



Is My Core Insurance Platform Ready for Agentic AI?

A Checklist for Insurance CEOs and CIOs

Why We Designed This Checklist

The emergence of agentic AI represents the next significant technology leap for insurers. Much has been written about its transformative power. However, the implications for insurance core platforms to successfully scale agentic AI have not yet received much coverage.

To harness the full potential of agentic AI, insurance leaders must ensure that their core insurance platforms possess the essential capabilities to integrate agentic AI safely, efficiently and effectively. Achieving this readiness extends beyond the responsibilities of enterprise IT leaders alone—it encompasses the broader organization, including the CEO and business unit leaders.

We designed this checklist along key thematic dimensions as a reference for insurance CEOs and CIOs. This checklist is a shortened, executive version of the first part of a two-part series on **Insurance Core Platforms in the Age of Agentic AI**.

Please see the separate publications for details:

- [Part 1: Insurance Core Platform Requirements to Win in an AI-First Future](#)
- [Part 2: Modernization Approaches to Prepare for an AI-First Future](#)

How to Use This Checklist

We recommend that insurance CEOs and CIOs use this checklist twofold:

- To assess their current insurance core platform's readiness and gaps in safely, efficiently and effectively scaling agentic AI and to identify key investment needs.
- To evaluate the degree of "future proofness" and agentic AI readiness of vendors' insurance core platforms during sourcing processes and procurement decisions.

Core technology decisions affect several strategy cycles to come. Accordingly, it is critical to be thorough in evaluating forward-looking capabilities.

Checklist Summary

1. Robust Integration & Interoperability



2. Dynamic Workflow Orchestration



3. Dynamic Scalability & Performance



4. Stringent Security, Access Control & Observability



5. Unified Data Management & Governance



6. Flexible Product & Pricing Engine



7. Other Key Requirements



The Agentic AI Core Platform Checklist

1. Robust Integration & Interoperability

Agentic AI systems must integrate seamlessly within an ecosystem of applications, processes and data sources. Integration and interoperability capabilities ensure that AI agents can interact smoothly with the various modules, functionalities, and data housed within the core platform.

| Checklist Item | Why It Matters |
|--|--|
| APIs Does the core platform offer comprehensive APIs following interoperability standards? | The core platform needs to comprehensively expose its functionality and data through APIs, allowing AI agents to invoke all core functions (e.g., underwriting, claims). ¹ APIs should adhere to interoperability standards (e.g., OAS) to streamline integration efforts and make it easier for AI agents to interact with the core platform. ² |
| Event-Driven Architecture Does the core platform follow event-driven architecture? | Event streams or an event mesh architecture enable AI agents to subscribe to and asynchronously react to business events. AI agents, subscribing to these event streams (e.g., “new claim created”), can process this information in near real-time without human triggers (e.g., enabling them to swiftly triage claims or flag potential fraud for further investigation). |
| Microservices Is the core platform comprehensively built using microservice architecture? | An AI agent must interact seamlessly with the core platform. This necessitates an open and accessible architecture, benefitting from microservices. In addition, an adequate microservices architecture delivers further benefits, such as higher fault tolerance, faster innovation cycles, flexible scaling —all benefitting agentic AI adoption. |

¹ Graphical user interface (GUI) based agents are less efficient, scalable and effective, failing to deliver deeper progression into agentic AI capabilities.

² The next integration evolution, Model Context Protocol (MCP), represents an open protocol standardizing how applications provide data and context to AI applications. Regardless, core platforms require robust APIs to integrate with MCP servers.

2. Dynamic Workflow Orchestration

Agentic AI must seamlessly embed within core insurance processes, driving value either by automating entire workflows or augmenting specific workflow steps and decisions. Core platforms require advanced workflow and orchestration capabilities to effectively integrate AI-driven actions in the context of insurers’ complex, rules-driven sequences with multiple interactions.

| Checklist Item | Why It Matters |
|---|--|
| <p>AI-Ready Workflow Engine</p> <p>Can the core platform’s workflow engine be configured to trigger AI-driven decisioning and task handling?</p> | <p>Unlocking the potential of agentic AI necessitates dynamic workflow logic capable of accommodating insights and branching decisions, allowing agents to invoke other processes and agents, such as an up-sell process during a claim interaction. The workflow engine must facilitate external API calls, event-driven waits and dynamic routing based on AI responses, incorporating necessary timeouts and fallback procedures.</p> |
| <p>Human-in-the-Loop Coordination</p> <p>Can the core platform’s workflow engine be configured to gracefully include human oversight and intervention?</p> | <p>Many agentic AI workflows will, for the foreseeable future, require human oversight or intervention, requiring the graceful integration of human steps (e.g., if the AI agent’s claims approval limit is exceeded). Such orchestration includes creating tasks in human work queues, managing notifications, and seamlessly merging human inputs back into the workflow.</p> |
| <p>Explainability & Governance</p> <p>Can the core platform enforce process adherence and document decision paths?</p> | <p>AI-driven workflow orchestration must prevent deviations from standard processes and ensure compliance, transparency and auditability. Core platforms must provide mechanisms enabling AI agents to transparently explain their recommendations or actions (e.g., human-readable justifications, clearly identifying decision factors and maintaining chain-of-thought paths).</p> |

3. Dynamic Scalability & Performance

Successfully deploying AI agents at scale requires insurance core platforms to be prepared for the substantial loads and performance demands imposed by agentic AI workflows. Robust scalability and performance capabilities ensure the core platform remains responsive, reliable, and cost-efficient as AI integration expands.

| Checklist Item | Why It Matters |
|---|---|
| Real-Time Processing Can the core platform deliver low latency, real-time processing? | Agentic AI necessitates real-time responsiveness instead of typical overnight batch processing common for most core platforms. This delivers better, real-time customer service and prevents incorrect decisions by AI agents based on outdated information. |
| Elastic Scaling Can the core manage high concurrency and elastic scaling? | AI agents can significantly amplify transaction volumes. A single task traditionally executed by a human might involve multiple specialized AI agents concurrently performing numerous actions, such as checking databases and calling internal and external APIs. Cloud-native architecture enables seamless addition of computing power and storage to enable the required parallel processing. |
| Performance Management Is the core platform able to monitor and apply performance policies to manage performance-cost tradeoff? | Effective management of the performance-cost tradeoff is essential. Interactions between core platforms and AI agents can be resource-intensive, particularly when invoking complex AI computations, such as large models for reasoning. Limiting the frequency of expensive AI services or scheduling resource-heavy tasks as batch inference can realize significant cost savings. |

4. Stringent Security, Access Control & Observability

With great autonomy comes great responsibility—and risk. When AI agents trigger actions across systems and deal with sensitive data, robust security, authority management and access control frameworks become critical. Core platforms must extend and reinforce security and auditing mechanisms to manage AI-driven operations effectively.

| Checklist Item | Why It Matters |
|--|--|
| <p>Identity and Authority Management</p> <p>Does the core platform support fine-grained authority management and identity propagation that can extend to AI agents?</p> | <p>To manage the risk of autonomous agentic AI systems, every AI agent's action must be confined within defined permissions and authority limits (e.g., restricting agents from deleting records or limiting their actions to specific APIs or to specific approval thresholds). The core platform must maintain the initiating agent's authentication context throughout the entire chain of interactions and provide clear safeguards, ensuring that AI cannot “go rogue”.</p> |
| <p>Real-Time Monitoring</p> <p>Can the core platform support real-time monitoring and logging of all AI activity?</p> | <p>Core platforms must treat AI components as first-class citizens within their monitoring infrastructures, ensuring abnormalities can be quickly detected and performance deterioration addressed. Furthermore, stringent audit requirements for AI-driven transactions can be expected, making it critical that the core platform can comprehensively log every AI action and trace transactions across distributed systems and ensuring detailed audit trails for forensic analyses (e.g., what information the AI possessed and how it acted upon it).</p> |
| <p>Cybersecurity</p> <p>Does the core platform have a strong cybersecurity framework to address traditional and AI-specific threats?</p> | <p>Agentic AI integration inherently expands potential security vulnerabilities by exposing more platform functionalities. Regular vulnerability assessments, penetration testing, and intrusion detection systems become essential proactive measures. New AI-specific risks emerge, such a prompt injection, requiring additional safeguards (e.g., content filtering).</p> |

5. Unified Data Management & Governance

Agentic AI effectiveness hinges on the quality and accessibility of data. Insurance carriers maintain extensive datasets—the diversity and volume of data are continually expanding with evolving business models like embedded and usage-based insurance. Thus, insurance core platforms must adopt robust data management and governance practices to ensure high-quality, compliant, and ethically managed data is available to AI agents.

| Checklist Item | Why It Matters |
|---|---|
| <p>Data Unification</p> <p>Does the core platform provide unified access to data (e.g., through a built-in data warehouse covering all modules)?</p> | <p>Insurance data frequently suffers from fragmentation across different schemas, data models, and systems. It is often also not enriched with external data sources. Standardized data schemas across the organization enable AI agents to interpret and utilize multi-source data effectively (e.g., for complex tasks such as upselling or fraud detection).</p> |
| <p>Data Quality & Tracking</p> <p>Can the core platform enforce data quality standards and data lineage tracking?</p> | <p>Effectiveness of AI hinges on data quality. Poor data quality emerges as a critical obstacle for many insurers—reinforcing the adage, "garbage in, garbage out." In line with previous comments about auditing, it is critical to systematically capture data lineage of agent-generated outputs and transformations.</p> |

6. Flexible Product & Pricing Engine

The product engine is the brain of the insurance core platform, feeding all insurance product data to AI agents. Agentic AI offers the potential for increased personalization, leveraging diverse data assets to tailor propositions to customers. However, to achieve personalization in the product offering (and not just in customer communication), the core platform needs a truly flexible product engine.

| Checklist Item | Why It Matters |
|--|--|
| <p>Flexible Configuration</p> <p>Does the core platform’s product engine support flexible, multi-tier product structures?</p> | <p>To achieve personalization, a flexible product engine must support dynamic coverage options and real-time policy tailoring, facilitating rapid and secure innovation and seamless integration with analytics pipelines and AI-powered pricing and decision engines. Multi-tier product structures allow the separation of technical products and marketing products, allowing AI to assemble tailored marketing products and bundles on top of pre-approved, regulatory-compliant technical products.</p> |
| <p>Robust Product Governance</p> <p>Can the core platform’s product engine ensure product governance?</p> | <p>If AI can assemble future marketing products, then robust versioning capabilities, as well as parent-child product inheritance, are needed to ensure product updates remain compliant and traceable despite frequent revisions.</p> |

7. Other Key Requirements

The ability to efficiently scale AI in a world of accelerated change hinges on additional requirements, ranging from multi-tenancy/multi-country scalability to easy maintainability and upgradeability to native AI capabilities within the core platform itself.

| Checklist Item | Why It Matters |
|--|--|
| Multi-Tenancy & Multi-Country Does the core platform support multi-tenant and multi-country deployments? | Investments into agentic AI, whether the actual agents or the supporting core platform and data infrastructure, can be significant. The best investments are those that can be shared and re-used across country units. Different countries running on the same tech stack and a similar data architecture will not just make it easier to build reusable AI agents but also to train and improve them on larger data sets. |
| Maintainability & Upgradeability Can the core platform be easily upgraded and is it supported by a strong product roadmap? | Agentic AI is evolving rapidly. Underlying concepts, models, technologies and regulations frequently iterate and advance. Consequently, AI-first core insurance platforms must be designed for easy maintainability and, most importantly, upgradeability. Insurers do not want to end up in the trap of customized platforms without a roadmap to keep up with rapidly evolving technological and regulatory changes. |
| AI-Native Capabilities Does the core platform provide native capabilities for agentic AI (e.g., built-in AI agent orchestration capabilities)? | AI agents may be defined, managed and orchestrated from within the core application or using an external application that interacts with the core platform. Where all or most tasks of an AI agent are related to core insurance processes (i.e., processes happening predominantly across the different modules of the core platform), native AI capabilities in the core platform can offer benefits of lower deployment and maintenance costs due to pre-integration and improved governance and control. |

Recommendations in Summary

While AI is evolving quickly and nobody can confidently predict the scope and timing of new breakthroughs, it is critical to opt for a core platform providing maximum long-term optionality. The right architecture and technology stack are critical in this regard, facilitating quick and agile adoption to new changes.

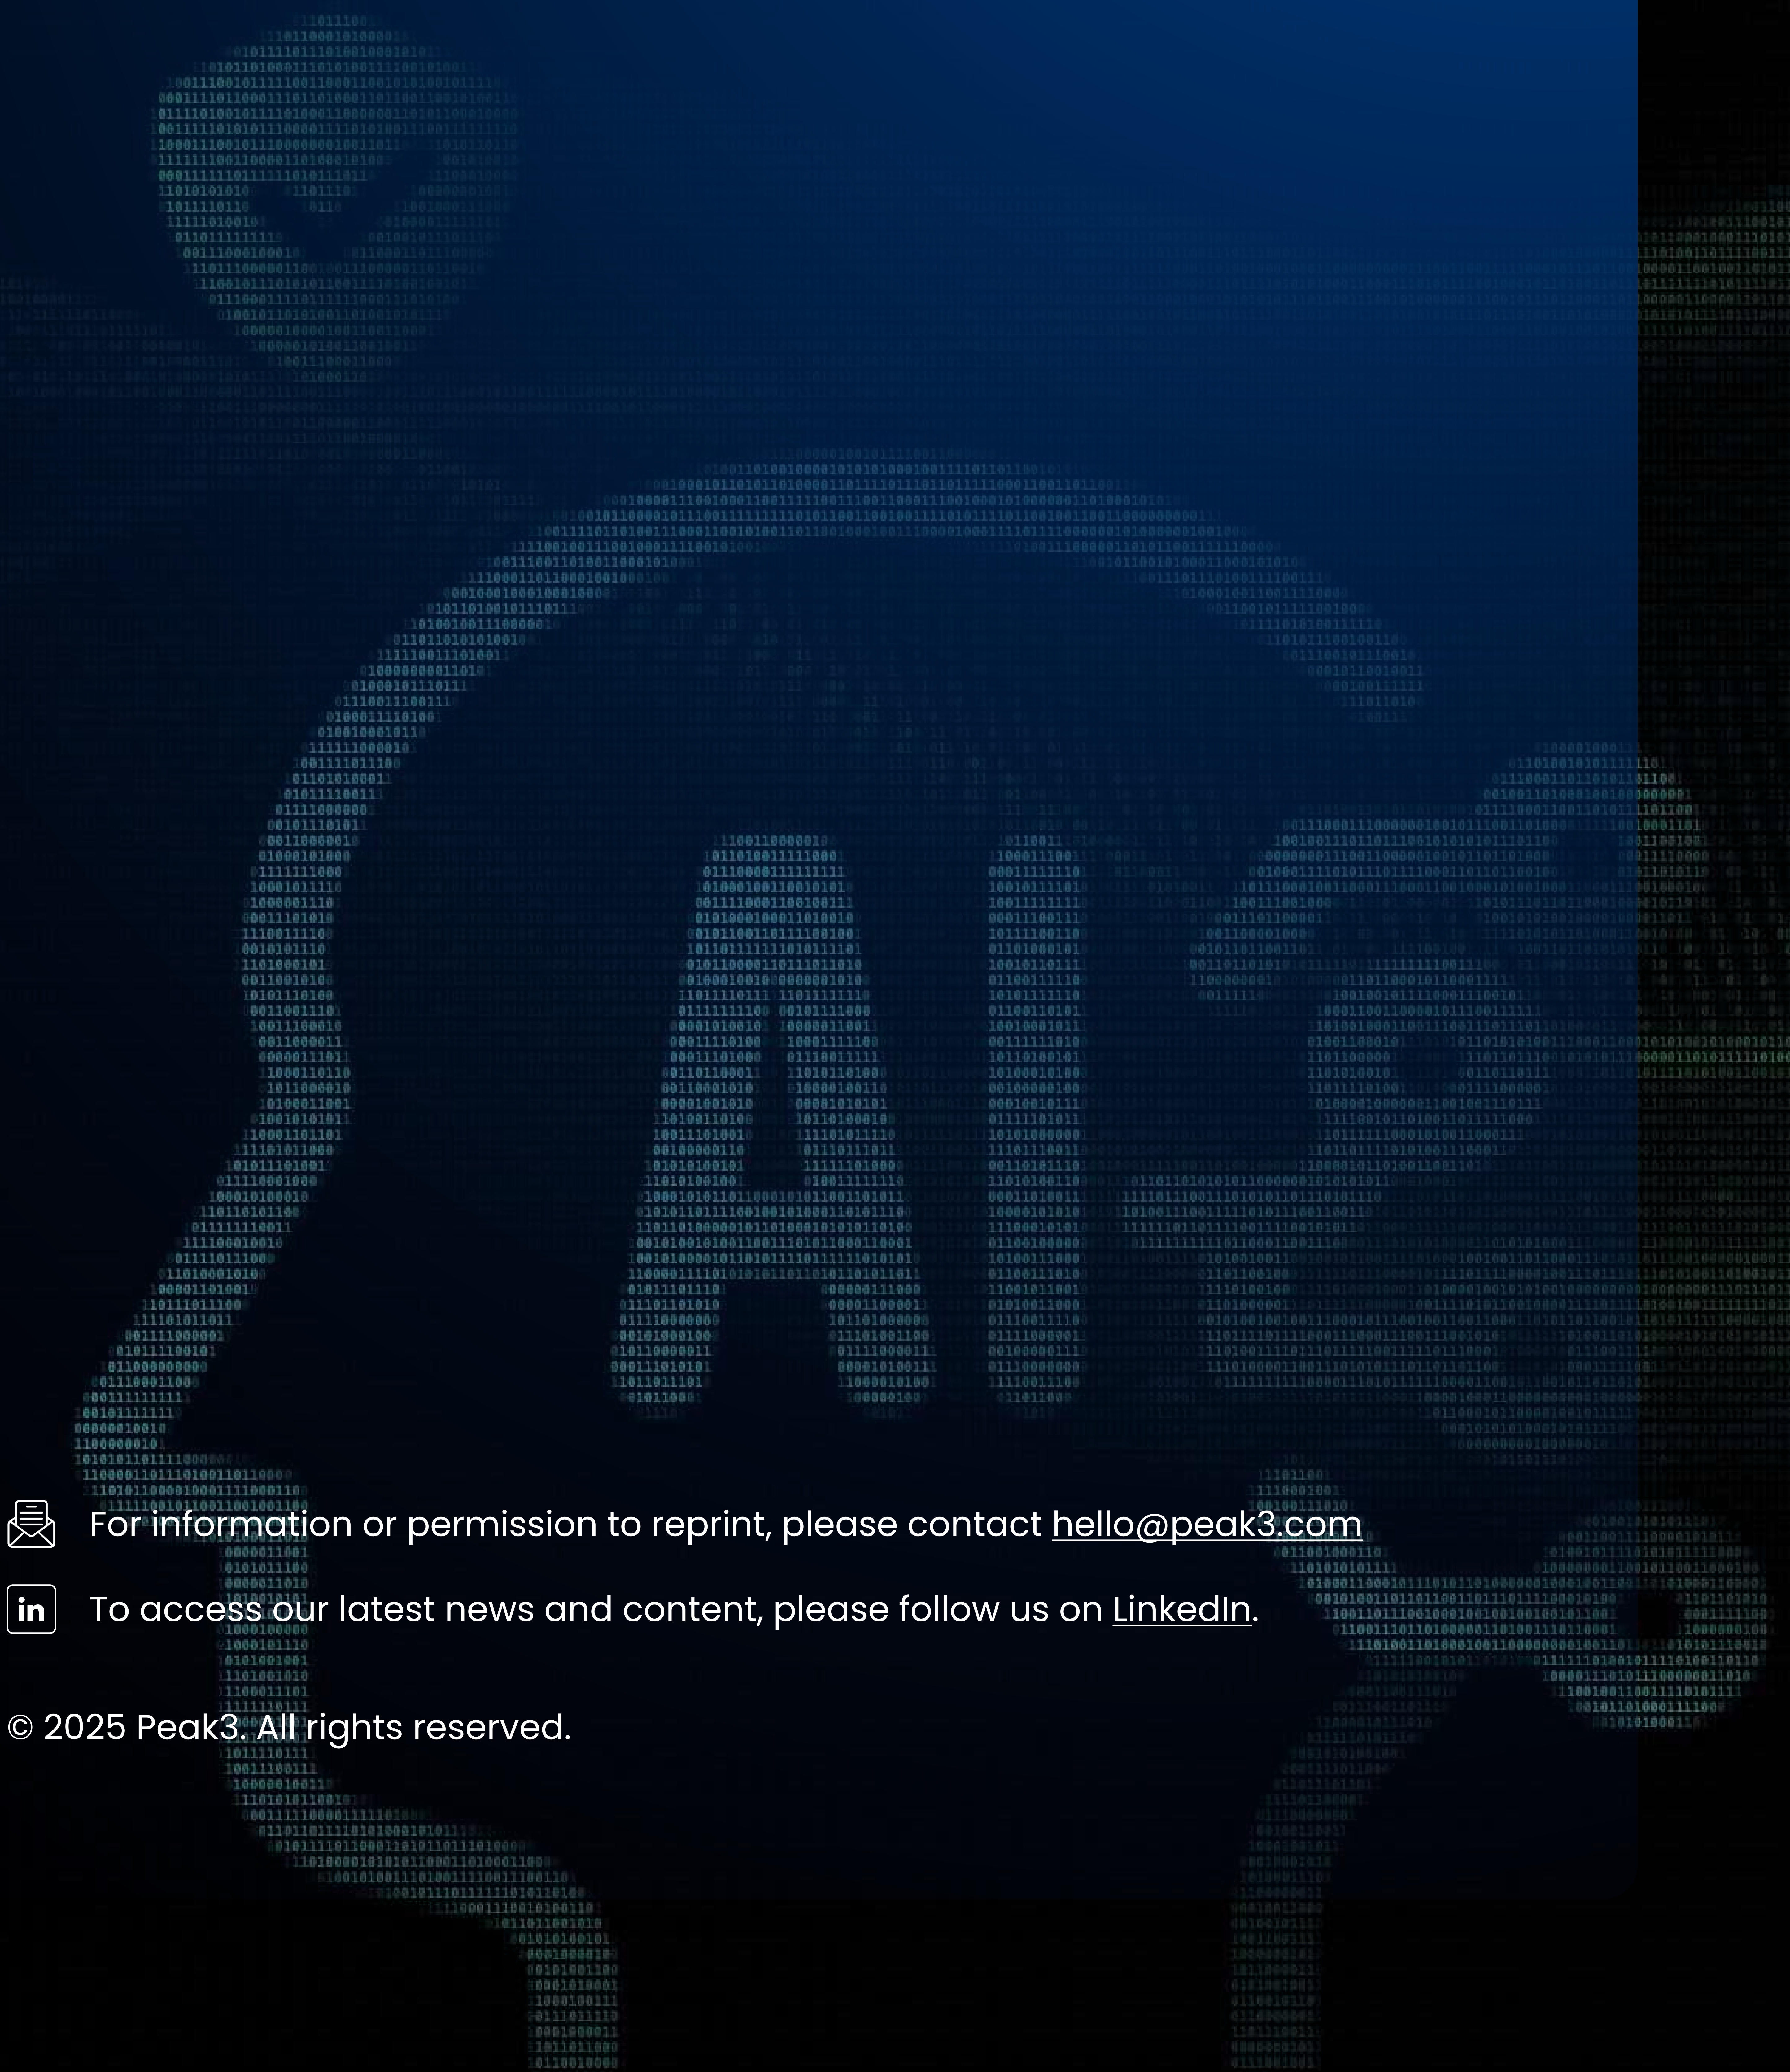
Insurance leaders should use this non-exhaustive checklist to evaluate their current insurance core platform(s) readiness to scale agentic AI and to identify key investment needs. If deciding to close identified gaps with third-party vendor solutions, insurance leaders should thoroughly evaluate not just how vendors currently perform against the aforementioned requirements, but also how future proof their underlying architecture and product strategy are.

Insurance leaders should make their core platform an asset and not a liability in their agentic AI journey—an asset that accelerates the adoption and value realization of agentic AI, while providing clear safeguards to mitigate risks.

About Peak3

Founded by insurance, digital and technology experts, Peak3 powers the digital operating system of the global insurance industry. We combine insurance core, distribution, and AI solutions to deliver a step change in performance for insurers, MGAs, and insurance intermediaries.

From greenfield embedded insurance ventures to digital-first, multi-country core modernization programs, our cloud-native SaaS solutions power top customers across life, health, and P&C insurance.



For information or permission to reprint, please contact hello@peak3.com



To access our latest news and content, please follow us on [LinkedIn](#).

© 2025 Peak3. All rights reserved.