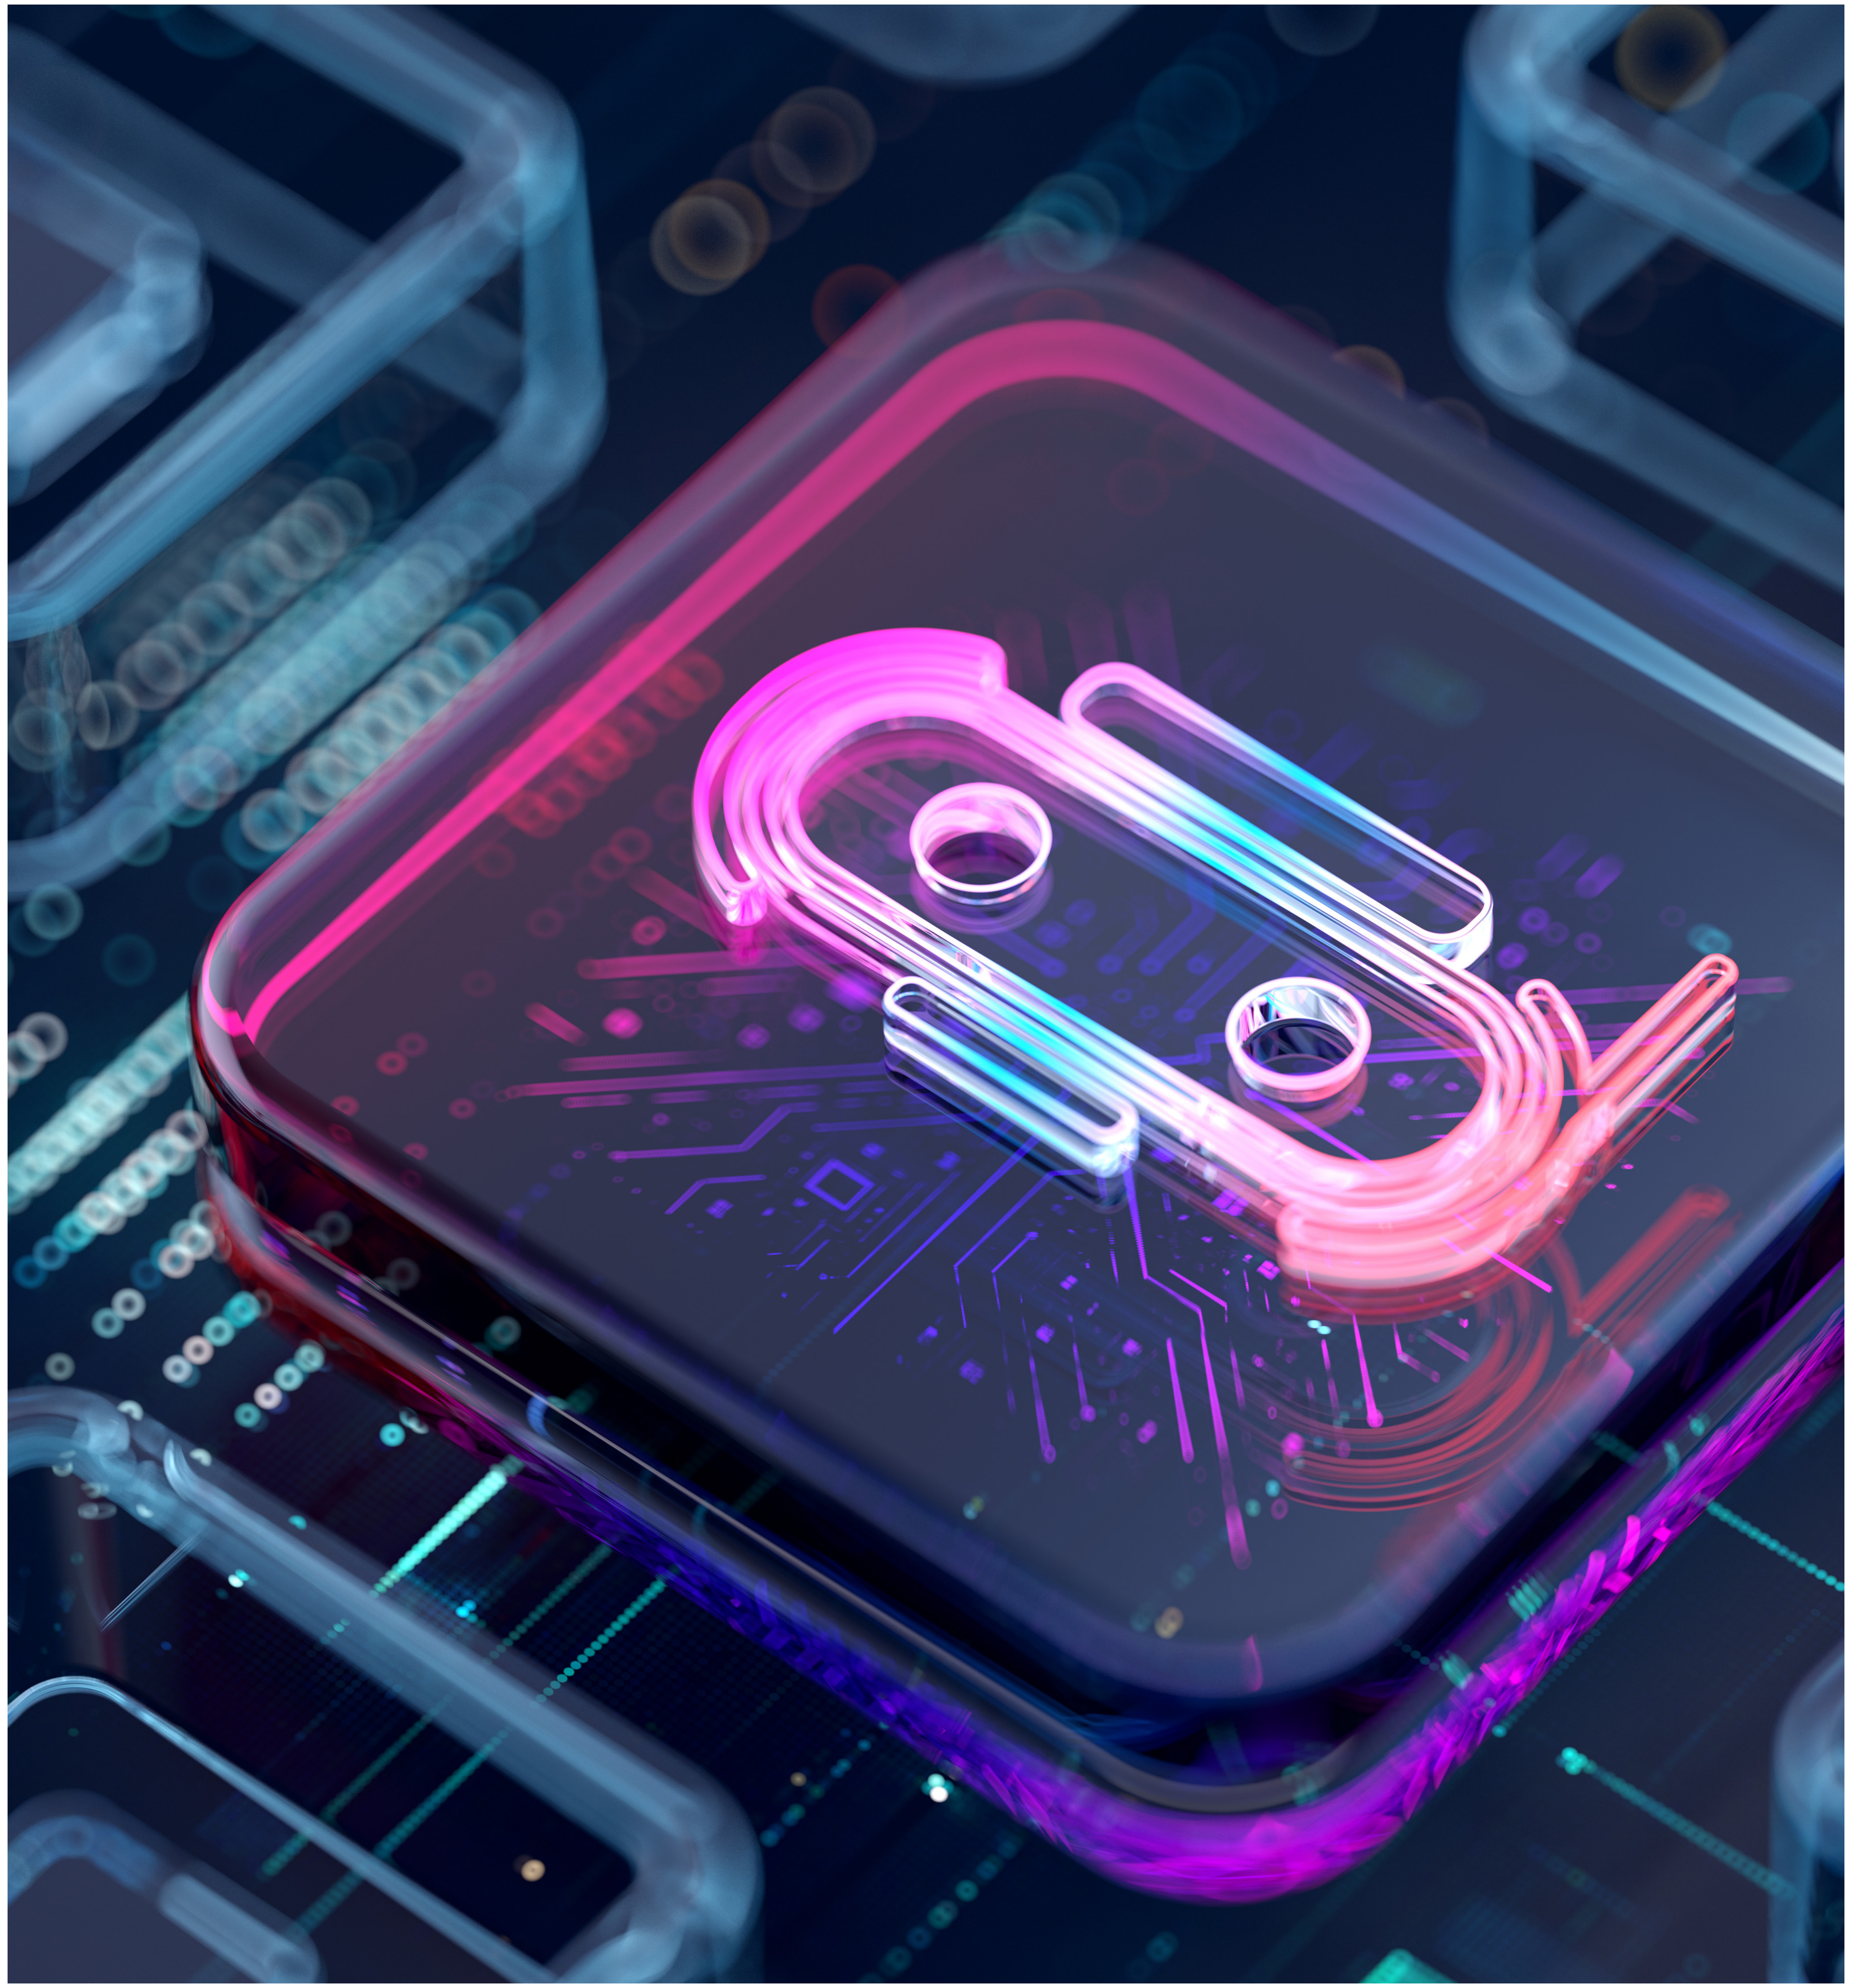

The Moment of Truth for Insurance AI

Something subtle but seismic is happening across insurance boardrooms.

For a decade, AI has been the star of slide decks, pilot programs, and digital-transformation goals. Yet here we are, in 2025, and fewer than 15% of insurers have scaled AI beyond isolated use cases (Gartner, Insurance CIO Agenda 2025). The rest remain trapped in proof-of-concept purgatory — flush with ambition, starved of traction.

The problem isn't vision. **It's trust.**

Most insurers know where AI can unlock value: faster claims, sharper underwriting, richer customer engagement. But without transparent, explainable, and regulator-ready AI, progress stalls.

And while regulators grow louder about explainability, policyholders grow less patient with opacity. The market now rewards insurers who blend speed and empathy, automation and assurance.

That's why Sutherland's Insurance AI Hub was built — to give carriers an end-to-end, enterprise-grade way to use AI responsibly, measurably, and at scale. This isn't another lab. It's an operating model for trustworthy transformation.

Breaking the Industry's Pilot Trap

Let's be honest: the industry's AI story is stuck in replay.

McKinsey estimates that while 60–70% of insurers have launched AI pilots, only one in seven has achieved measurable ROI. Why? Because the rules that made pilot success—isolated, vendor-led, time-boxed—are the same ones that make enterprise failure.

Why Pilots Fail

- **Lack of Explainability:** Black-box AI decisions expose carriers to regulatory risk.
- **Fragmented Experiments:** Disparate vendor-led pilots never connect across value chains.
- **Legacy Barriers:** Rigid policy administration systems block real-time AI integration.
- **Talent Gaps:** Lack of domain-trained AI specialists delays implementation.
- **Compliance Anxiety:** Fear of data exposure or bias halts deployment.

In other words, the problem isn't data science — it's the operating discipline.

To scale, insurers need AI that's explainable to regulators, auditable by risk officers, and understandable by underwriters.

They need AI that humans — not just algorithms — can trust.

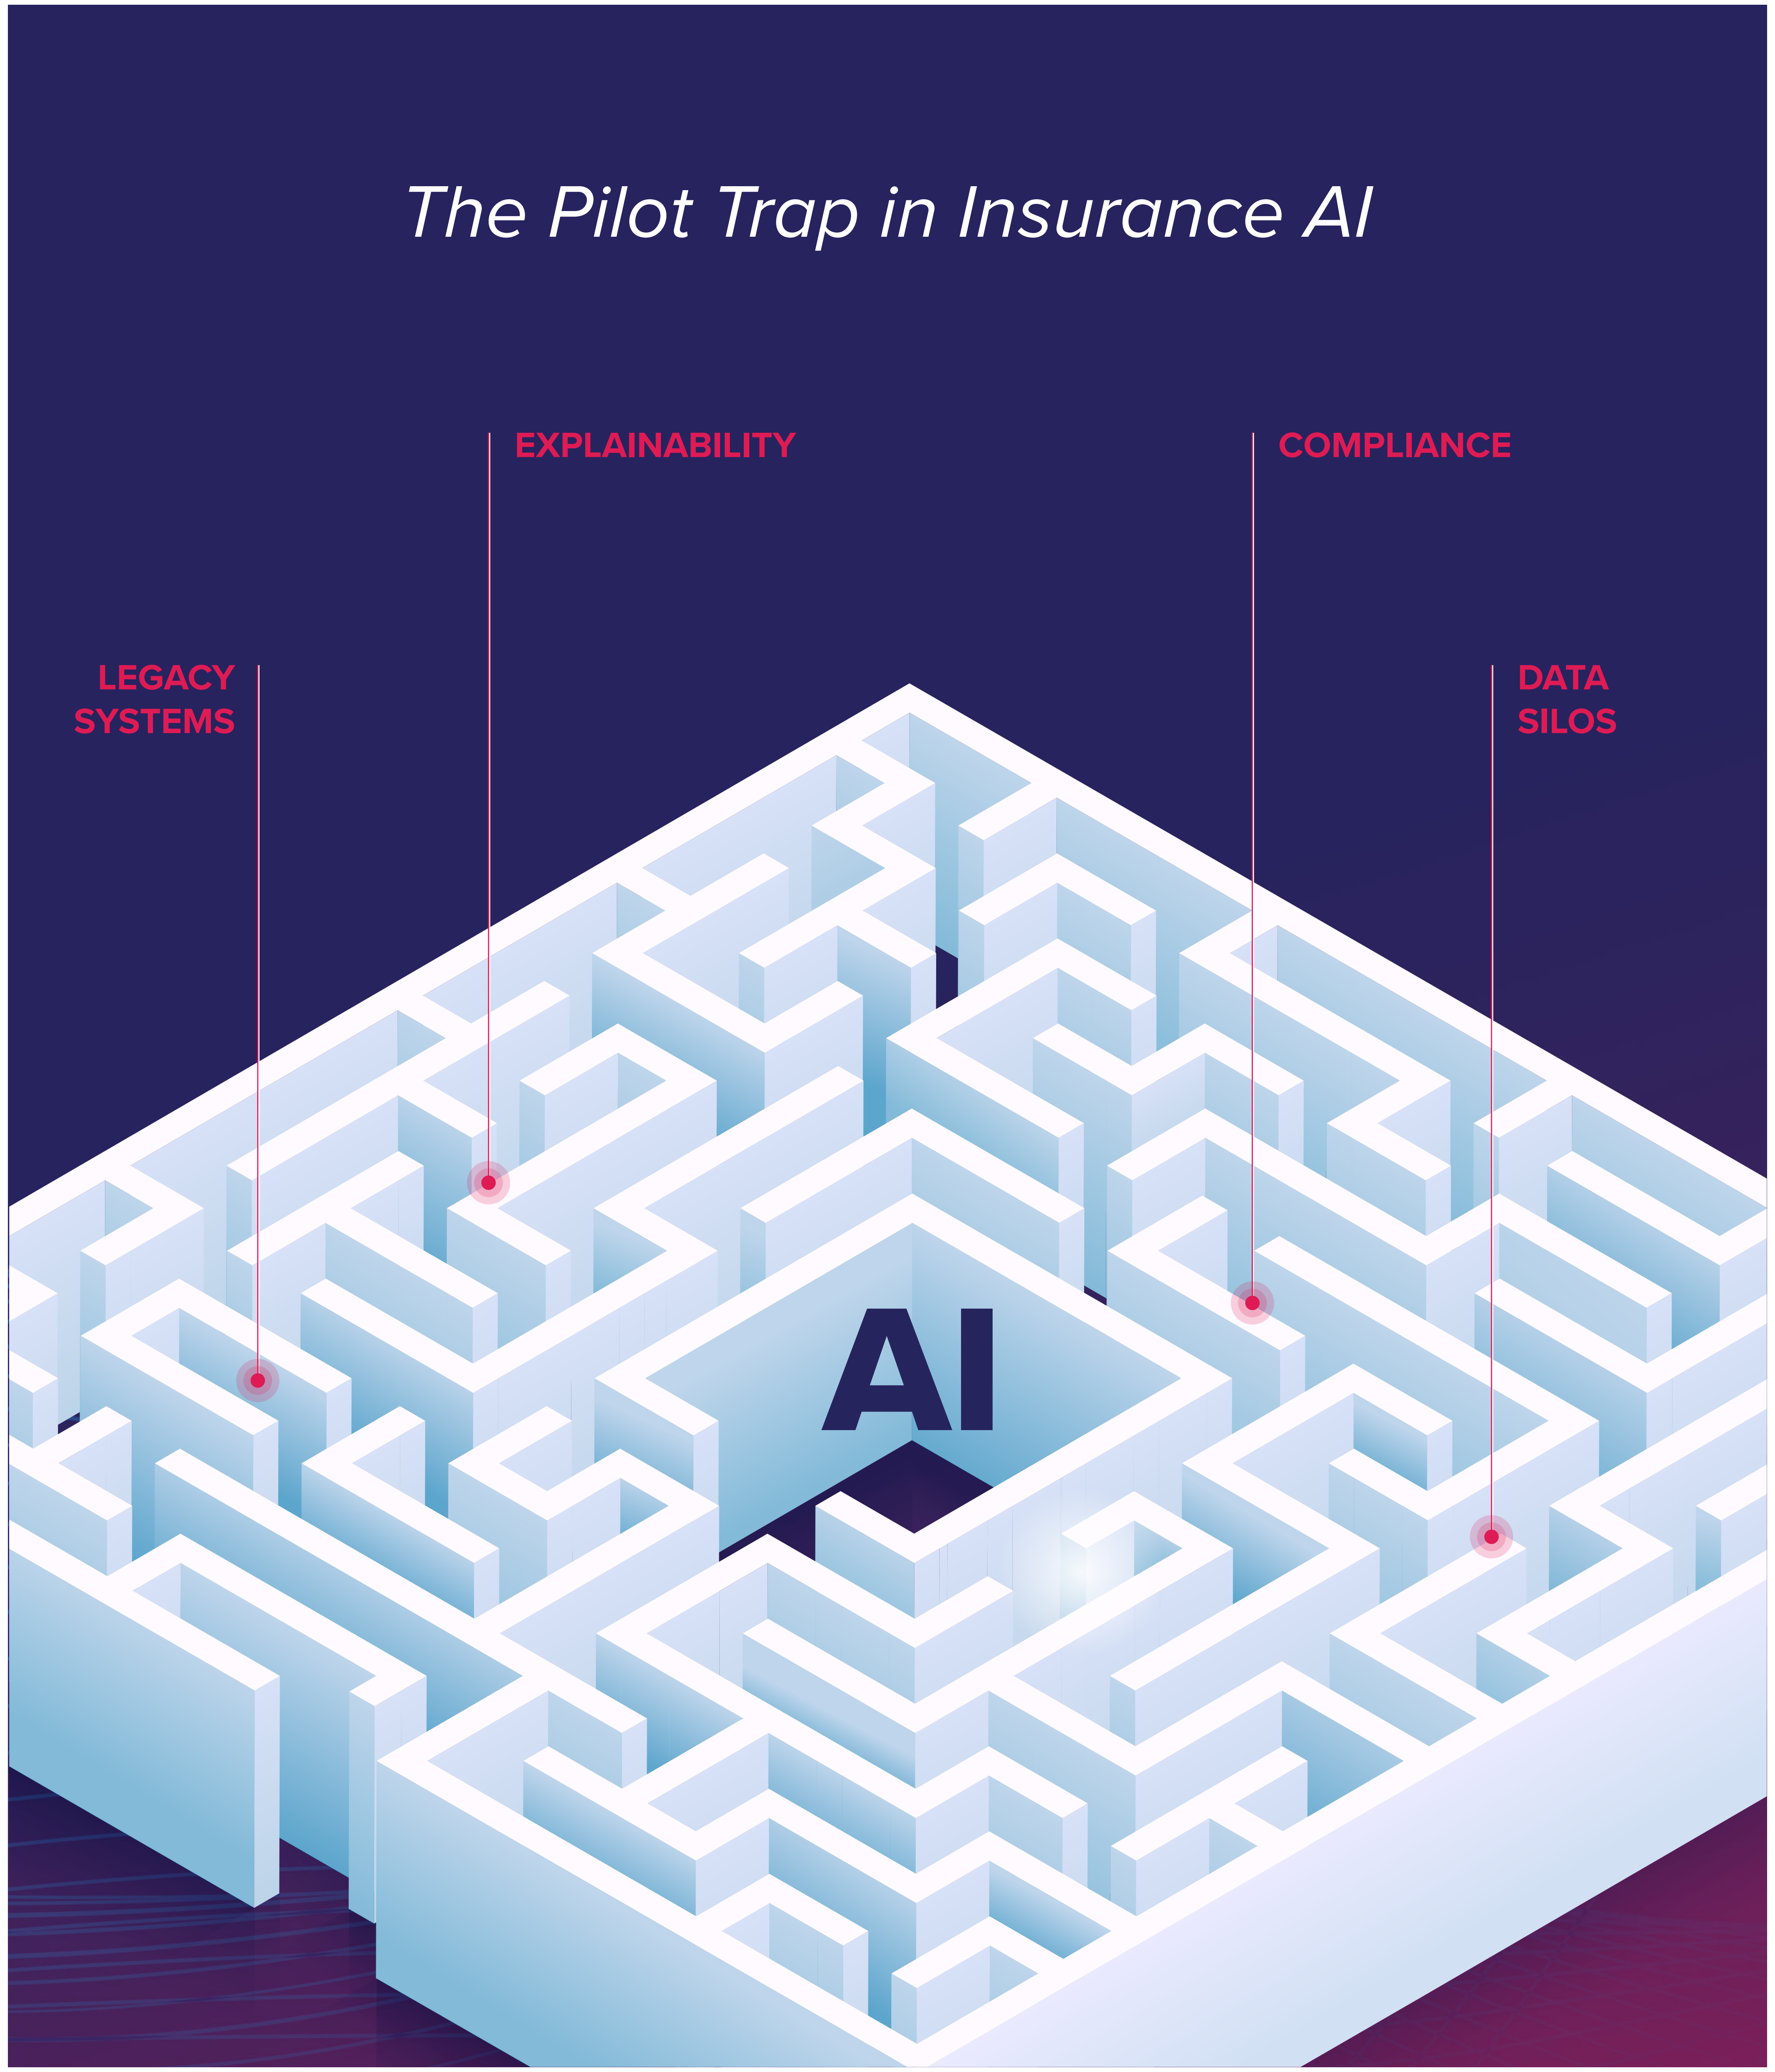

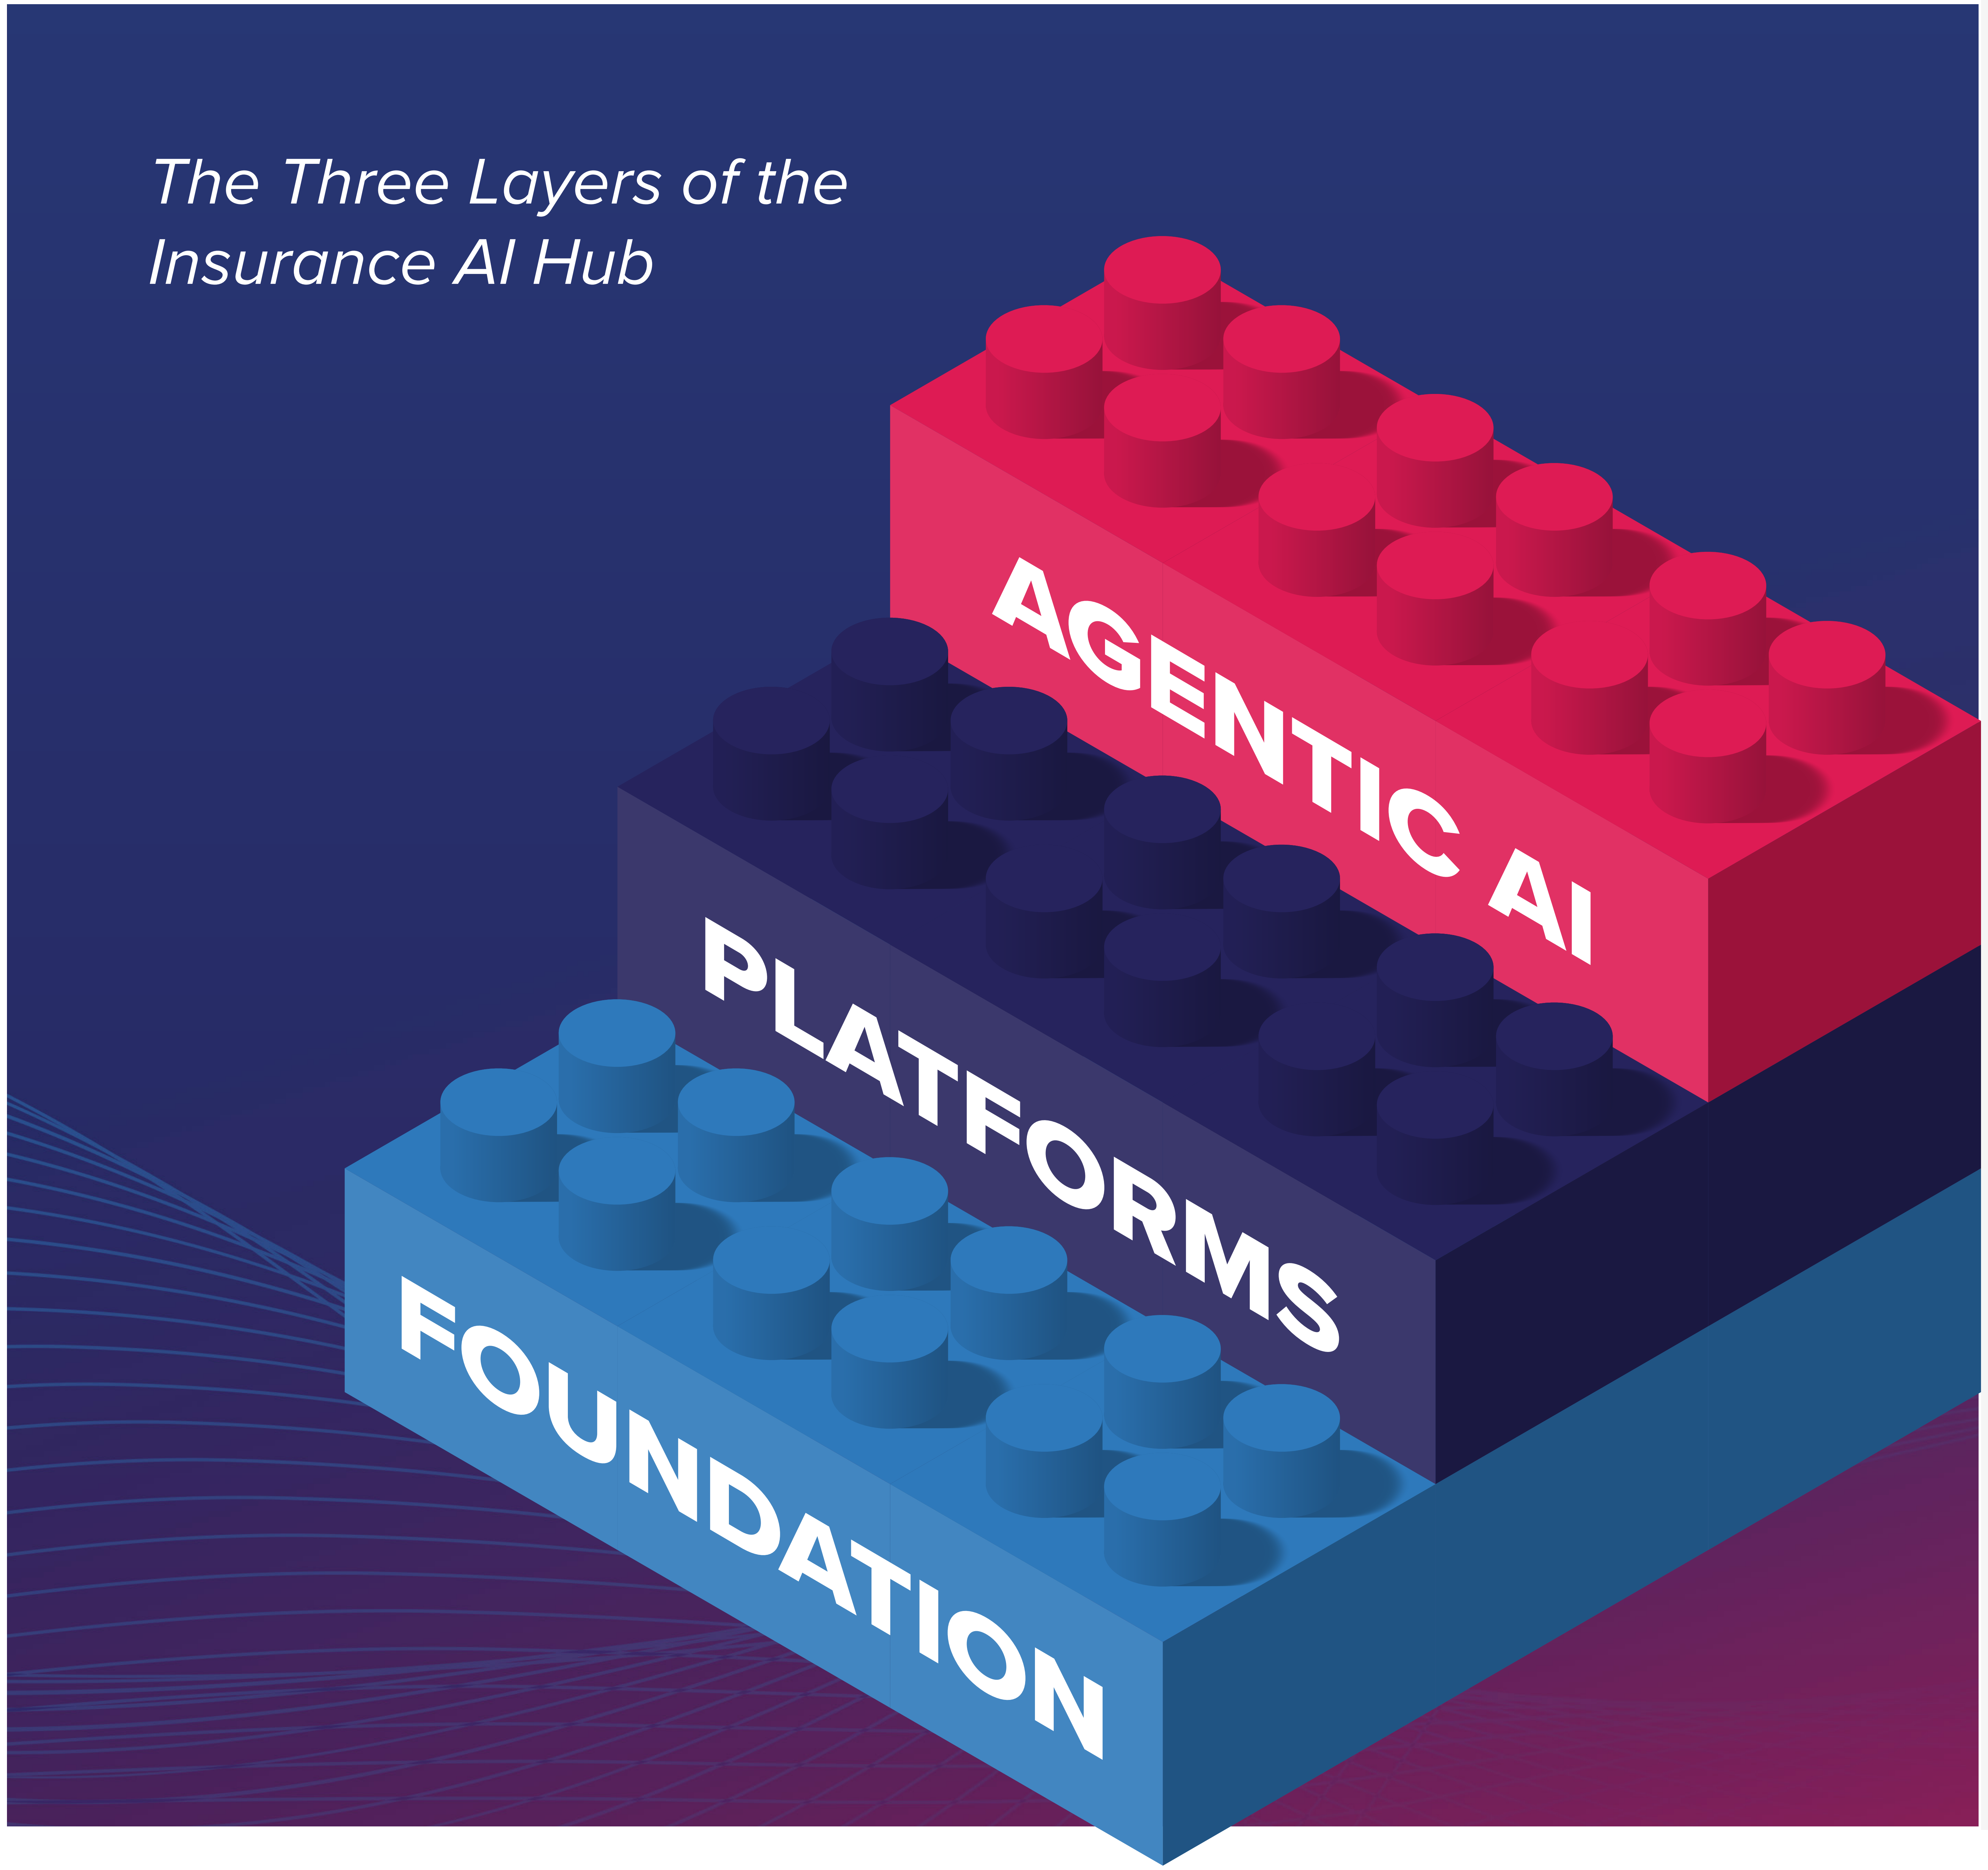

The Insurance AI Hub: From Point Solutions to a Platform for Trust

Overview

Every insurer trying to scale AI eventually runs into the same invisible wall: governance, interoperability, and context.

You can't scale what you can't control, and you can't trust what you can't explain. That's why the Insurance AI Hub isn't another collection of algorithms. The **Insurance AI Hub** is Sutherland's end-to-end AI ecosystem, purpose-built for insurers. It combines over **60 domain-trained AI Agents**, secure platform infrastructure, and an ethical AI foundation.

Each layer — **Foundation, Platforms, and Agentic AI** — works as part of a unified, explainable AI operating model, enabling insurers to achieve enterprise-grade outcomes safely and at scale.

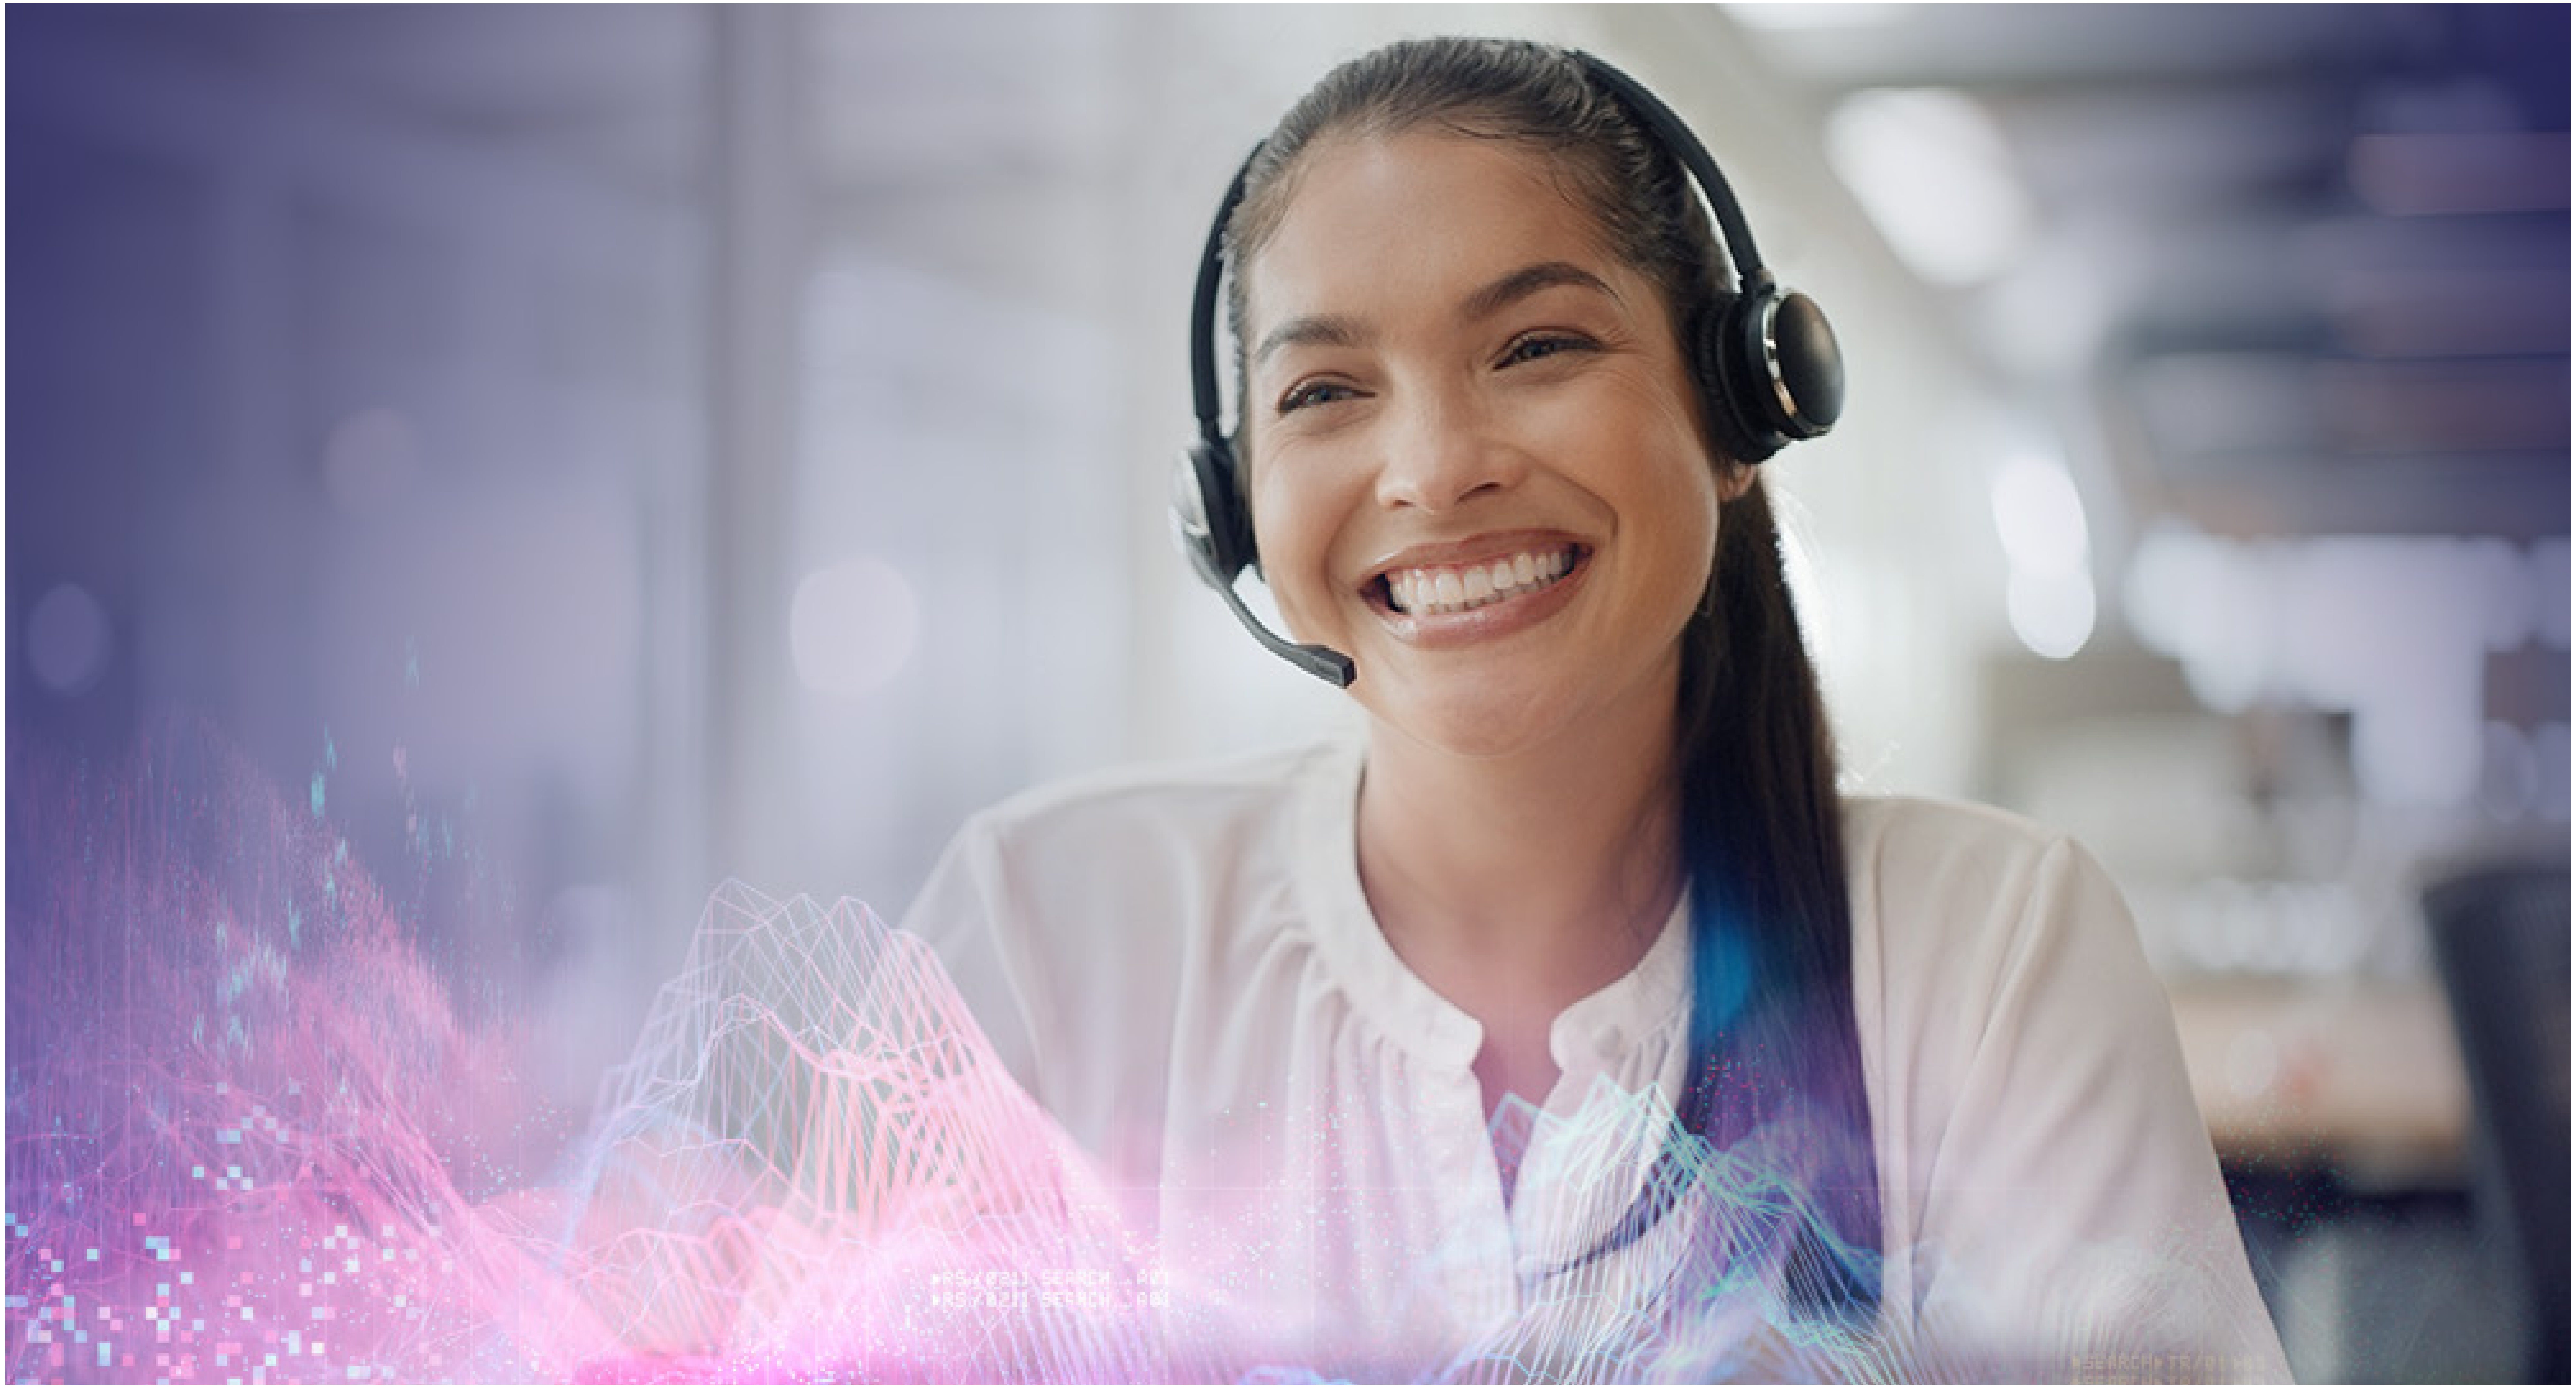

The Three Layers of Trustworthy Transformation

3.1 The Foundation: The Bedrock of Responsible AI

If AI is to become the new nervous system of an insurer, then the Foundation is its compliance DNA. This base layer encodes how transparency, fairness, and resilience are engineered in — not patched later.

Component	Purpose
Studio	A creation workspace for low-code and pro-code AI solutions trained on insurance data.
Orchestrator	The command layer that connects humans, models, APIs, and workflows.
Trust Center	Bias testing, fairness dashboards, and explainability reports — regulator-ready by design.
Test Center	Runs stress tests on real claims and policy data before deployment.
Knowledge Center	Creates a living knowledge graph that gives agents context and reasoning power.
Cloud Center	Ensures secure, multi-region deployment with enterprise governance.

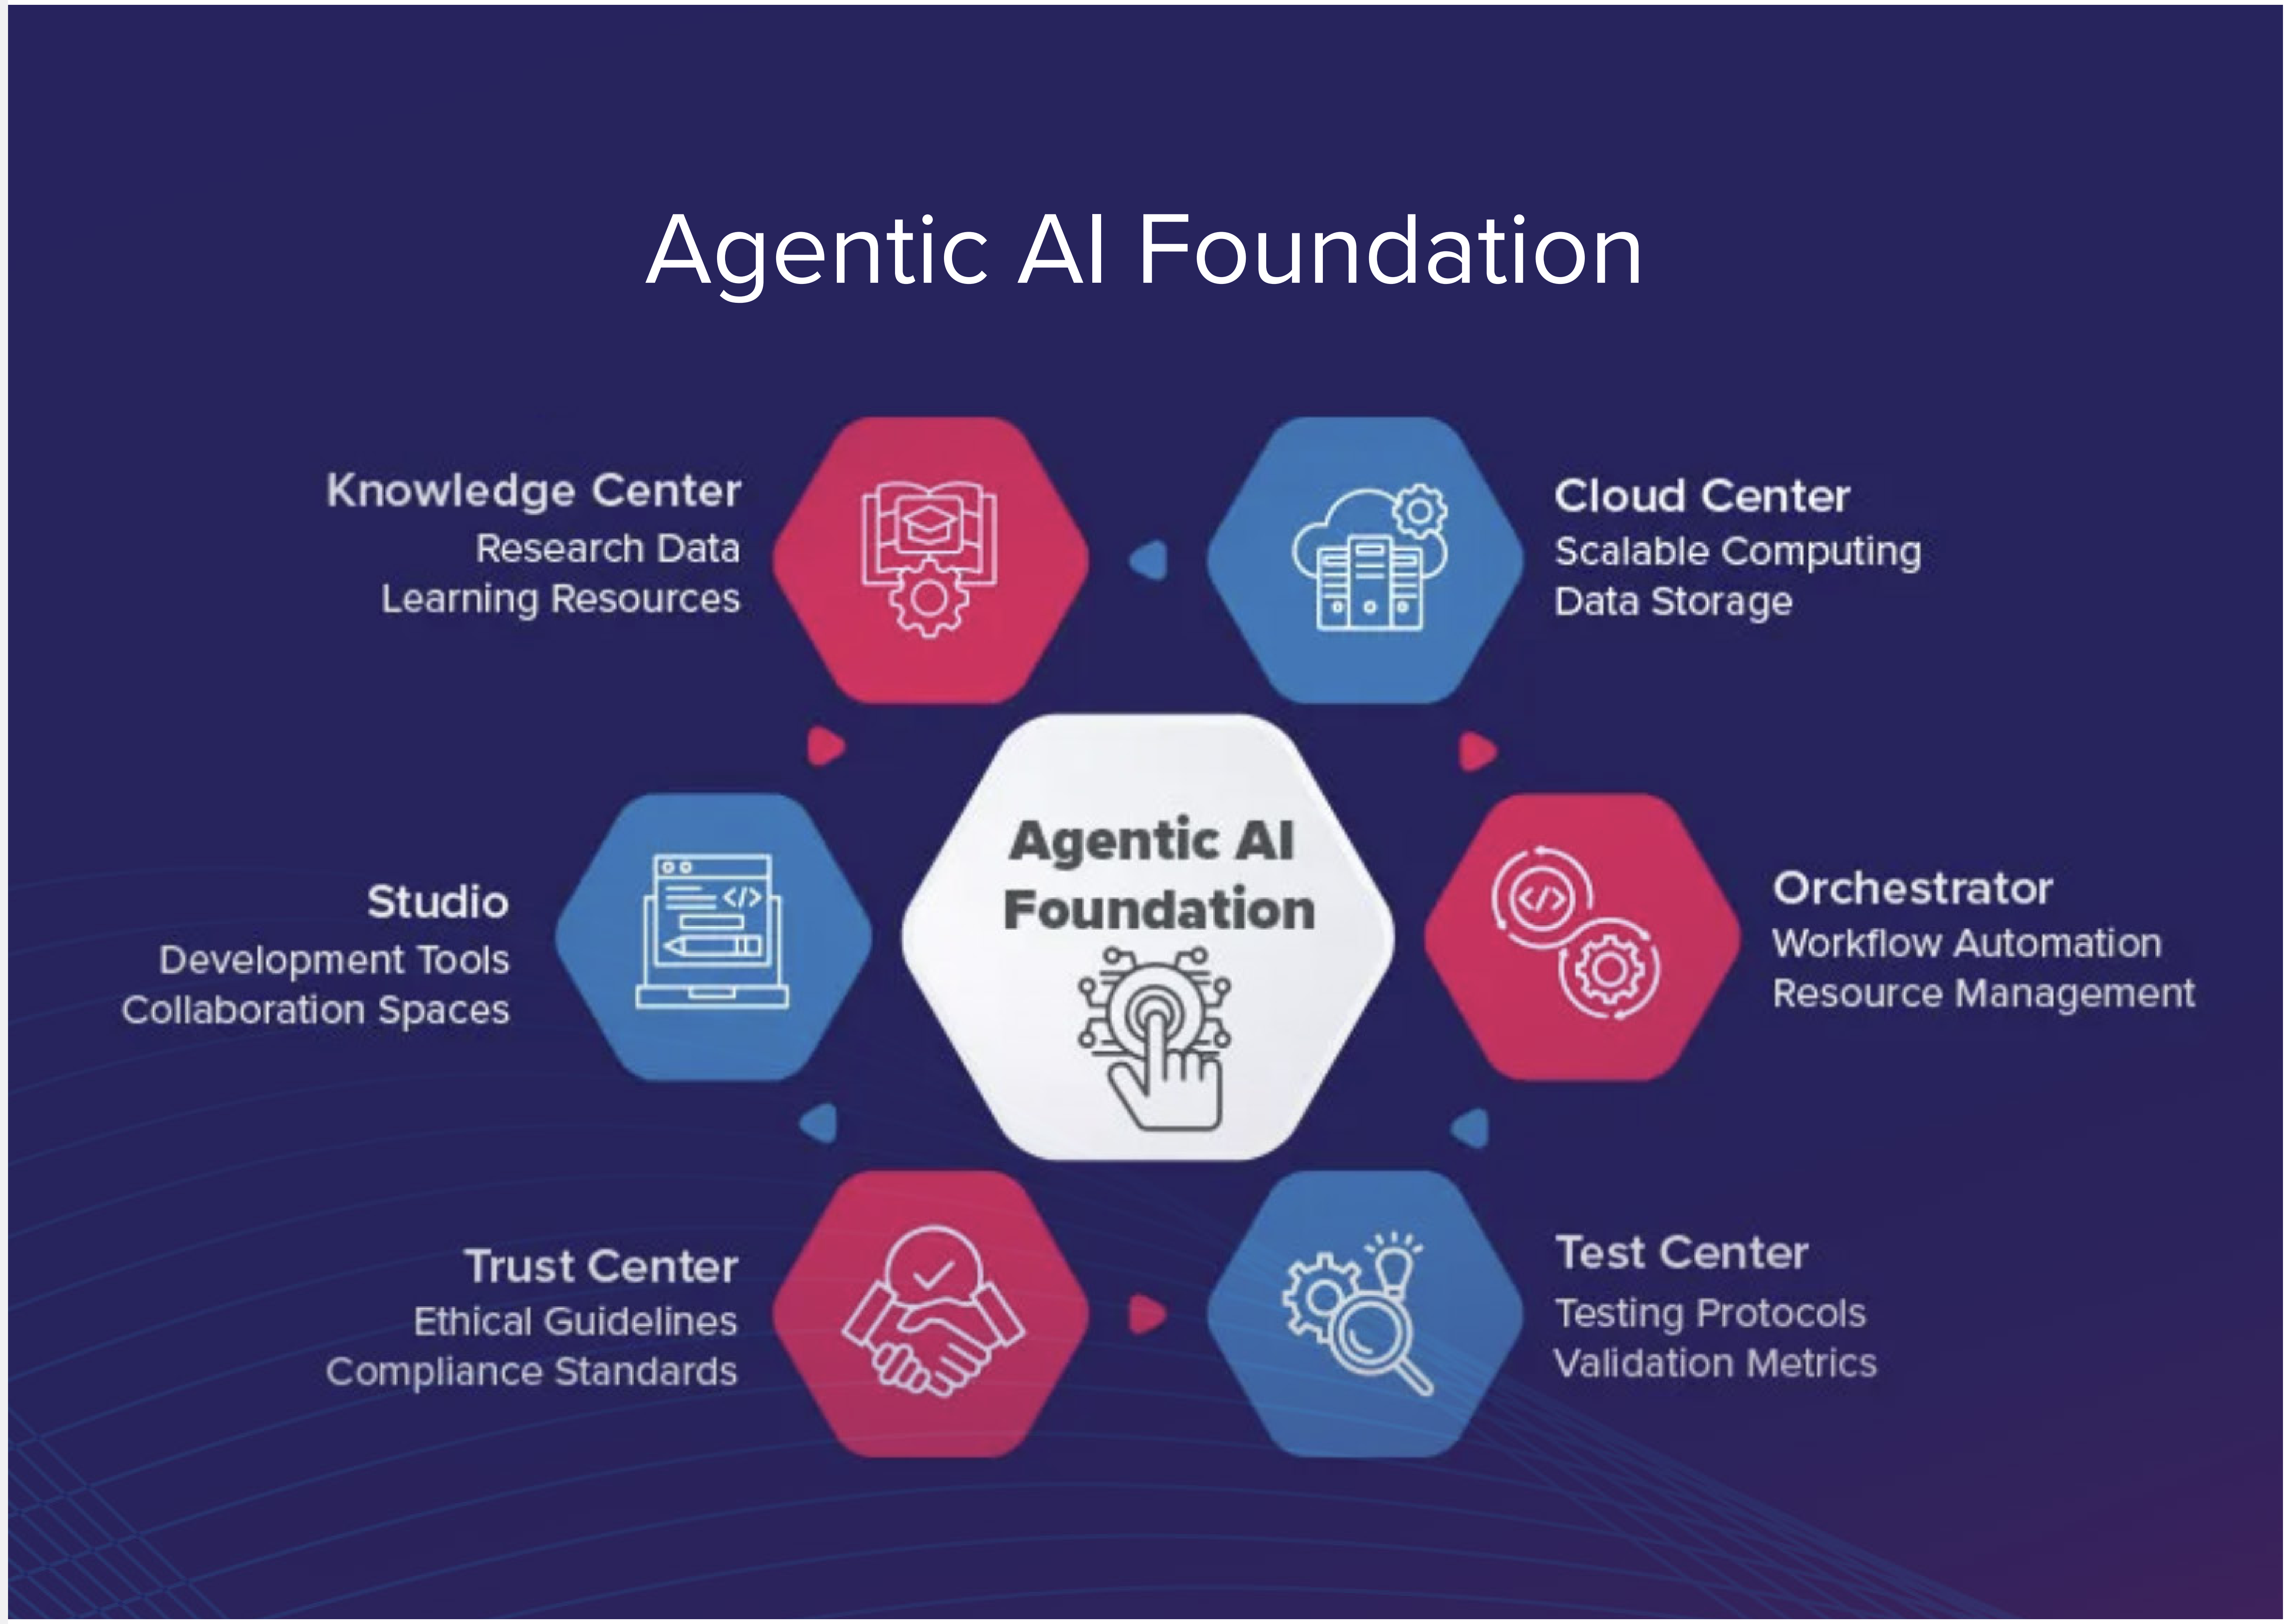

This foundation mirrors what Gartner calls the "Responsible AI Core" — the combination of governance, security, and model oversight that separates operational AI from experimental AI (Gartner, AI Governance Market Guide 2025).

It's not just policy — it's infrastructure for accountability.

The 8 Guardrails of Responsible AI

1

Custom Compliance Guardrails

HIPAA, PCI-DSS, SOC 2, and state laws built into every workflow.

2

Private Model Hosting

AI runs in your private cloud or air-gapped data centers.

3

Red Teaming & Adversarial Testing

Continuous simulation to detect vulnerabilities before deployment.

4

Secure Data Residency

Customer and claims data never leave the client environment.

5

Granular Access Control

Role-based permissions and least-privilege principles.

6

Zero Retention Model

No third-party model retains insurer data.

7

Automated Redaction & PII Masking

All sensitive identifiers are masked pre-processing.

8

Auditability & Traceability

Every decision and prompt is logged for regulator-ready audit trails.

3.2 The Platforms: The Engine of Scalable Transformation

Most carriers don't fail because AI models misfire; they fail because the plumbing can't carry the pressure. Legacy stacks built for batch processing, multiple policy-admin systems, and siloed data lakes choke when asked to deliver real-time personalization or claims decisions. The Platform layer solves this by introducing an ecosystem that's modular, open, and connected — so insurers can integrate human expertise, partner innovation, and automation without rewriting the enterprise each time. Sutherland's Platform layer powers enterprise-grade scalability through modular, interoperable systems.

- **Experience Platforms** — CX360, Media Labs: Deliver natural, conversational engagement across FNOL to renewals.
- **Automation Platforms** — Robility®, Extract, Strala: Digitize enrollment, endorsements, and servicing at scale.
- **Partner Ecosystem** — Federato, Five Sigma, Solvarys: Federate specialized AI engines for underwriting and claims.

This architecture gives insurers plug-and-play interoperability, eliminating the decades-old friction of disjointed tools.

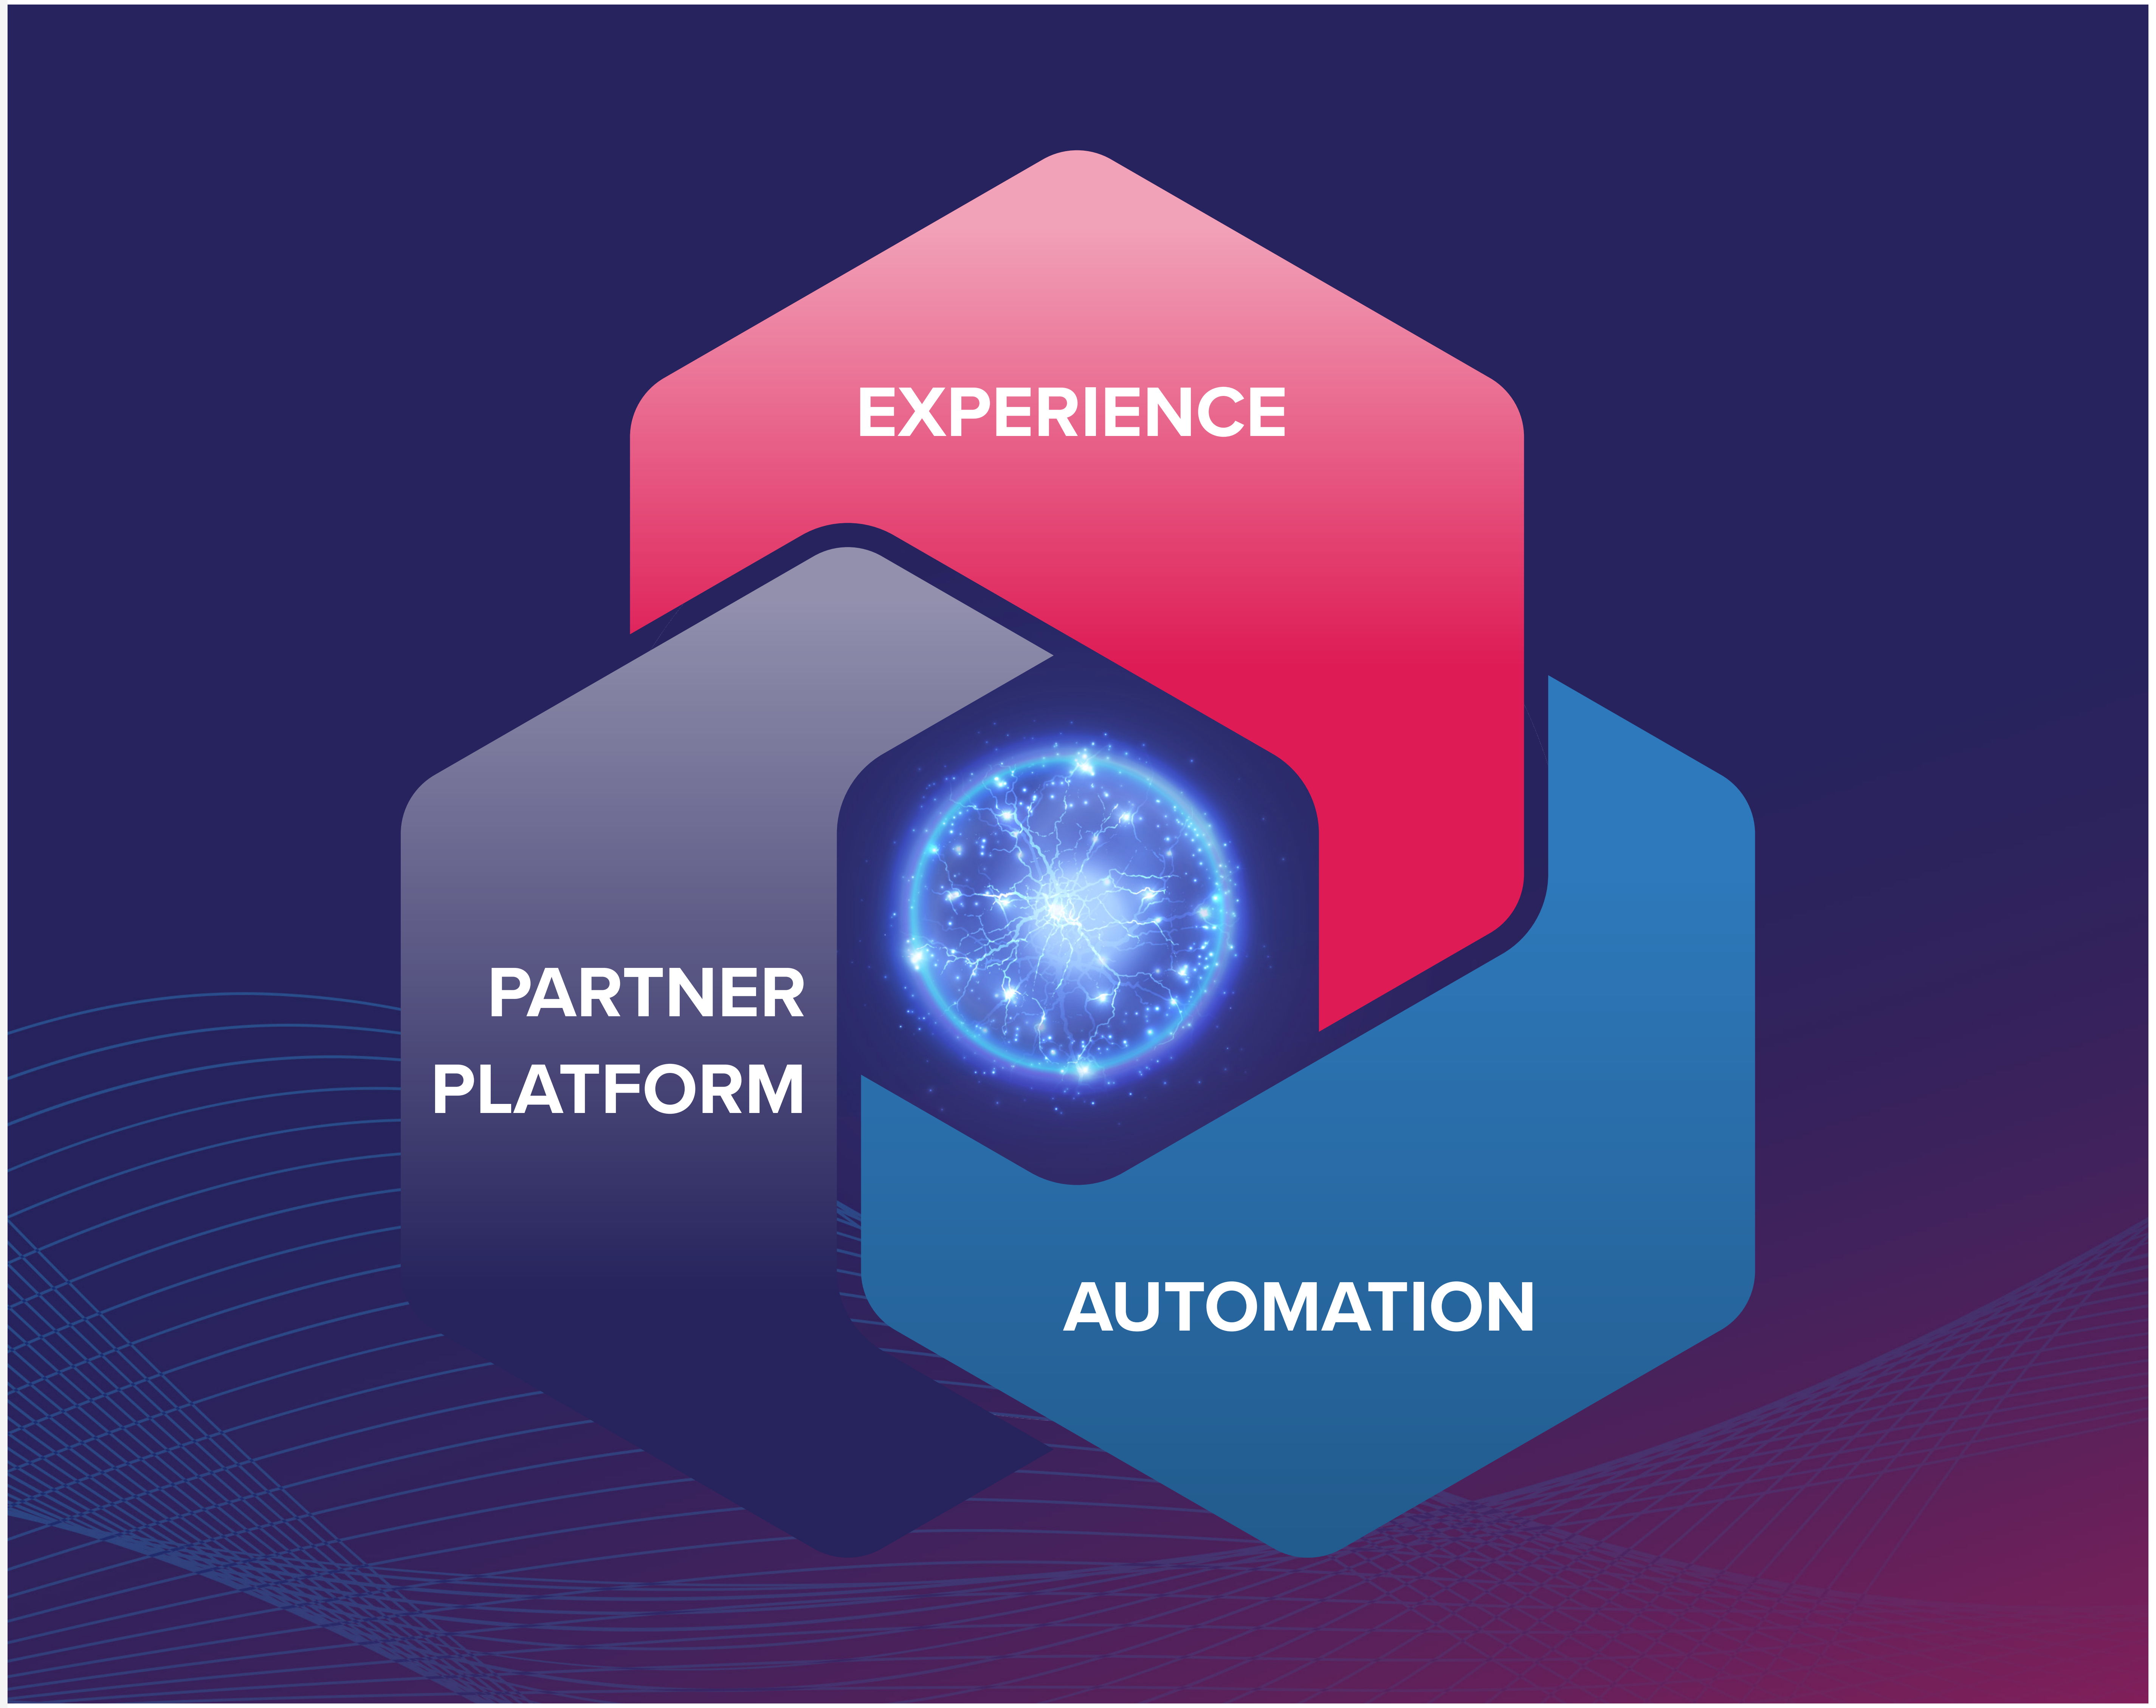

3.3 The Agentic AI Layer: The Workforce of the Future

The final layer brings intelligence to life.

Agentic systems — autonomous yet accountable — are the fastest-rising trend on Gartner's 2025 Hype Cycle. For insurers, they represent the moment AI stops being a tool and starts acting like a colleague.

Each of the Hub's 60+ agents is trained on insurance semantics, policies, and compliance rules. They collaborate, hand off tasks, and document every decision, so nothing disappears into a black box. Collectively, they redefine productivity: faster claims, cleaner underwriting, and customer conversations that feel genuinely human.

Meet Your 60+ AI Agents

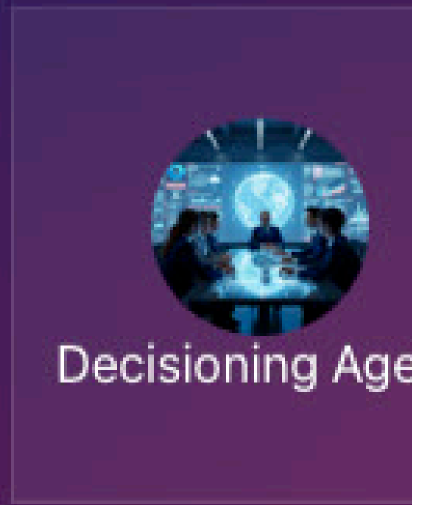

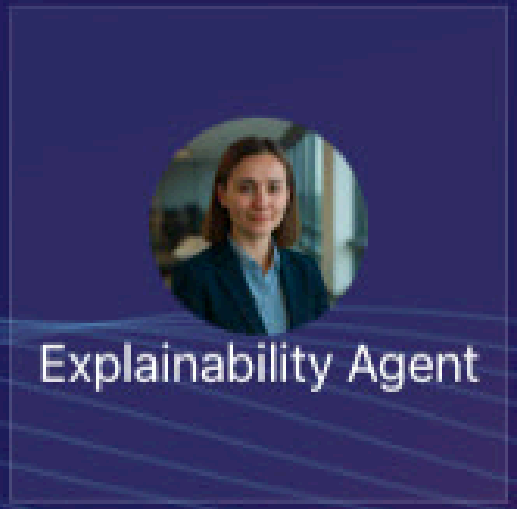

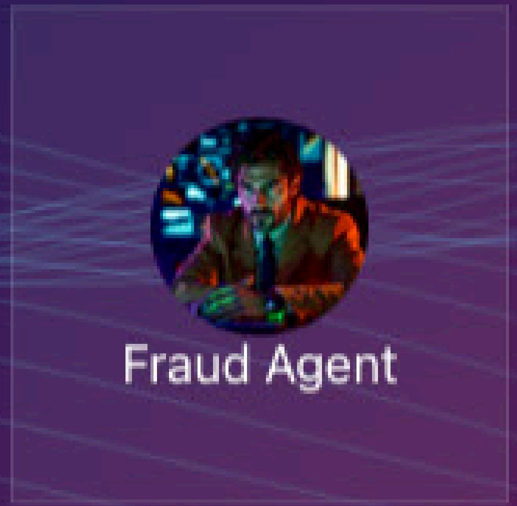

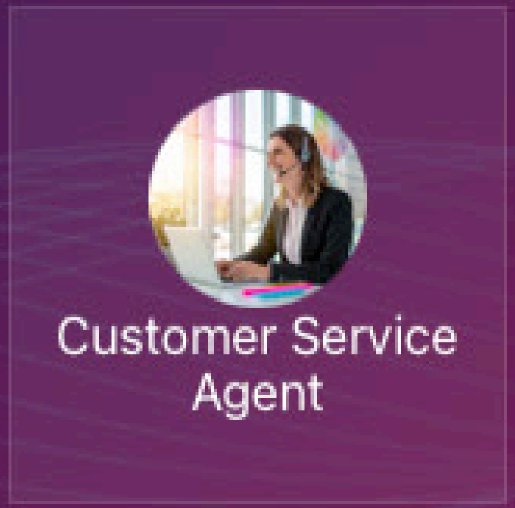

Agent Families and Example Use Cases

Function	Example AI Agents	Outcomes
Claims & Servicing	Claims FNOL Agent, Death Claims Validator, Fraud Detection Agent, Voice FNOL Agent	**30% faster claims** **12% leakage reduction**
Underwriting	Connected Underwriting Agent, Eligibility Agent, Decisioning Agent	**16% higher win rates** **3.7× growth in target business**
Customer Experience	Conversationalist Agent, Renewal Negotiator, Troubleshooter Agent	**+10–15 NPS uplift**
Operations & Policy Admin	Policy Conversion Agent, Enrollment Processor, Clerk Agent	**50% efficiency gain**
Governance & Risk	Regulator Agent, Narrator (Explainability Agent)	**100% audit traceability**

These agents can be orchestrated together into multi-agent systems that execute complex workflows across claims, underwriting, and servicing — all while maintaining human-in-the-loop oversight.

Insurance AI in Action: Proof, Not Promise

For years, AI in insurance has been a story of potential. 2025 is the year it became a story of proof.

Across carriers using Sutherland's Hub, agentic systems have moved beyond pilots to measurable P&L impact — compressing claim cycles, trimming leakage, and lifting NPS.

This is how the shift looks in numbers and behavior:

- Decisions once bottlenecked by manual review now flow through explainable pipelines.
- Front-line staff spend less time reconciling data and more time advising customers.
- Compliance officers receive auto-generated audit trails instead of weekend fire drills.

McKinsey's 2024 State of AI report found that firms scaling AI enterprise-wide achieved three times faster EBIT growth than those stuck in pilots — and insurers are no exception. Scale and trust now move in tandem.

Here's what Sutherland's Insurance AI Hub has delivered:

Agentic AI for Claims & Enrollment

30% faster cycle time

12% less leakage

10+ NPS increase

Connected Underwriting

30% efficiency gain

16% higher win rate

Cognilink Document Intelligence

30% OPEX reduction

1–1.5% indemnity improvement

Voice AI for Customer Engagement

+10 NPS points

20% higher contact center efficiency

What once required months of systems integration and governance reviews can now be orchestrated in weeks. Each deployment follows Sutherland's **"Think Big, Start Small, Scale Fast"** model — enabling carriers to transition from limited pilots to measurable, enterprise-scale impact.

What Insurance Leaders Need Now

Persona	What Keeps Them Up at Night	What the Hub Delivers	Quantified Value
CIO / CTO	Integrating AI securely; compliance risk	Explainable, private, zero-retention AI deployment	30–40% faster rollout, zero breaches
COO / Head of Ops	Manual inefficiencies; data silos	60+ AI Agents automate complex workflows	50% productivity gain
Head of Claims	High leakage, low satisfaction	Voice + Agentic AI	30% faster claims, +10 NPS
Head of Underwriting	Slow triage and low hit ratio	Connected Underwriting	16% win rate improvement
Chief Risk Officer	Auditability, bias	Trust Center + ethical guardrails	100% traceability, 0 violations

This isn't about selling automation. It's about restoring confidence — to regulators, to boards, to customers.

Every leader faces a different flavor of sleepless night.

For the CIO, it's security and sprawl; for the COO, throughput; for the CRO, auditability.

Yet their endgame is shared: confidence — in data, in decisions, in compliance.

That's why the Insurance AI Hub resonates in the boardroom: it translates AI into the language of value creation and risk mitigation. It allows technology leaders to promise innovation without compromising the regulator's patience.

Beyond metrics, what CXOs consistently say they value is control with clarity: the ability to trace a decision, defend it in an exam, and still move faster than last year. The Hub makes that possible.

Building Responsible AI at Scale

Scaling responsibly is not optional; it’s existential. As of 2025, the NAIC, NYDFS, and Colorado Division of Insurance have all issued AI governance expectations requiring documented explainability, bias testing, and data lineage.

Sutherland’s Trust Center operationalizes those standards through a lifecycle every carrier can audit.

A Governance Framework for Insurance AI

1

Design

Start with transparent intent and ethical design review.

2

Build

Validate bias, data integrity, and fairness metrics

3

Test

Real-world simulation under multiple risk scenarios.

4

Deploy

Private hosting, compliance alignment, explainability dashboards.

5

Monitor

Continuous audit, red teaming, retraining cycles.

Beyond Automation — The Human Dimension of Trust

It's tempting to think of AI as the replacement for human judgment.

But the real advantage comes when it amplifies it.

Front-line adjusters, underwriters, and service representatives using agentic AI report a different kind of productivity — one that is less about speed alone, and more about confidence. They can see the model's reasoning, challenge it, and learn from it.

That transparency builds a culture where humans stay in charge and technology earns its keep.

As one chief claims officer told us, **"We didn't automate empathy — we freed time to deliver it."**

"Human-in-the-loop is how trust scales."

Trust, in the end, is not coded — it's co-created.

The Road Ahead — Scaling Responsibly

Scaling AI in insurance now resembles building a regulated enterprise platform, not an experiment. Those who treat it that way are already pulling ahead.

The near-term roadmap unfolds in four disciplined phases:

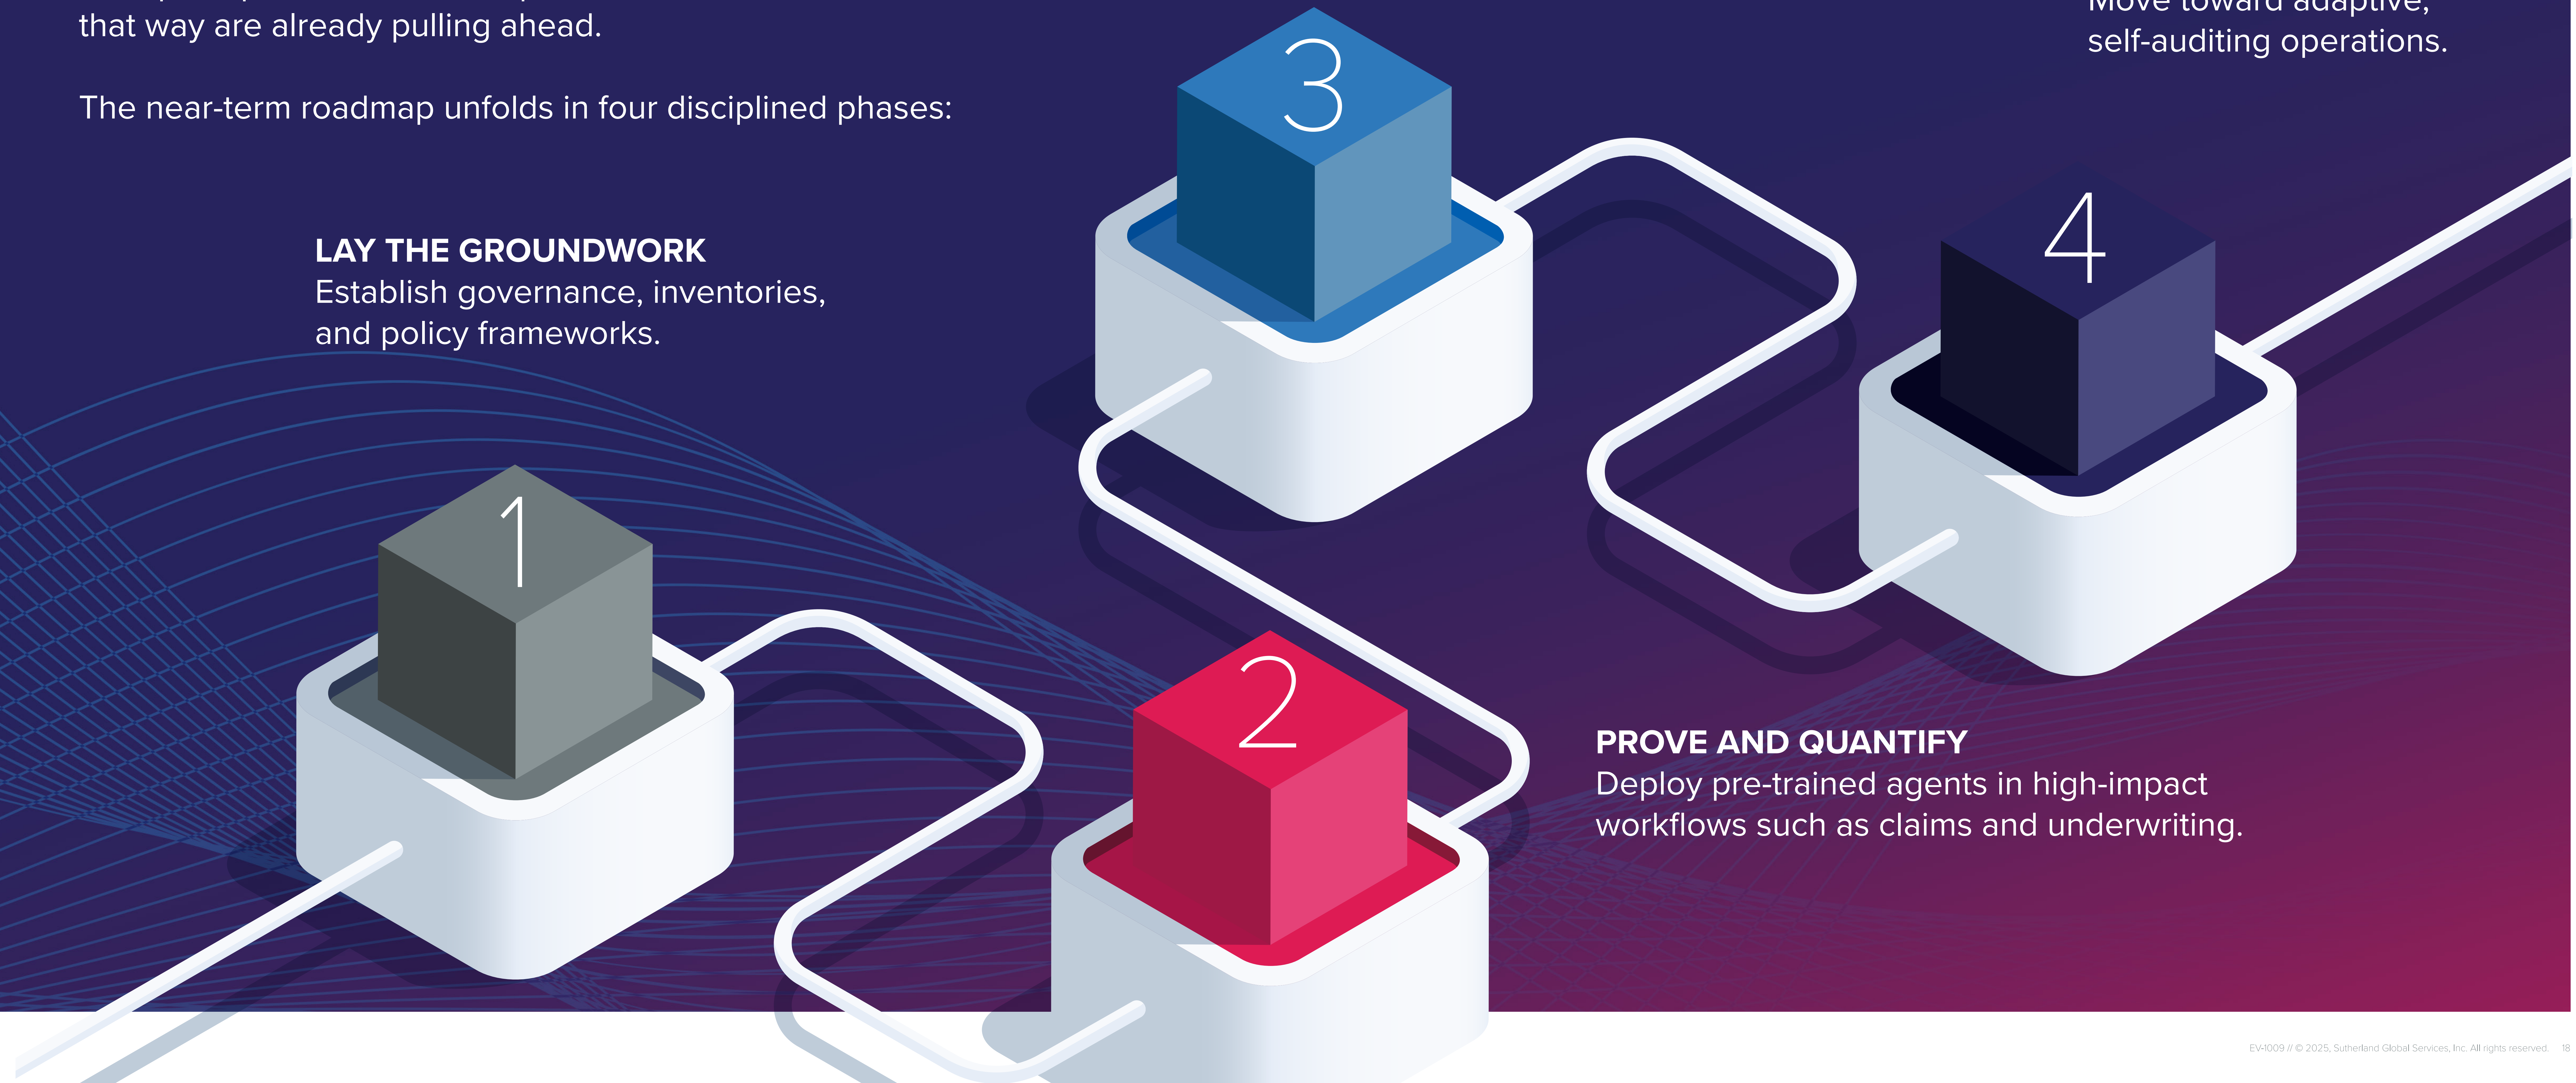

LAY THE GROUNDWORK
Establish governance, inventories, and policy frameworks.

PROVE AND QUANTIFY
Deploy pre-trained agents in high-impact workflows such as claims and underwriting.

INTEGRATE
Connect partner ecosystems, data lakes, and orchestration layers.

EVOLVE
Move toward adaptive, self-auditing operations.

Trust Is the New Currency of Insurance AI

The insurance industry has always been built on trust — between carriers and customers, brokers and underwriters, regulators and society. AI doesn't change that truth; it magnifies it.

Speed and efficiency win quarters; trust wins decades. The carriers that master explainable, compliant, and human-centered AI will own the next era of advantage.

Sutherland's Insurance AI Hub is more than a technology stack. It's a way to operationalize integrity at scale — where every automated decision can be explained, every model can be audited, and every customer interaction can still feel human.

Because in the end, the question isn't *"How intelligent is your AI?"* It's *"Can you trust it with your reputation?"*

GenAI powered platforms and 60+ domain-trained AI agents across P&C, Life, and Group Benefits to power seamless, personalised experiences and smarter growth from underwriting to claims and operations.

https://www.sutherlandglobal.com/industries/insurance/insurance-agentic-ai-hub

Artificial Intelligence. Automation. Cloud Engineering. Advanced Analytics. For Enterprises, these are key factors of success. For us, they're our core expertise.

We work with global iconic brands. We bring them a unique value proposition through market-leading technologies and business process excellence. At the heart of it all is Digital Engineering – the foundation that powers rapid innovation and scalable business transformation.

We've created over 200 unique inventions under several patents across AI and other emerging technologies. Leveraging our advanced products and platforms, we drive digital transformation at scale, optimize critical business operations, reinvent experiences and pioneer new solutions, all provided through a seamless "as-a-service" model.

For each company, we provide new keys for their businesses, the people they work with, and the customers they serve. With proven strategies and agile execution, we don't just enable change – we engineer digital outcomes.

SUTHERLAND®

digital outcomes.

sutherlandglobal.com