cognizant | venbrook
Boldly Go Forward

Reinforcement before autonomy

Engineering trustworthy autonomy in insurance AI

Overview

The insurance industry is undergoing a pivotal transformation, yet true AI autonomy remains out of reach. While a few leaders have successfully scaled AI, most insurers are still constrained by legacy infrastructure, regulatory caution and immature governance frameworks. Even among advanced adopters, fully delegating operations to autonomous AI agents is not yet feasible.

This paper explores two strategic paths forward: the long-term pursuit of artificial general intelligence (AGI) and the immediate application of reinforcement learning (RL). We introduce a novel reinforcement-switch framework which combines continuous learning with proactive human-AI control transfers to enable accountable autonomy. This model ensures resilience in dynamic environments by embedding trust, reversibility, and oversight into AI operations. It represents a fail-forward approach to engineering safe, scalable autonomy in insurance.

About the authors

Dr. Venkatesh Upadrista
Global Head of Transformation, BFSI-IOA
Drvenkatesh.upadrista@cognizant.com

Follow www.Cognizant.com

Dr. Venkatesh Upadrista leads global transformation for the BFSI-IOA vertical at Cognizant. In this role, he is responsible for driving AI-led transformation across the unit, ensuring customer success and enhancing the delivery of modern business operations within the financial services and insurance sectors.

Justin Slaten
Chief Information Officer, Venbrook Companies
Jslaten@Venbrook.com

Follow www.Venbrook.com

Justin Slaten is an accomplished technology executive with over 25 years of management experience. He is recognized for his innovative approach to product development, business process improvement, and scaling teams and operations. His leadership consistently drives enhanced productivity and sustainable growth.

Table of contents

Background—Industry view

The state of the industry—A fragmented landscape of maturity

The insurance sector, long defined by fiduciary duty and regulatory conservatism, now faces a crossroads. A growing body of research reveals a stark truth: While a handful of global insurers are exploring AI as a transformative tool, most companies remain early in their journey.

According to BCG, only 4% of companies have matured to the point of running AI as a full-fledged value engine, while 22% are scaling AI to generate impact. The remaining 74% are still in the early phases—either doing very little with AI or limited to proof-of-concept efforts. Meanwhile, research from Forrester indicates that fewer than 5% of insurers will realize significant, direct AI-driven revenue gains—highlighting widespread experimentation overshadowed by legacy systems and fragmented governance models.

Our own journey reflects this broader reality. Legacy platforms, built for a different era, have inhibited data fluidity, delayed decision-making and limited our ability to generate insights at the speed today's market demands. Before activating meaningful AI capabilities, we had to modernize the enterprise data landscape—integrating disparate systems into a unified, trusted and accessible foundation. This isn't merely a technical upgrade; it is a strategic imperative.

The lesson is clear: Ambition alone does not unlock AI's potential. Enterprises must first re-architect their core—data, infrastructure and governance—before AI can scale with impact.

Even among insurers who have streamlined legacy systems and begun experimenting with foundational models, few have embedded the transparency, auditability and governance required to inspire confidence in regulated environments. Perhaps because, in a domain where risk is existential, the margin for error is virtually zero.

Until the industry aligns its architectural readiness with its strategic aspirations, AI will remain a promise partially fulfilled. The call to action is not just to innovate—but to prepare, to govern and to lead with purpose.

Autonomous AI agent control

Most organizations aren't ready yet - even the most advanced players, those who have invested heavily in modernization, established robust governance frameworks, embedded ethical safeguards and deployed human-in-the-loop models, are not ready to let go of control.

Because the truth is, no organization today is ready to hand over complete operational control to an autonomous AI agent. AI may assist. It may recommend. But autonomous action, without any human checkpoint, remains a leap too far. Having said that, if true intelligent autonomy is our goal, we can't simply preserve human-in-the-loop approach indefinitely—we must ensure that safeguards remain.

In our own experience, this pattern holds true. AI is widely used in claims fraud detection to identify anomalies and surface suspicious patterns. However, they do not make the final determination. Flagged cases are still referred to experienced claims professionals who understand the nuances of policy language, client history and regulatory context.

The same is true in underwriting. AI can help surface risk indicators and suggest pricing based on historical loss data or exposure trends. But ultimately, it is the licensed underwriter, often working closely with the broker, who weighs those insights against market conditions, client relationships and the judgment developed over years of experience in the field. The technology supports the decision but does not replace the human element.

What we consistently observe is that AI systems today are advisory, not autonomous. They inform human decision-makers but do not replace them. This is not merely a limitation of current technology—it is a deliberate, principle-driven design, rooted in the insurance industry's deep responsibility to manage risk with care, accountability and human oversight.

As we look ahead, this balance must evolve. We cannot indefinitely preserve the human-in-the-loop as a constraint on scale or innovation. Nor can we remove it without confidence in the safeguards, explainability and governance required for responsible autonomy.

This raises a fundamental strategic question: How do we build systems that are autonomous, but remain accountable?

Understanding artificial general intelligence

Artificial general intelligence (AGI) refers to a type of AI that possesses the ability to understand, learn, and apply knowledge across a wide range of tasks—much like a human. Unlike narrow AI, which is trained for specific functions (e.g., fraud detection or language translation),AGI can:

- Generalize knowledge across domains and unfamiliar problems
- Adapt dynamically to new and unpredictable environments
- Reason and make decisions with minimal human guidance
- Learn continuously from experience and feedback
- Exhibit creativity and innovation, going beyond preprogrammed logic

Understanding reinforcement learning

Reinforcement learning (RL) is a branch of machine learning where AI agents learn by interacting with their environment. Instead of learning from static datasets, RL agents receive rewards for correct decisions and penalties for errors—continuously refining behavior through feedback loops.

Two paths forward in solving the autonomy challenge

This introduces an urgent strategic imperative: If autonomy is the goal, how do we get there safely, reliably and with the current constraints? We believe there are two distinct—but complementary—pathways.

The AGI frontier—A long-term bet

AGI represents the pinnacle of machine intelligence—systems capable of applying knowledge across domains, adapting to novel contexts and inventing new ideas. AGI would not just automate tasks; it would innovate. Like Einstein, who looked beyond physics and redefined our understanding of gravity, time and space, AGI systems would synthesize insights across disciplines, unlocking breakthroughs no single dataset could anticipate.

AGI is not about learning from the past—it's about creating new futures. Consider, for example, the COVID-19 pandemic. The early signals were scattered across public health records, climate patterns, mobility trends and genomic data—each meaningful in isolation, but not actionable in aggregate. A system with AGI-like capabilities could, in theory, have connected those signals early, providing foresight where fragmented human systems failed to anticipate the scale or impact. For the insurance industry, such anticipatory intelligence would be transformative. It could enable early-stage risk modeling, dynamic policy repricing, proactive loss mitigation and capital reallocation well before catastrophic events unfold—turning reactive operations into predictive ecosystems.

Yet, we must acknowledge the timeline. Industry leaders like OpenAI, DeepMind and Meta are pouring billions into AGI, but mainstream applicability is likely 10–15 years away.

The insurance sector—bound by present-day regulatory frameworks and fiduciary obligations—cannot afford to wait for a hypothetical leap. So, what now?

The reinforcement-switch model with a pragmatic blueprint

While AGI is a distant moonshot, RL offers a near-term, operational bridge to intelligent autonomy—rooted in feedback, adaptability and context.

Here's how this translates into insurance workflows:

As depicted in Figure 1, the reinforcement framework for autonomous agents progresses through a structured journey from initial learning to delegated autonomy.

1. Model initiation: AI is trained on historical data and deployed in production.
2. Coworker mode: The agent makes predictions; humans validate and act.
3. Reward loop: If the AI's recommendation aligns with the human's decision, it earns a reward. If not, a penalty is applied.
4. Continuous learning: The model is part of every interaction, constantly improving its decision-making ability.
5. Delegated autonomy: Once consistent accuracy is demonstrated (e.g., 100 correct decisions over six months), control is handed over to the AI agent.

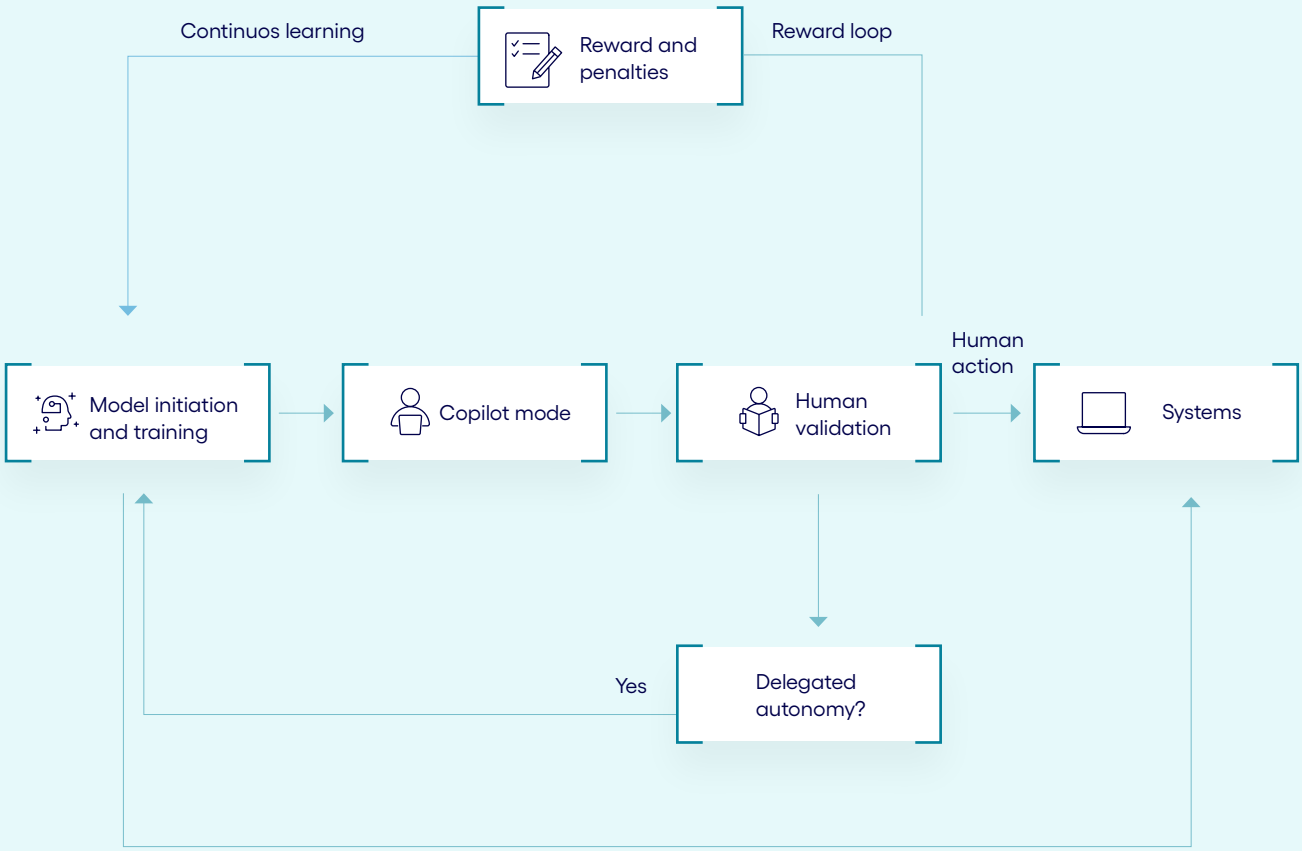

The reinforcement learning loop

Now true autonomy in production environments isn't linear, it's cyclical—even with reinforcement. AI agents operate confidently until something changes: A regulatory framework is updated, a data pipeline is restructured, or the business logic shifts.

And when change hits, most agents—trained on past patterns—fail to adapt in real time. This lack of adaptability introduces significant risk, as the agent may continue executing actions that are no longer contextually valid, potentially leading to critical errors.

This pattern isn't an exception, it's the norm. To manage this rhythm without friction, insurers need an intelligent orchestration layer that enables seamless switching between humans and machines.

That's the essence of the AI-human switch framework—a model engineered for resilience through reversibility.

How the reinforcement-switch model works

This reinforcement-switch model combines reinforcement learning with human-centered control systems to deliver accountable autonomy.

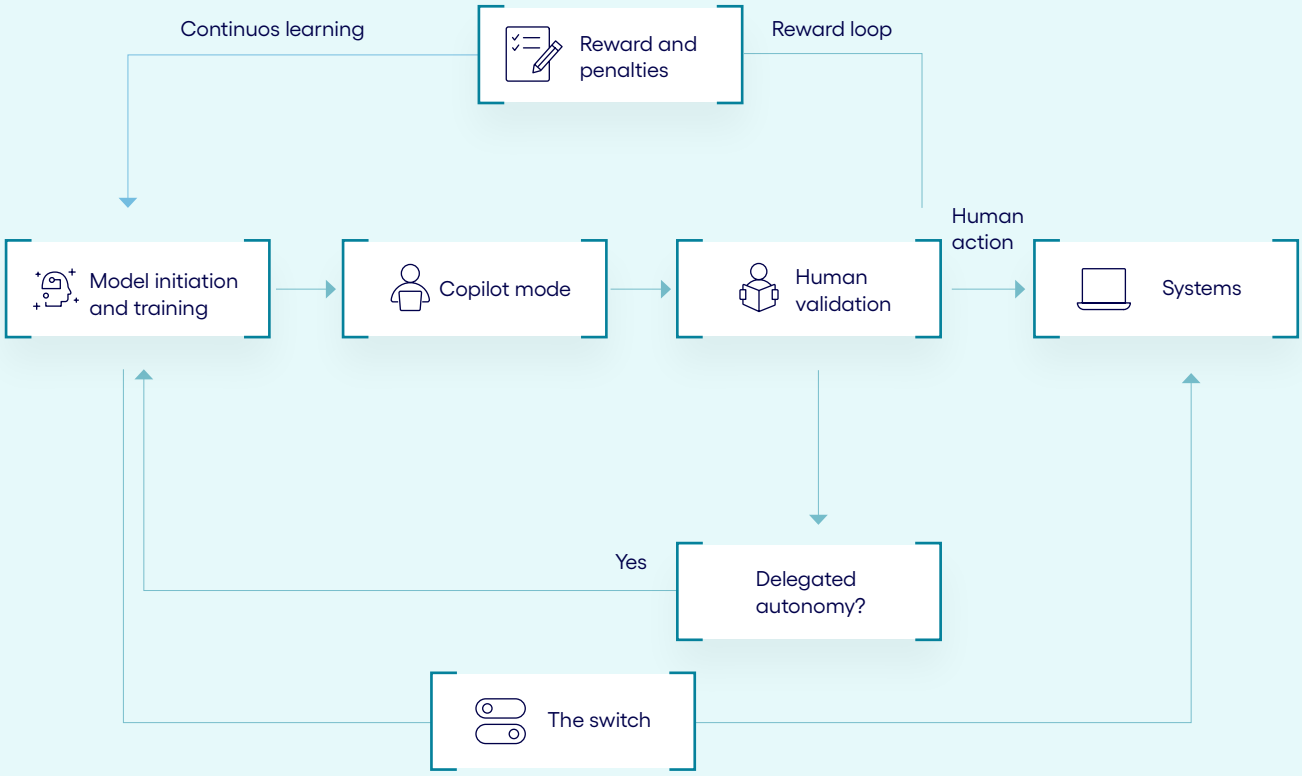

The reinforcement-switch model

The reinforcement-switch model introduces a dynamic, phased approach to human-AI collaboration—designed not as a binary transition to autonomy, but as a negotiated continuum of control. It offers a structured path from human-supervised intelligence to conditional autonomy, anchored in trust, oversight and resilience.

1. Prediction and Reinforcement phase

AI agents begin by generating predictions or actions. Human operators validate outcomes—reinforcing successful decisions or applying corrective feedback. This feedback loop is essential to shape behavioral baselines and establish trust.

2. Confidence accumulation

As agents deliver consistently accurate outcomes—e.g., 100 validated decisions across varied scenarios—they earn graduated autonomy. Trust is not assumed; it is earned through performance.

3. **Delegated autonomy with embedded switch triggers**
   Upon crossing the confidence threshold, agents operate independently, but not indefinitely. Autonomy is bounded by embedded "switch triggers," including:
   - Emergence of unfamiliar data or novel variables
   - Breach of predefined risk thresholds
   - Regulatory anomalies or black swan events
4. **Switchback protocol**
   When a trigger is activated, control reverts automatically to human operators. This is not a fallback—it's a built-in safeguard. Switchbacks are designed into the system architecture to ensure real-time accountability without disrupting operational continuity.
5. **Emergency reclaim protocol**
   In parallel, a manual override allows human agents to forcibly reclaim control when necessary—for example, if an AI agent demonstrates erroneous decision logic. This ensures strategic authority always remains with the organization. Autonomy, in this model, is conditional, not unconditional.
6. **Reinforcement continuum**
   Even after a switchback occurs, the AI agent does not go offline—it continues to operate in a recommendation-only mode, offering insights and suggestions to the human operator. Control remains firmly with the human, but the agent stays engaged in the decision-making process. During this phase, the original prediction and reinforcement phase resumes where every recommendation made by the agent is either rewarded or penalized based on human validation. These continuous feedback loops enable the agent to refine its models, learn from edge cases and gradually rebuild the trust required for future autonomy.

This is not just a safety net—it's a strategic design for fail-forward AI: Systems that evolve through friction, recover through rhythm and strengthen through supervised learning.

The reinforcement-switch model is purpose-built for environments where the cost of imperfection is acceptable and the value of continuous learning compounds over time. It shines in medium- to low-criticality and high-frequency use cases, the kind where intelligent autonomy can drive real-world impact without compromising trust.

Consider AI agents trained to detect early signs of water leakage. These agents can proactively shut off valves before minor issues escalate into major losses. Already, forward-thinking insurers are subsidizing such technologies as part of smart home ecosystems—not just to reduce claims, but to reposition themselves as proactive risk partners. And if the AI agent makes an incorrect decision, the impact is minimal: Control can be immediately returned to human operators, the learning loop continues and the system is strengthened over time.

In such scenarios, autonomy is no longer a liability—it becomes a strategic advantage, one that can be dialed up or down based on risk, context and maturity.

While AGI remains on the horizon, autonomous agents today must operate within clear, governed boundaries, guided by human intent and institutional accountability. The goal isn't to replace human control—it's to scale it with intelligence, adaptability and trust.

The reinforcement-switch model offers a pragmatic blueprint for making that future both achievable and safe—but only if we act decisively. The insurance industry cannot afford to wait for AGI to arrive before pursuing intelligent autonomy. The window of opportunity to modernize infrastructure, embed governance and pilot controlled autonomy is rapidly closing. By adopting a reinforcement-driven, switch-enabled approach today, organizations can build trust, scale innovation and create the operational resilience needed to thrive in an unpredictable future. The choice is clear: Start engineering accountable autonomy now, or risk being left behind as the market evolves rapidly.

References:

BCG (2025): How Insurers Can Supercharge Their Strategy with AI
How insurers can supercharge strategy with artificial intelligence | BCG

Predicting AI trends in insurance for 2025
Predictions 2025: Can AI Deliver on its promises For insurance?

Gartner (2025): The 2025 Hype Cycle for Artificial Intelligence Goes Beyond GenAI
The Hype cycle for Artificial Intelligence 2025 | Gartner

Stanford HAI (2024): The Road to AGI: Benchmarks and Barriers
The 2024 AI Index Report | Stanford HAI

cognizant®

Cognizant helps engineer modern businesses by helping to modernize technology, reimagine processes and transform experiences so they can stay ahead in our fast-changing world. To see how Cognizant is improving everyday life, visit them at www.cognizant.com or across their socials @cognizant.

World Headquarters
300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Fax: +1 201 801 0243
Toll Free: +1 888 937 3277

European Headquarters
280 Bishopsgate
London
EC2M 4RB
England
Tel: +44 (01) 020 7297 7600

India Operations Headquarters
5/535, Okkiam Thoraipakkam,
Old Mahabalipuram Road,
Chennai 600 096
Tel: 1-800-208-6999
Fax: +91 (01) 44 4209 6060

APAC Headquarters
1 Fusionopolis Link, Level 5
NEXUS@One-North, North Tower
Singapore 138542
Tel: +65 6812 4000

WF3713550