TATA CONSULTANCY SERVICES

TATA

An AI-First Insurance Blueprint:

From pilot sprawl to compounding advantage: designing insurance for the next era

Preface

Many insurers today are deploying AI in meaningful ways—extracting data from broker submissions, prioritizing risks, surfacing underwriting insights, and automating repetitive tasks. These efforts accelerate existing workflows but leave the underlying operating model intact. The business logic, governance structures, and systems architecture were built for rules, not intelligence.

This article argues that the **AI-first model is materially different from AI-enabled** and that failing to account for that difference early on can create long-term constraints.

An AI-first insurer doesn't just layer AI onto legacy processes. It changes how decisions are structured, how intelligence flows across functions, and how systems interact dynamically. These changes require a shift in the underlying architecture—toward platforms that support adaptability, real-time context, and coordinated action across multiple agents and domains.

The article explores:

- The **AI-first operating model,** in contrast with AI-enabled—and why insurers need to design for it with intention, **starting now** not down the road
- The **foundational technology architecture** required to support learning, context-aware, and orchestrated systems
- The **transformation path,** with practical guidance on where to start and which capabilities to prioritize so the foundations for AI-first are in place—even if use cases begin small

AI-first won't emerge on its own. Without deliberate design choices, today's architectures may entrench patterns that limit tomorrow's potential.

Content

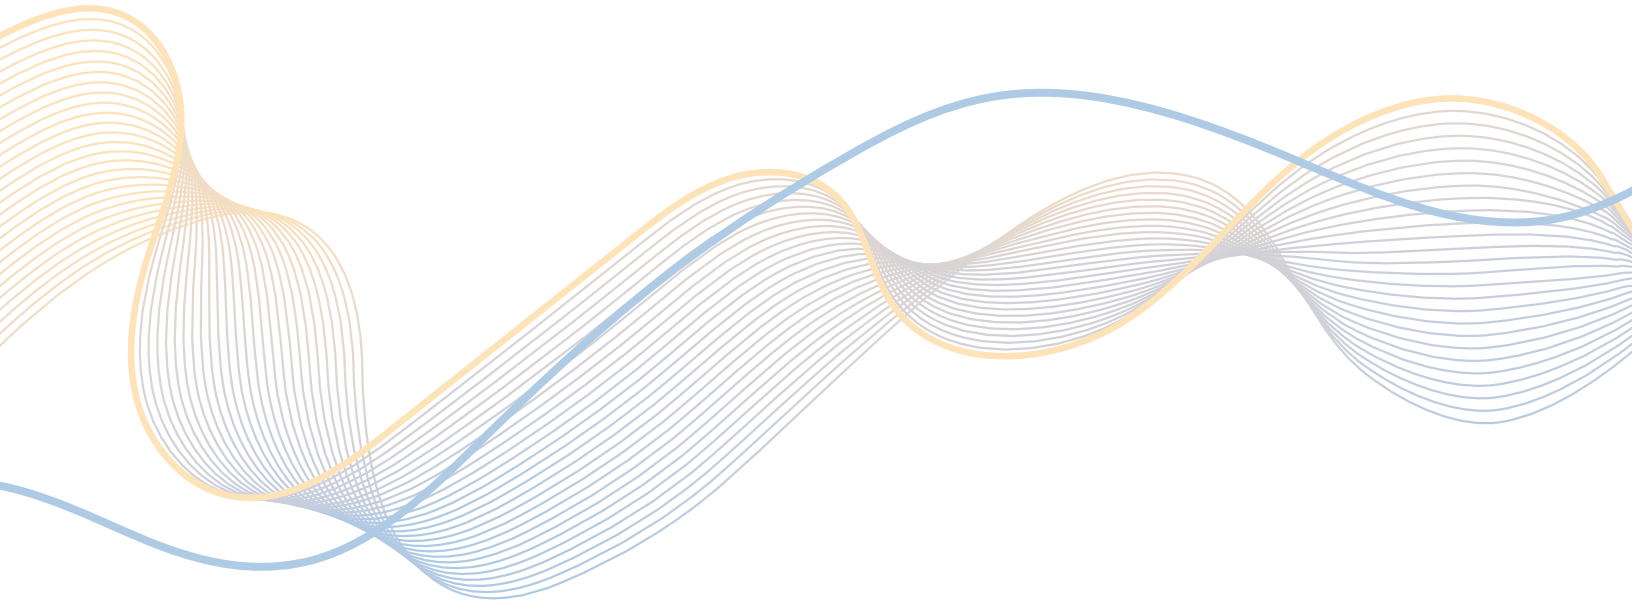

I. Why Insurers Need to Design for AI-First Now

I. Why Insurers Need to Design for AI-First Now

Painting the Picture of AI-First Insurance

When a new commercial submission arrives from a broker, most insurers handle it through a familiar sequence: intake, manual data entry, underwriter review, pricing, and approval—each step passing through different systems, often with delays. An AI-enabled approach might automate one or two of these steps, like extracting data from broker emails or flagging high-risk submissions.

In an AI-first model, as soon as a submission arrives, multiple intelligent AI agents launch in parallel:

- Data ingestion from submission portals, broker emails, internal records, and external databases
- Triage based on completeness and complexity
- Simultaneous scoring of appetite fit, win likelihood, and portfolio exposure
- Fraud detection runs alongside these processes rather than after bind, scanning for anomalies before quotes are even generated (see figure 1)

Submissions are then routed intelligently with straightforward cases going to underwriters with the right authority, complex or high-risk ones escalated with comprehensive summaries, and incomplete ones triggering automated broker follow-up. Post-bind, pricing models update portfolio exposure dynamically and fraud checks continue to catch subtle patterns.

Instead of disconnected handoffs, the process becomes an intelligent, self-reinforcing system where every insight informs future decisions in real time, creating enterprise-wide intelligence rather than departmental automation. This is the key difference between AI-first and AI-enabled operations.

Figure 1. Illustrative AI-Driven Submission Workflow

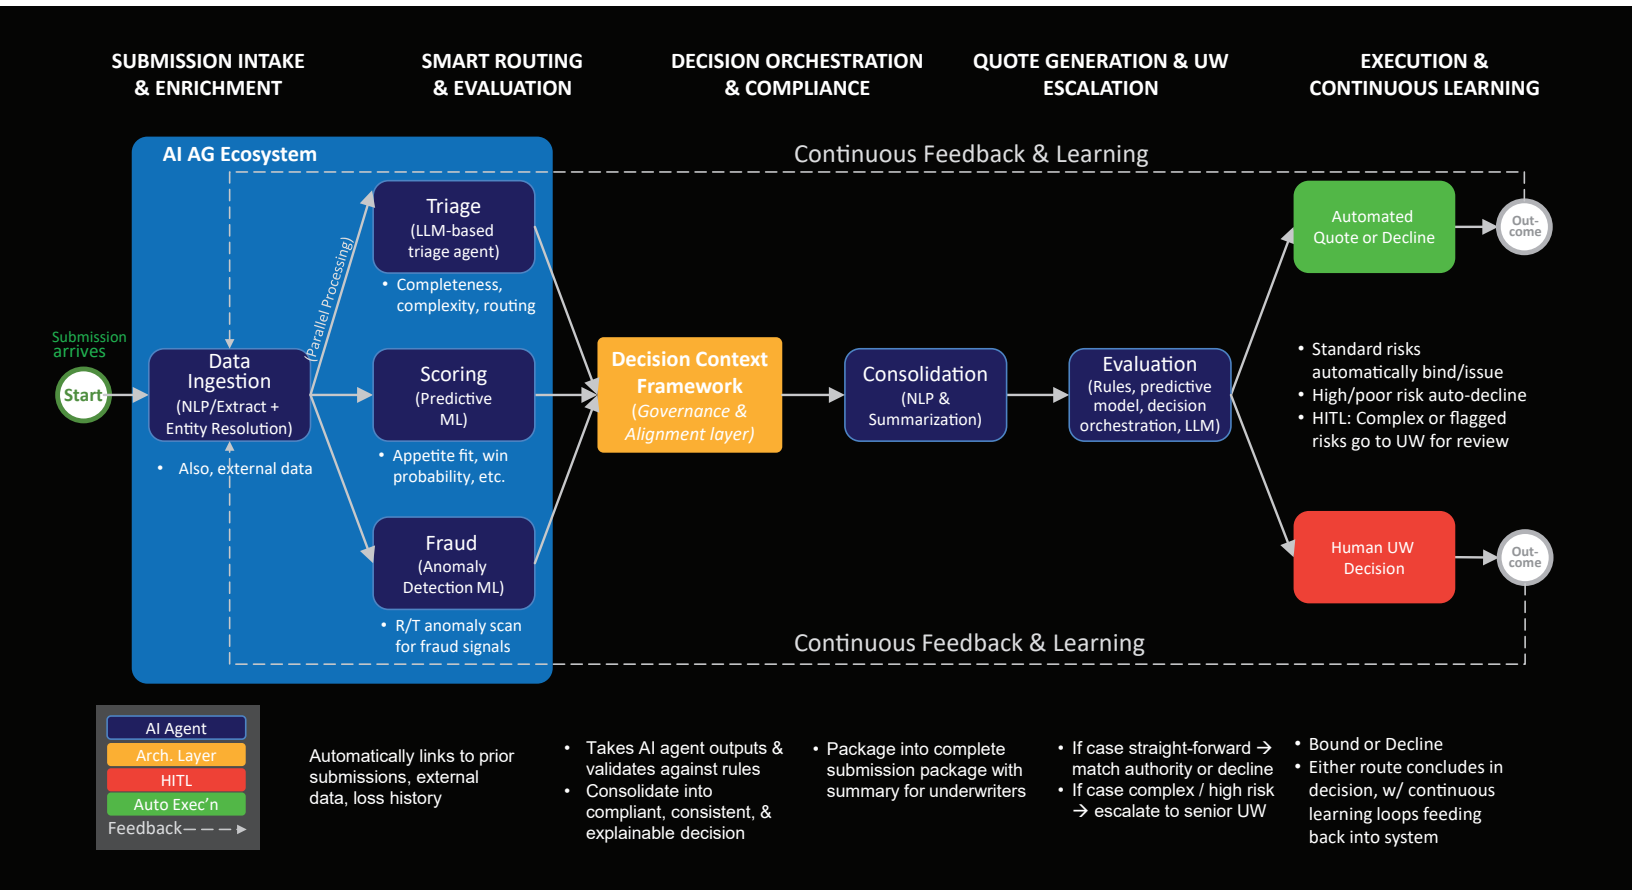

It is important to note that AI-first does not mean complete automation that supplants human decision-making. Rather, it embeds human-in-the-loop (HITL) oversight as a foundational design principle (see figure 2). While intelligent automation drives efficiency and enhances customer experience, HITL guardrails ensure model integrity by incorporating human judgment where it is most valuable. It is a deliberate design choice, not a manual workaround for technology limitations as in the past.

Figure 2. AI-First Design Principles

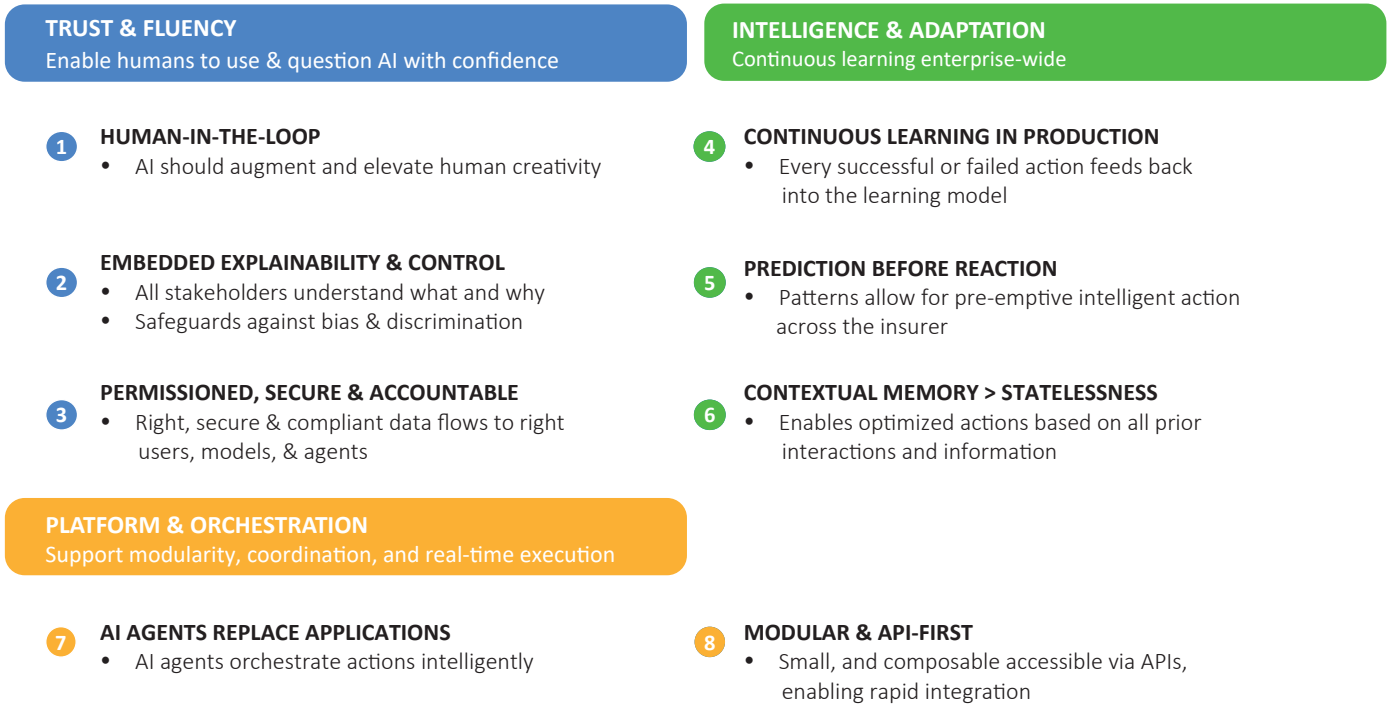

The AI Maturity Progression: Where Are You Today?

Understanding your current position helps chart the path to AI-first operations. The maturity map below illustrates the transformation journey that lies ahead for insurers as their adoption of AI evolves (see figure 3).

Most insurers today operate between Experimental and Capable stages, with a handful of leaders transitioning into early Operational. AIG reported that its Underwriter Assistance Generative AI platform now reviews 100% of private company and non-profit financial lines submissions without adding underwriters , which will fuel significant growth in new E&S business by 2030. Allianz UK has introduced BRIAN, a generative AI tool that helps underwriters quickly search and retrieve information from lengthy documents. Within nine months of launch, it reportedly is saving an estimated 135 working days by answering over 13,000 questions in seconds rather than hours. These are examples of leveraging AI to amplify the bandwidth and efficiency of skilled workers with the potential to achieve a step-change improvement in business performance through lower expense and loss ratios, stronger submission flow and higher hit rates.

Such early achievements show what ambitious AI deployment looks like within a business function. As carriers extend this level of AI integration across all business functions—connecting underwriting insights with claims, pricing, and customer management—they lay the foundation for an adaptive, continuously learning operating model: an AI-first insurer. This shift echoes the broader vision of the

Perpetually Adaptive Enterprise described by TCS , where modularity, intelligence, and orchestration enable organizations to evolve dynamically in response to internal signals and market forces.

Figure 3. AI Maturity Model

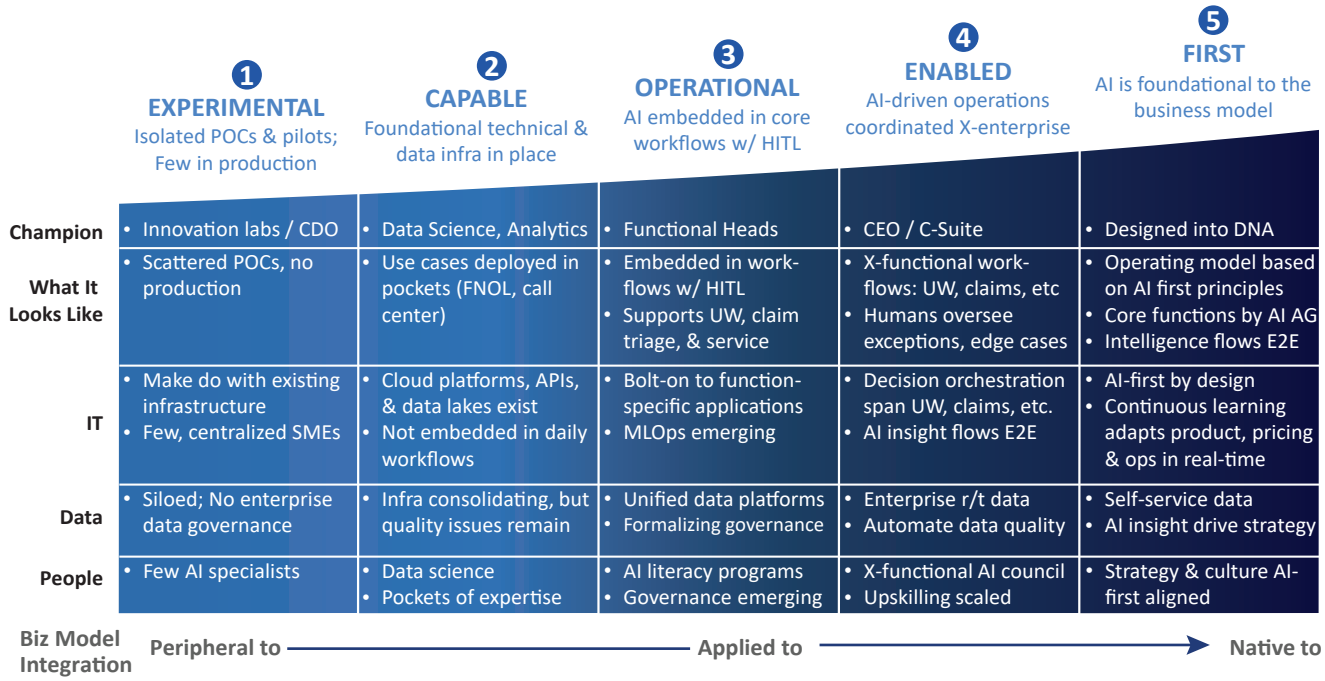

The Trap of 'Good Enough' AI

In the late 1990s, many dismissed e-commerce by focusing on its early shortcomings—fraud risks, clunky user experiences, the hassle of returns. The idea of immersive "virtual malls" felt far-fetched, especially when online sales volumes were minimal. It seemed inconceivable that shoppers would abandon the tactile, social experience of browsing books at Barnes & Noble for a faceless website called Amazon.

The big mistake of underestimating e-commerce stemmed from a subtler one. Namely, underestimating how better experiences led to incremental shifts in customer behavior that compounded into structural change. People didn't start by buying everything online—they started with books. When Amazon made that easier, faster, trustworthy, and more convenient, consumer expectations shifted, and online retail took off.

Today, insurers face a similar inflection point. There is a tendency to view AI-first as a future ambition while focusing on AI-enabled efficiencies that feel more practical and achievable today. This underestimates how quickly applied AI is advancing and already capable of reshaping underwriting, claims, and distribution. Stanford's Massive Multitask Language Understanding (MMLU) benchmark already shows AI systems matching or exceeding average human performance on many professional tasks in medicine, law, and finance. That's more than enough for insurance workflows where speed, recall, and consistency are paramount.

At the same time, carriers face mounting structural changes to the industry: climate volatility, cyber risk, social inflation, and shifts in reinsurance capacity. Recent profitability has leaned on the usual levers—tightening underwriting, raising rates, absorbing more risk. These are business-as-usual tactics deployed to fix a model under pressure. AI-enabled improvements can extend these levers, but they still may not be enough going forward.

That's why insurers must begin rethinking the model, not just adopt new tools. Building for AI-enabled efficiencies alone risks unwittingly embedding constraints that will be costly to unwind. Designing for AI-first should better position insurers to adapt to volatility, deliver better value to customers, and establish sustainable competitive advantage.

The Flywheel Effect

The flywheel effect illustrates what makes AI-first different: reinforcing loops that create advantages which compound over time and become increasingly difficult for competitors to match.

Claims patterns automatically inform underwriting appetite adjustments. Shifts in underwriting appetite guide how brokers target and qualify risks. Customer interactions update risk profiles in real-time across all touchpoints and make subsequent interactions smarter and feed directly into product and service design. Market responses to product changes trigger immediate pricing and coverage optimizations. Each learning loop is valuable on its own, but when connected they form a system that keeps learning and improving across the enterprise, creating an impact that is more than the sum of its parts. This elevates AI from merely an efficiency play to a growth and resilience strategy, creating the possibility of countering the structural pressures noted above.

In the short run, AI-enabled and AI-first carriers can look alike. But as underlying models advance—benchmarks like MMLU showing gains in months rather than years—the compounding loops widen the gap. Those that design for AI-first will be better positioned to convert rapid model progress into enterprise advantage. Those that don't may find themselves limited by fixes that can't scale or age quickly and increasingly unable to close the gap with competitors.

The next step is to explore what this model could look like in practice, through the architecture that makes AI-first possible.

II. AI-first Architecture and the Challenges of Building It

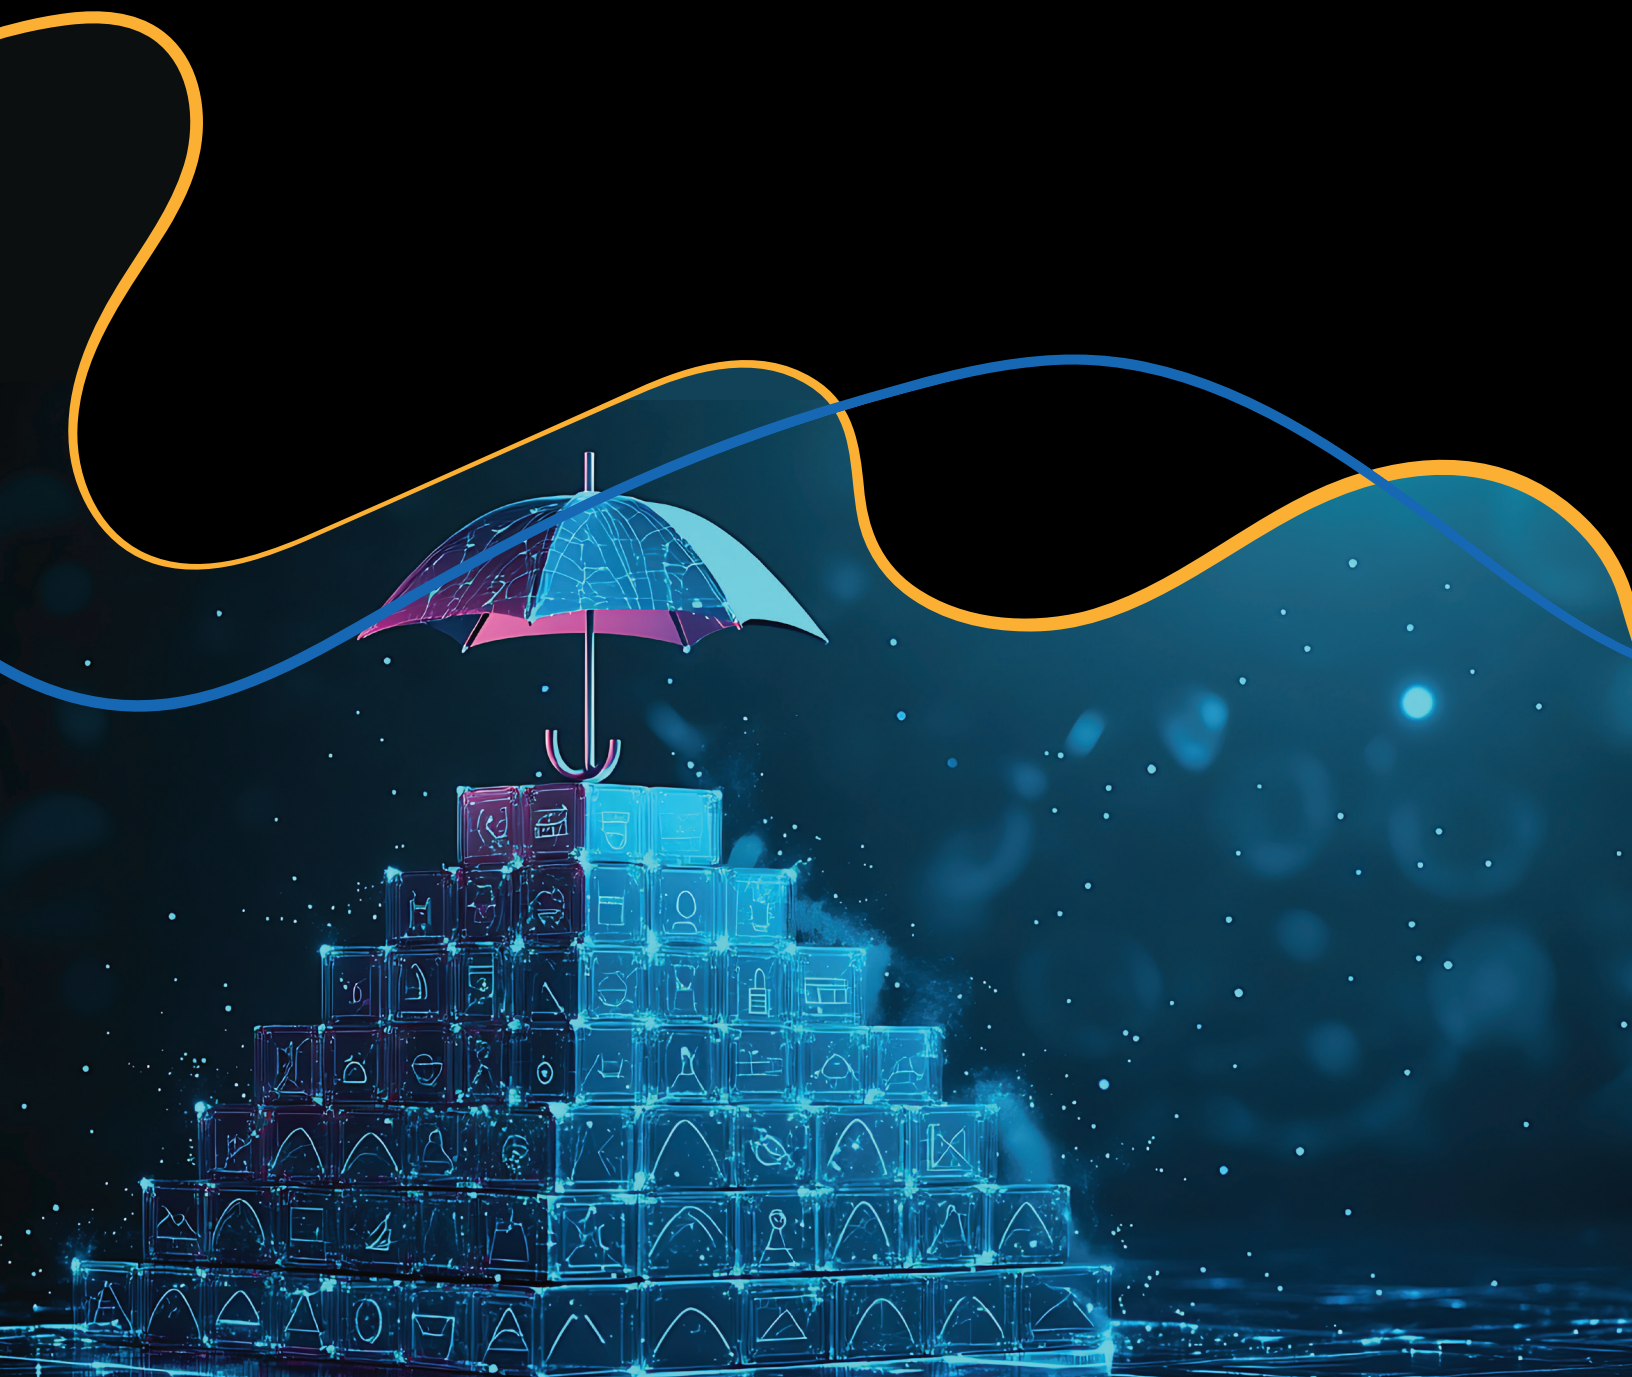

II. AI-first Architecture & the Challenges of Building It

From Value Chains to Living Systems

To better grasp AI-first Insurance, a new mental model may be helpful. The notion of insurance operations as a value chain is deeply ingrained in industry practitioners: marketing leads to application, then underwriting, policy issuance, billing, and eventually claims. Each step hands off to the next, like an assembly line.

AI-first operations work more like biological systems, where multiple organs work simultaneously and share information continuously. Just as your circulatory system carries nutrients to every part of your body while your nervous system coordinates responses, an AI-first insurer has information flowing everywhere while intelligent systems coordinate decisions across all functions. This biological metaphor helps explain why AI-first is profoundly different from AI-enabled: instead of automating individual steps in a chain, you're building an interconnected system where every component strengthens the others. (See sidebar)

MAPPING ANATOMY TO AI-FIRST

Circulatory System → Data
- Pumps clean information to every function

Neural System → Decision Context
- Processes signals and trade-offs for coordinated choices

Nervous System → Orchestration
- Routes tasks between specialized AI agents

Connective Tissue → Integration
- Links legacy systems, reinsurers, and partners

Skeletal System → Governance
- Provides structure and ethical boundaries

Immune System → Feedback
- Detects threats (drift, bias) and strengthens the system

The AI-first Reference Architecture

Building an AI-first foundation is a complex undertaking and no single architecture will fit every organization. The core design principles introduced in the previous section provide the lens for understanding the key layers of the architecture described here (see figure 4). This reference model is not intended as a one-size-fits-all blueprint, but rather as a guide to help leaders see how the layers fit together and anticipate the implementation challenges of each.

1. Data & Integration. This layer aggregates, cleanses, and streams data from all sources (policy systems, claims platforms, external data providers, customer interactions) into a single, real-time, accessible repository to ensure every AI component works with the same current information. Unlike a traditional Operational Data Store (ODS), which mainly consolidates copies of transactional data for reporting and light decision support, this repository is designed as a real-time data hub. And unlike most data warehouses or lakes, which support analytics and historical reporting, it powers live AI workflows, delivering API-ready data to agents and applications and capturing outcomes back into the system. Challenge: Data resides in systems that were never designed to work together, making the principle of access-based, secure data flows particularly challenging. Often the bottleneck isn't building the models but stitching together fragmented, low-quality data.

2. Intelligence. The intelligence layer hosts the AI models that power predictions, classifications, and insights across business functions. These models encode business rules, regulatory requirements, risk appetite, and portfolio constraints in machine-readable form, and must be constantly tuned to avoid drift. This is the "brain" of the architecture, but does not operate alone.

To remain relevant and accurate, these models depend on the Context & Knowledge layer, which supplies contextual information and memory. Here, technologies such as knowledge graphs, retrieval-augmented

generation (RAG), and the Model Context Protocol (MCP) apply the right context into each model's decision process. Together, Intelligence and Context & Knowledge ensure that AI decisions are both consistent and contextually relevant.

Challenge: Decision rules are the same business and underwriting rules that drive eligibility, pricing, and approvals. However, in most carriers they're not only scattered across manuals and institutional memory, they're embedded in ancient systems and outdated languages (e.g., COBOL in legacy PAS). Extracting, converting, and standardizing those hard-coded rules into a modern, machine-readable framework can be as complex as building models. But without this step, AI decisions lack consistency and auditability.

3. AI Agent Orchestration. This is the coordination layer where specialized AI agents (underwriting, claims, fraud detection, customer service) communicate and share context. Rather than isolated tools, these agents work together through an orchestration manager that routes information and coordinates decisions.

Beyond interactions between agents, this layer also defines how agents and humans interact. For example, if a fraud detection agent detects an anomaly while the pricing agent produces a competitive quote, the orchestration logic determines whether to escalate to an underwriter, request more data, or take some other action. By design, human judgment—rather than being sidelined—is brought into the loop where it creates the most value.

These orchestrated interactions take place in the Experience layer, where well-designed interfaces help humans engage with AI—consuming the right information, in the right context, at the right time—and act on it effectively.

Challenge: Without agent-level orchestration, insurers end up with fragmented tools. This is why many AI efforts manifest as "pilot sprawl" — isolated proofs of concept that don't evolve into tangible advantages for the enterprise. Orchestration turns them into a cohesive system.

4. Integration Infrastructure. The connectivity layer that links AI components to existing policy administration systems, external data sources, regulatory reporting tools, and partner platforms. This layer connects modern AI capabilities to legacy insurance systems and makes new AI capabilities usable in practice.

Challenge: Many insurers still rely on COBOL-era batch systems or SOAP APIs. Wrapping these with modern API orchestration platforms (MuleSoft, Apigee, etc.), or using middleware and staged migration, is often required to bridge these legacy assets with AI-first layers.

5. Security. This protection layer ensures compliance, risk management, and resilience across the architecture. It safeguards data flows, model outputs, and decision processes from breaches or manipulation, providing the foundation for trust without which no system can operate.

Security also connects downward into the Infrastructure layer. Without hardened compute, storage, and networking, protective measures cannot scale. This layer also extends upward into Governance & Continuous Learning, since secure audit trails are critical for explainability and regulatory compliance. Challenge: AI introduces new vulnerabilities beyond traditional IT. Systems can be tricked by manipulated inputs, probed to expose sensitive training data, or fall behind rapidly changing regulations. Without robust security, even well-designed AI solutions face unacceptable operational and reputational risks.

6. Governance & Continuous Learning. This oversight/observability layer monitors how models and agents perform over time. It records decisions, confidence levels, pass/fail or exception rates, then links those outcomes to business results and human overrides. By creating audit trails that show who made which decision, when, and on what basis, this information ensures explainability, bias checks, and accountability while also feedback into retraining and rule updates.

Governance and continuous learning are paired because they are inseparable: monitoring outcomes provides the feedback loops needed to retrain and update models. Every output becomes an input that improves the operation. There is also a direct tie to the Experience layer, where user feedback is captured, and to the Infrastructure layer, which provides scalable pipelines for retraining and redeployment. Collectively, these loops keep the system adaptive.

Challenge: Without structured oversight and learning processes, enterprises have no reliable way to detect model drift, explain decisions, or build trust over time. Operationalizing this layer requires building governance and learning processes that are rigorous enough to satisfy regulators and risk managers, but lightweight enough that the organization actually adopt them. Without this balance, governance becomes a bottleneck or a box-checking exercise.

To revisit the living organism metaphor, these layers work together like biological systems: the data layer circulates information like a bloodstream, the decision framework coordinates choices like a brain, orchestration routes tasks like a nervous system, integration connects to external systems like sensory organs, governance maintains accountability like a skeleton, and continuous learning adapts and improves like an immune system.

Figure 4. AI-first Reference Architecture

LAYER	*ROLE*	*WHY IT MATTERS*
EXPERIENCE	User interfaces where humans interact with AI	Your AI is only as good as the experience it delivers
ORCHESTRATION	AI agents that coordinate tasks and make decisions	Intelligent automation that knows when to escalate to humans
INTELLIGENCE	The AI models that power predictions and insights	Brain of your system—needs constant care & improvement so models continue to perform
CONTEXT & KNOWLEDGE	Contextual information & memory for AI	Gives AI the context it needs to be accurate and relevant
DATA & INTEGRATION	Pipelines that feed clean data from all sources	Garbage in, garbage out—quality data is everything—for trust and accessibility
GOVERNANCE & CONTINUOUS LEARNING	Monitoring and governance for AI operations	You can't improve if you can't measure how your AI is performing
SECURITY	Protects data, models, & decision flows from manipulation, breaches, & misuse.	Even the most advanced AI systems cannot operate safely or at scale without this foundation.
INFRASTRUCTURE	The compute, storage, and networking foundation	Powers everything above it—must be scalable and cost-effective

It's As Much About the Talent As the Tech

Building an AI-first insurer isn't just a technology challenge, it also requires creating an organization that thinks and acts as a connected system. In an AI-first insurer, a claims tweak can influence underwriting appetite, capital allocation, or even product strategy overnight. Traditional hierarchies—with departmental boundaries, approval ladders, and sequential hand-offs—are designed to minimize interdependence, not embrace it.

Early movers are treating AI implementation as a cultural catalyst by focusing on:

- **Cross-functional ownership:** AI systems span underwriting, claims, distribution, compliance, and IT. A single accountable owner—what Gartner terms a "Fusion team lead"—must own end-to-end outcomes, without which the risk of misaligned incentives or fragmented architecture becomes heightened.

- **Technical literacy for business users:** Product managers and underwriters don't need to become data scientists, but they must understand enough about model behavior and feedback loops to spot anomalies and propose fixes.

- **Risk-aware agility:** AI will make mistakes. Leading carriers establish rapid detection, correction, and learning processes, treating failures as opportunities for model improvement rather than assigning blame.

Siloed teams are antithetical to a business model and IT architecture premised on connecting intelligence across the enterprise. It's easy to imagine disjointed teams pulling in different directions, preventing a thoughtful AI-first architecture from delivering its full potential. Given that technologies are evolving rapidly, a collaborative, cohesive organization will be better equipped to adapt and solve challenges, even with a moderately sophisticated architecture.

The AI-first architecture highlights the layers that make continuous learning possible—data delivered with the right context at the right time, intelligence that adapts, orchestration that connects decisions, and governance that sustains trust. For insurers, the challenge is less about understanding the blueprint and more about choosing the right starting point.

The next section explores practical entry points and how insurers can take those first steps toward becoming AI-first.

III. From Vision to First Steps

III. From Vision to First Steps

Most insurers operate across a patchwork of technology and cultural maturity. A cloud-native, SaaS front-end may sit in front of a decades-old mainframe with a hybrid data hub in the middle. The policy administration landscape may be composed of modern vendor solutions and outdated in-house platforms depending on the line of business. This uneven mix makes the choice of where to begin critical. Legacy environments often have deeply ingrained processes and risk-averse cultures that slow adoption, while modern platforms may face competing priorities that dilute focus. With these trade-offs in mind, insurers tend to fall into three broad archetypes:

- **Modern cloud-native insurers** (mostly MGAs and insurtechs): Can integrate AI capabilities directly through APIs and scale quickly.
- **Hybrid cloud carriers:** Need adapter layers between new AI tools and legacy cores, but can move incrementally.
- **Legacy mainframe operations:** Require event streaming or staged data migration before AI integration becomes feasible.

These archetypes are directional, not definitive. In practice, the level of maturity will differ across entity, business line, and functions, making it essential to begin where the foundation can support progress and build momentum.

Promising Starting Points

Understanding where your organization sits on the maturity continuum is only the initial step. The next step is to identify which use cases can serve as effective starting points. The most promising ones (1) deliver tangible business value and (2) reinforce key AI-first capabilities. Based on early industry pilots and implementations, a few starting points offer a favorable balance of impact and feasibility.

1. Intelligent Submission Triage

- **What it does:** AI scores incoming submissions for routing to appropriate underwriters or automated processing
- **Why it works:** Clear routing decisions with measurable impact on cycle times
- **What you need:** Submission data access, risk scoring models, workflow integration
- **What success looks like:** Meaningful portion of submissions qualify for streamlined processing with substantial cycle time reduction

2. Document Intelligence and Data Extraction

- **What it does:** AI extracts structured data from broker submissions, applications, certificates, and supporting documents
- **Why it works:** High-volume, repetitive task with clear accuracy metrics applicable across all lines of business
- **What you need:** Document management integration, OCR capabilities, quality assurance workflows
- **What success looks like:** Majority of documents processed without human review with substantial improvement in processing speed

3. Risk Pattern Detection and Fraud Prevention

- **What it does:** AI identifies anomalies and suspicious patterns across submissions, claims, and policyholder behavior

- **Why it works:** Scales human expertise to detect subtle patterns across large volumes
- **What you need:** Historical data, pattern recognition models, investigation workflow integration
- **What success looks like:** Earlier detection of problematic risks with reduced investigation overhead

Potential Challenges

Scope creep: Starting with "simple" underwriting automation that expands to include pricing, portfolio management, and regulatory reporting. The result is diluted focus and mediocre solutions for a multitude of challenges. Pick one workflow, mature the relevant AI agents to execute it well before adding orchestration / agentic AI capabilities.

Data quality assumptions: A common pitfall is assuming existing data is "good enough" for AI. In reality, inconsistent or incomplete data can derail projects and extend timelines. Projects plans should factor in adequate time (very roughly 30-40%) for data cleansing, validation and standardization.

Change management underestimation: Technical teams may build strong models that business users don't trust or adopt. Involve business stakeholders in AI agent development and testing, not just requirements gathering.

Vendor promises: AI vendors often demo capabilities on clean, prepared datasets or sandboxes that don't reflect real conditions. The risk is disappointment and minimal ROI when solutions are launched with imperfect production data. Insist on proof-of-concepts using your actual data and edge cases.

Phases of Capability Building & Big Rocks: Building in the Right Order

Momentum in an AI-first transformation comes not only from choosing good entry points but also from sequencing capability building and foundational investments together. Capabilities mature over phases; foundational investments are the "big rocks" that must be in place so later initiatives can scale. Some activities will start in one phase and continue into the next, but this framework highlights when each capability should reach maturity.

Phase 1: Foundation – Establish the Essentials

Goal Lay down the core capabilities that make AI-first adoption possible.

Big Rocks
- Event streaming & data pipelines
- Data quality & integration processes
- Baseline governance & security controls
- Basic monitoring & logging

Why? Without these elements, early AI use cases risk becoming isolated pilots that cannot scale or be trusted. This phase also creates working AI capability that demonstrates value while building reusable infrastructure.

Phase 2: Connected Intelligence – Orchestrate Across Functions

Goal Move beyond isolated AI models to orchestrated agents and workflows across functions.

Big Rocks
- Agent orchestration layer
- Cross-functional data flows
- Enhanced governance & security
- Intermediate monitoring & audit trails

Why? Ensures intelligence flows end-to-end within and across functions, unlocking multi-use-case value.

Phase 3: Optimize at the Enterprise Level

Goal	Expand AI-first practices across the enterprise and embed continuous improvement.
Big Rocks	• Enterprise-wide observability & monitoring • Advanced orchestration across domains • Resilient governance & compliance • Continuous optimization loops
Why?	AI-first becomes an operating model that continuously learns, adapts, and scales securely. Faster time-to-value for new capabilities, enterprise-wide optimization.

This two-dimensional approach—big rocks capabilities and implementation phases—ensures you build the right things in the right sequence for your specific situation.

Investment Planning Framework

Rather than specific budget figures, consider these cost factors that vary significantly by carrier:

- Infrastructure complexity multiplier: Legacy systems without APIs require substantially more integration work than modern platforms. Your existing architecture is the primary cost driver.
- Data readiness factor: Clean, accessible data accelerates projects. Poor data quality can double or triple implementation timelines as teams spend time on data cleansing rather than AI development.
- Organizational change requirements: Carriers with strong cross-functional collaboration and change management capabilities move faster. Those requiring significant cultural shifts need longer timelines and more training investment.
- Regulatory environment: States with established AI guidelines and supportive examiners reduce compliance overhead compared to jurisdictions still developing their approach.
- These factors interact in complex ways, making generic cost estimates misleading. A thorough technical assessment and pilot project provide more reliable planning data than industry benchmarks.

Measuring What Matters

Traditional metrics tell you how well the machine runs. In an AI-first insurer, what matters more is how fast the system learns, how precisely it acts, and how effectively it adapts. To avoid defaulting to superficial measures—like how many models are deployed or how many workflows are automated—carriers should focus on outcomes across five dimensions:

1. Business Value. Are AI-driven decisions actually improving core business performance? Examples:

- Underwriting margin uplift from AI-assisted submissions
- Increased quote-to-bind conversion driven by triage precision or faster turnaround

2. Learning & Adaptation. Is the enterprise getting smarter with each new signal, submission, or claim? Examples:

- Time from new signal to model update or rating logic change
- Percentage of AI workflows that self-adjust based on observed outcomes

3. Decision Quality. Are AI decisions calibrated, contextual, and better than manual benchmarks? Examples:

- Loss ratio improvement on risks selected by AI agents
- Underwriter override rate—and how often the override was more accurate

4. Operational Responsiveness. Is the system reacting faster and more intelligently at scale? Examples:

- Reduction in cycle time from submission to bind
- Percentage of submissions autonomously triaged or routed by agents

5. Trust & Governance. Are decisions transparent, fair, and audit-ready?
- Share of decisions with embedded explainability artifacts reviewed by users
- Number of detected model drift or bias incidents resolved within governance SLAs

Ultimately, "measuring what matters" means measuring how much smarter, faster, and more adaptive your enterprise is becoming rather than how much AI you've deployed.

Regulatory Considerations

AI-first insurers face a new kind of regulatory challenge. These systems make decisions faster, operate across functions, and adapt over time. Traditional, static governance frameworks are built for rules-based systems that change slowly and episodically. They are not well equipped to oversee the more dynamic AI-first operation.

Mistakes like model bias can spread rapidly. Compliance, therefore, must be built into how models are designed, trained, and deployed, with safeguards for fairness, transparency, privacy, and accountability. The framework below outlines the safeguards that AI-first insurers need to meet regulatory expectations. It is not exhaustive, but provides a practical framework for understanding the key considerations.

1. Explainability & Auditability. Can decisions made by AI systems be clearly understood, traced, and explained?
- Build models and agents that produce not only decisions but decision rationales.
- Maintain auditable logs for agent actions, overrides, and feedback loops.

Example: For each quote generated by an AI agent, retain a transparent summary of inputs, risk scores, and decision paths.

2. Fairness & Bias Mitigation. Are AI systems producing equitable outcomes across demographics, geographies, and customer segments?
- Monitor for disparate impact across protected classes—even unintentionally.
- Use bias detection tooling in model training and in production environments.

Example: Ensure that property insurance pricing models don't systematically penalize zip codes correlated with race or income.

3. Data Rights & Consent. Is data collected, used, and shared in ways that comply with privacy and consent rules?
- Classify all inputs into regulatory-grade data categories (e.g., consented vs. inferred vs. public).
- Respect jurisdictional boundaries: a claim image used in Texas may need different handling than in the EU or Quebec.

Example: NLP models trained on broker-submitted emails must ensure no unauthorized use of sensitive customer information.

4. Autonomy & Oversight. Are there appropriate controls over what AI is allowed to decide—and when humans intervene?
- Define decision thresholds for full autonomy vs. human-in-the-loop vs. human-in-command.
- Ensure that business rules, not just models, enforce governance logic.

Example: An AI agent may approve low risk renewals autonomously, but must escalate new risks or exceptions.

Addressing explainability, fairness, and control as core design principles rather than afterthoughts reduces regulatory risk and makes AI systems more adaptable over time.

Conclusion: Compounding Advantage by Design

AI-enabled insurers can make today's processes faster, but efficiency gains alone don't reshape how the enterprise learns. AI-first design goes further: it creates systems where intelligence flows across underwriting, claims, distribution, and service—so each new capability reinforces the next.

Because these approaches are still in their early days, no one can say for certain how durable the lead of early movers will be. But there are strong reasons to believe that feedback loops, enriched data, and team cohesion will compound over time, creating advantages that late adopters may struggle to match. Insurers that begin designing for AI-first now give themselves the best chance not just to keep pace, but to shape the industry's direction.

Authors

Howard Kahai Pang
TCS Americas Consulting Partner – Insurance

Vince Ofrecio
TCS Americas Managing Partner – Insurance

Lino Vivas
TCS Americas Consulting Partner – Insurance

i Susanne Sclafane, "AIG: Turning One Human Underwriter Into Five, 'Turbocharging' E&S", Carrier Management, April 28, 2025
https://www.carriermanagement.com/features/2025/04/28/274588.htm

ii Saumya Jain, "Allianz UK Introduces generative AI tool BRIAN", Reinsurance News, September 15, 2025
https://www.reinsurancene.ws/allianz-uk-introduces-generative-ai-tool-brian/

iii Tata Consultancy Services, https://www.tcs.com/what-we-do/perpetually-adaptive-enterprise

iv Ng S.T. Chong, "AI in 2024: Surpassing Human Capabilities, Revolutionizing Industries, and Facing Key Challenges", January 21, 2025.
https://c3.unu.edu/blog/ai-in-2024-surpassing-human-capabilities-revolutionizing-industries-and-facing-key-challenges

v "Fusion Teams: A Proven Model for Digital Delivery", Gartner, March 17, 2023

About Tata Consultancy Services

Tata Consultancy Services (TCS) (BSE: 532540, NSE: TCS) is a digital transformation and technology partner of choice for industry-leading organizations worldwide. Since its inception in 1968, TCS has upheld the highest standards of innovation, engineering excellence and customer service.

Rooted in the heritage of the Tata Group, TCS is focused on creating long term value for its clients, its investors, its employees, and the community at large. With a highly skilled workforce of over 600,000 employees in 55 countries and 202 service delivery centers across the world, the company has been recognized as a top employer in six continents. With the ability to rapidly apply and scale new technologies, the company has built long term partnerships with its clients – helping them emerge as perpetually adaptive enterprises. Many of these relationships have endured into decades and navigated every technology cycle, from mainframes in the 1970s to Artificial Intelligence today.

TCS sponsors 14 of the world's most prestigious marathons and endurance events, including the TCS New York City Marathon, TCS London Marathon and TCS Sydney Marathon with a focus on promoting health, sustainability, and community empowerment.

TCS generated consolidated revenues of over US $30 billion in the fiscal year ended March 31, 2025. For more information, visit www.tcs.com

Tata Consultancy Services: Building Perpetually Adaptive Enterprises
TCS, a global leader in IT services, consulting, and business solutions, leverages technology for business transformation and helps catalyze change. **Learn more.**