Peak3

Insurance Core Platforms in the Age of Agentic AI

PART 1

Insurance Core Platform Requirements to Win in an AI-First Future

Introduction to this Publication

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

In this Part 1 of a two-part series on **Insurance Core Platforms in the Age of Agentic AI**, we outline eight thematic functional and technological success factors and requirements of core insurance platforms for an agentic AI future. In **Part 2**, we discuss different modernization approaches and no-regret moves to meet these capability requirements.

This publication is not meant to deeply delve into the architecture and technological foundations of AI agents, the different types of agents, and the use cases. Rather, it focuses on guiding insurance CIOs and CTOs to think strategically about their core platforms within the context of the paradigm shift agentic AI presents.

Parts of this document are technical. For non-technical readers, we recommend the briefer publication: **Is My Core Insurance Platform Ready for Agentic AI? – A Checklist for Insurance CEOs and CIOs.**

Recommendations

- Align your long-term vision for (agentic) AI on senior-level and deduce the role of your core platform(s) to meet that vision.

- Beyond evaluating the technical architecture and use cases of AI agents, clarify what the dependencies and implications for your core platform(s) are.

- Evaluate your current systems and potential new vendor solutions against the requirements set out in this paper and the supporting checklist to identify gaps. Incorporate or adjust the requirements as AI technologies evolve.

- Be ready to pivot and focus investments on no-regret moves, such as robust integration capabilities.

- Insurers with operations across countries: Strive for multi-tenant and multi-country capable core platforms to multiply benefits by making it easier to build reusable AI agents and train and improve them on larger data sets.

Essential Core Platform Capabilities to Scale Agentic AI

To harness the full potential of agentic AI, insurance leaders must ensure 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.

In the sections that follow, we outline the key functional and technological capabilities necessary to achieve agentic AI readiness, systematically organized by thematic categories (see *Illustration 1* for an overview).

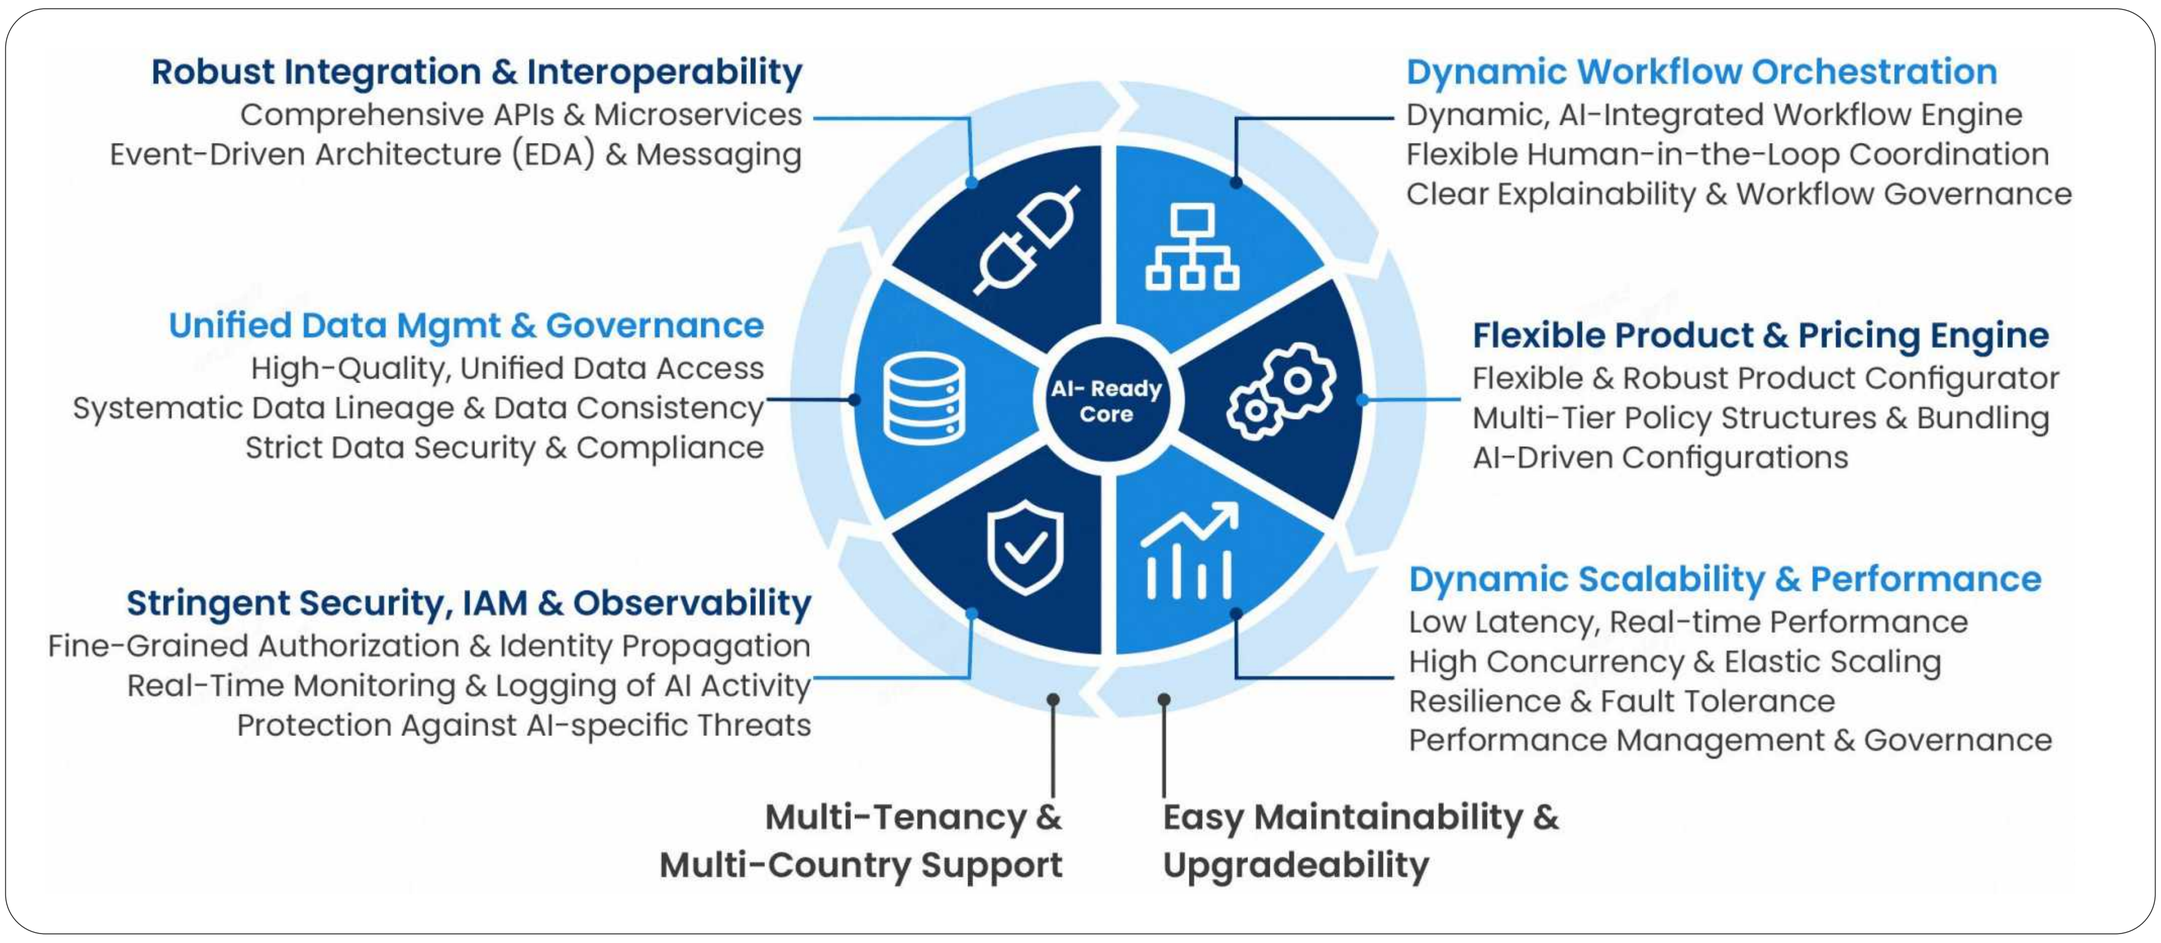

ILLUSTRATION 1: Essential insurance core platform capabilities to scale agentic AI

Some readers might question whether an insurance core platform remains necessary in the age of agentic AI. Could the business logic, rules, and calculations traditionally managed by core platforms be fully absorbed into an agentic AI application layer, allowing direct interactions with the database and potentially making core platforms obsolete? We firmly believe that the insurance core platform will continue to play an essential role and outlined this in a separate blog post.[1]

---

[1] See the Peak3 blog post: *AI Killed the Insurance Core System. Long Live the Insurance Core System*.

Exhibit: What are AI Agents?

Gartner® defines AI agents as "goal-driven software entities that have been granted rights by the organization to act on its behalf to autonomously make decisions and take action. These entities use AI techniques—combined with components such as memory, planning, sensing, tooling, and guardrails—to complete tasks and achieve objectives."[2]

AI agents can perceive their environments, reason about complex situations, and act to achieve specific objectives, often involving multiple steps and systems or other AI agents. Built on large language models (LLMs), AI agents mimic human cognitive functions like learning, problem-solving, and decision-making, but with the advantages of processing vast amounts of data more rapidly.

Agentic AI is positioned as a key enabler for creating more efficient, customer-centric, resilient and responsive insurance operations. The adoption of agentic AI will evolve from simple AI agents with restricted autonomy to execute specialized tasks to comprehensive AI agent ecosystems, where collaborative agents orchestrate across diverse task-specific agents with a high degree of autonomy.

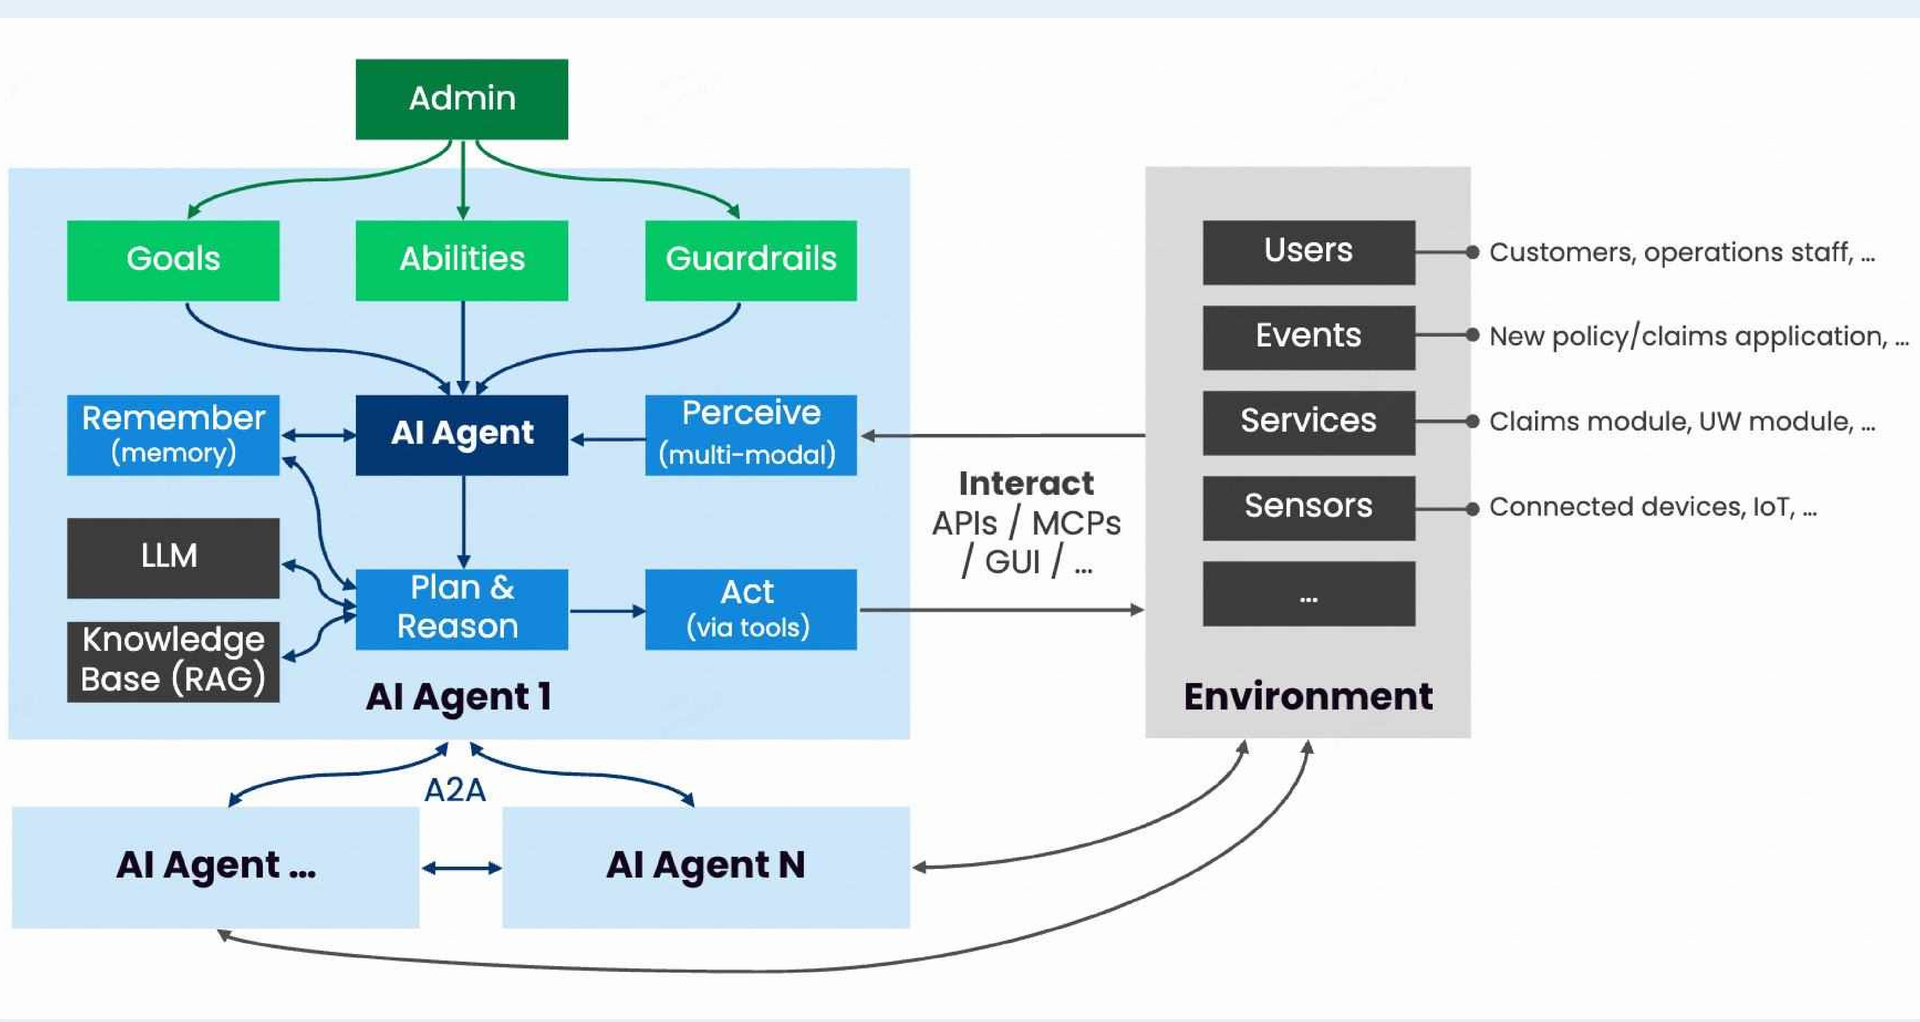

ILLUSTRATION 2: Simplified, illustrative architecture of an AI agent

---

[2] Gartner, *TSP 2025 Trends: Agentic AI — The Evolution of Experience* by Jim Hare, Tom Coshow, Mark McDonald, Radu Miclaus, and Sid Nag, 24 February 2025 (for Gartner subscribers only). GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

1. Robust Integration & Interoperability

Agentic AI systems must integrate seamlessly within a complex ecosystem of applications and data sources. Core insurance platforms typically manage numerous interconnected processes, including policy quoting, policy servicing, claims intake, settlement, and billing.

Integration and interoperability capabilities ensure that AI agents can interact smoothly with the various modules, functionalities, and data housed within the core platform.

Key Capabilities

Comprehensive APIs & Microservices

The core platform should expose its functionality and data through APIs, allowing AI agents to invoke business functions effortlessly. An LLM must interact seamlessly with the core platform, which necessitates an open and accessible architecture, benefitting from microservices.

For example, an AI claims triage agent might intelligently extract data from medical reports, proactively follow up with customers if documentation is incomplete, utilize policy APIs to verify eligibility and match policies, and invoke claims APIs to register the claim.

APIs should be thoroughly documented and adhere to interoperability standards such as the OpenAPI Specification (OAS), streamlining integration efforts and making it easier for AI agents to interact with the core platform.
The next evolution, Model Context Protocol (MCP), represents an open protocol standardizing how applications provide data and context to AI applications. Regardless, core platforms must maintain API-based integration with MCP servers.

> **If all your technology is not exposed through the right set of APIs and a flexible set of microservices, it'll be hard to deliver agentic experiences.** [3]
>
> - Chief Data Officer, Intuit

[3] See the CIO article: *What gives IT leaders pause as they look to integrate agentic AI with legacy infrastructure*.

While graphical user interface (GUI)-based agents are gaining some traction—interacting directly with applications through their interfaces—they offer only limited and temporary solutions. Compared to API integration, GUI interaction is less efficient, scalable, and effective, and it fails to facilitate deeper progression into agentic AI capabilities.

Event-Driven Architecture (EDA) & Messaging

Beyond traditional request-response APIs, event-driven integration significantly enhances agentic AI. Event streams or an event mesh architecture enable AI agents to subscribe to and asynchronously react to business events. Decoupling interactions through events simplifies the process of adding or updating AI agents without affecting existing systems.

For instance, a claims module might publish events such as “New Claim Created” or “Claim Closed”. AI agents, subscribing to these event streams, can process this information in near real-time, enabling them to swiftly triage claims or flag potential fraud for further investigation. This approach also benefits other areas across the policy lifecycle—for example, a "New Policy Bound" event could automatically trigger AI-driven cross-selling activities.

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. Insurance processes, such as claims handling and underwriting, involve complex, rules-driven sequences with multiple interactions. Thus, core platforms require advanced workflow and orchestration capabilities that effectively integrate AI-driven actions and adapt dynamically to evolving decision-making.

Key Capabilities

Dynamic, AI-Integrated Workflow Engine

While most core platforms support basic workflow and decision engines, they are traditionally optimized for human-centric processing. To be truly AI-ready, these workflow engines must support the integration of AI-driven decisions and triggers at designated points within workflows. They must facilitate external API calls, event-driven waits, and dynamic routing based on AI responses, incorporating necessary timeouts and fallback procedures—for example, defaulting to human intervention if an AI agent delays its response.

For example, depending on the claims type and amount for a hospitalization claim, the workflow may automatically trigger a fraud, waste and abuse (FWA) agent investigating patterns such as unusual patient/provider behavior or excessive treatment. Depending on the resulting FWA risk score and type, different workflow paths could apply, invoking other handling steps, decisions, agents or humans.

Unlocking the potential of agentic AI necessitates dynamic workflow logic capable of accommodating insights and branching decisions, allowing agents to invoke other processes. In practice, if an AI agent identifies new information mid-process—such as a customer indicating a potential sales opportunity or complaint during a claims interaction—the workflow should pivot dynamically to handle the new pathway, potentially initiating parallel subprocesses or alternative routes efficiently.

Flexible Human-in-the-Loop Coordination

Many agentic AI workflows will still require human oversight or intervention for specific tasks, requiring the graceful integration of human steps. Core platform orchestration must seamlessly incorporate pauses, escalations, and hand-offs for human actions. This approach aligns closely with the later outlined access controls and authority frameworks.

For instance, if an AI agent identifies that a claim exceeds its authorization limit, the workflow should automatically route the decision to a human with the appropriate authority. Such orchestration includes creating tasks in human work queues, managing notifications, and merging human inputs back into the workflow seamlessly.

Clear Explainability & Workflow Governance

AI-driven workflow orchestration must remain compliant, transparent, and auditable. Insurers generally seek deterministic outcomes following standard processes. Any deviation from standard processes should be routed to a human with appropriate authorizations and any alternative paths triggered by AI actions should be meticulously documented and governed by the same rigorous business rules applicable to traditional workflows.

Furthermore, core platforms must provide mechanisms enabling AI agents to transparently explain their recommendations or actions. This might include generating human-readable justifications, clearly identifying decision factors, or maintaining detailed logs of chain-of-thought paths. This could include the storage of intermediate prompts and reasoning the agent went through to reach a conclusion. Such mechanisms enhance transparency, facilitate compliance auditing, and maintain accountability within AI-driven processes.

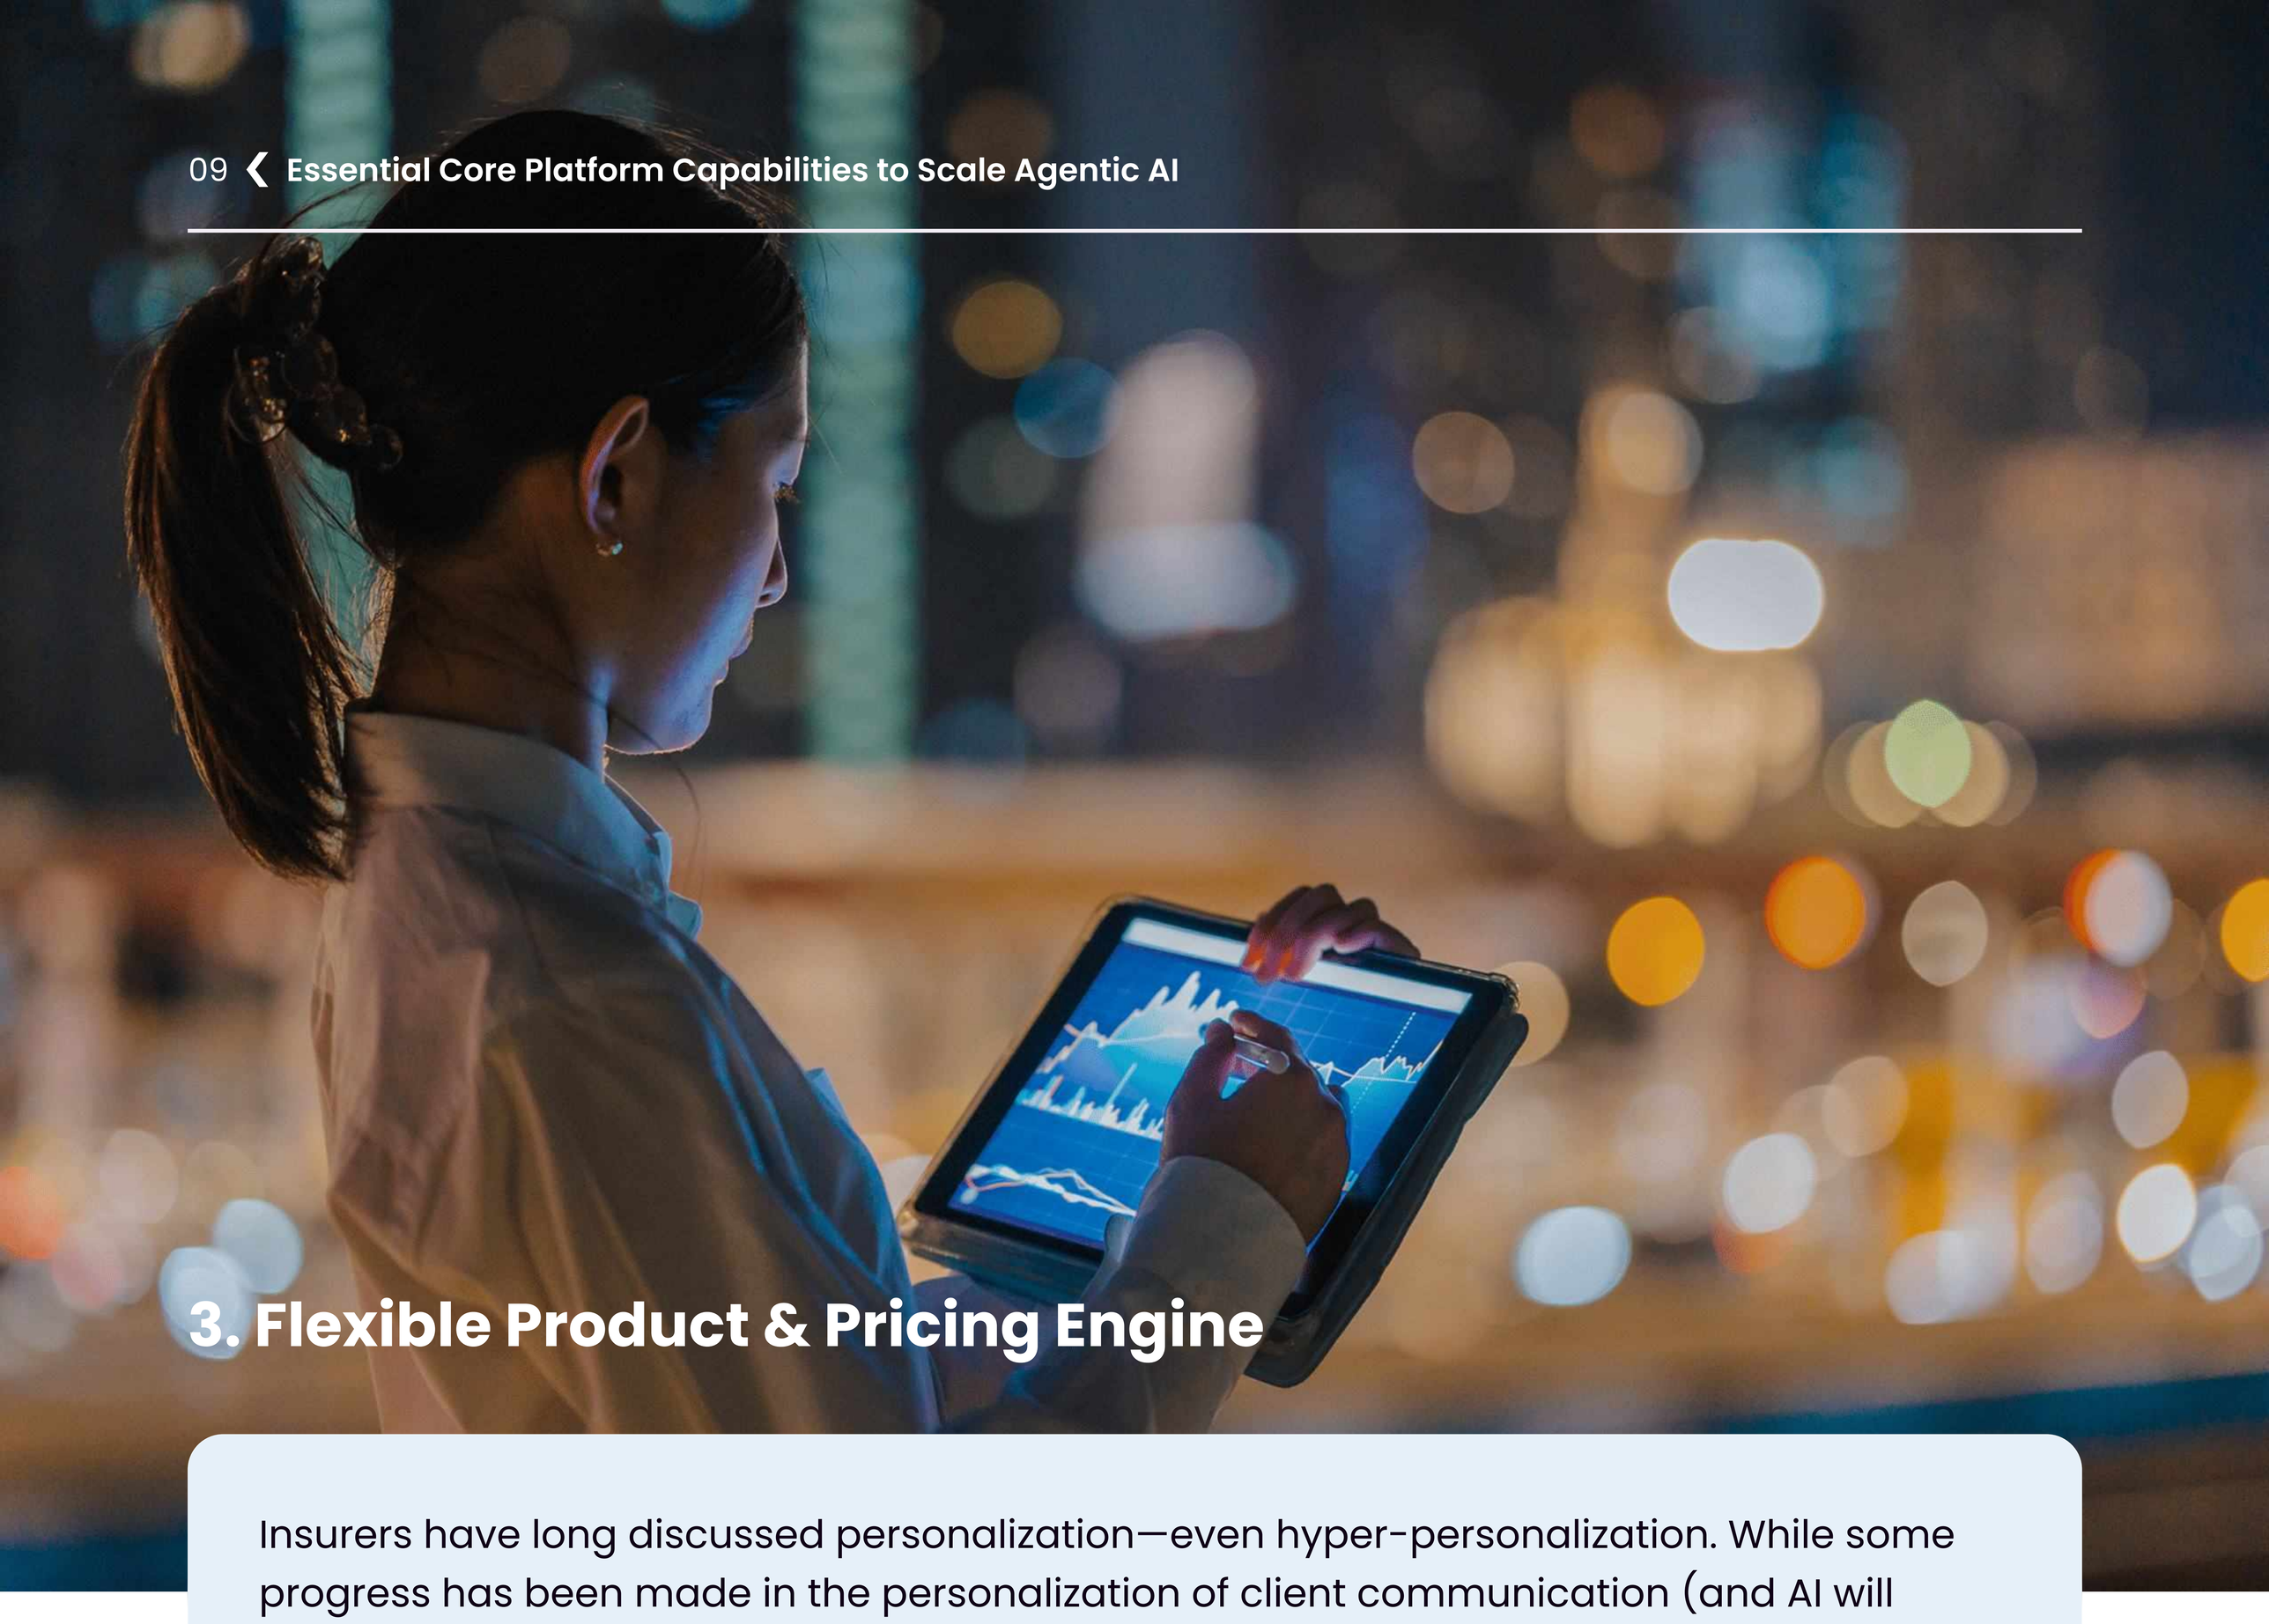

3. Flexible Product & Pricing Engine

Insurers have long discussed personalization—even hyper-personalization. While some progress has been made in the personalization of client communication (and AI will enable more), insurance products are often far from personalized. To thrive in an AI-first and agentic AI future, core insurance platforms require a flexible product engine.

Key Capabilities

Flexible & Robust Product Configurator

A flexible product engine must support dynamic coverage options and real-time policy tailoring, facilitating rapid and secure innovation. Traditional product frameworks cannot adequately address hyper-personalization demands due to their rigidity. The product engine should enable rapid product adaptations, such as integrating new rating factors or combining benefits across product lines, and seamlessly integrate with analytics pipelines and AI-powered pricing and decision engines.

Robust configuration tools and versioning capabilities, as well as parent-child product inheritance, ensure product updates remain compliant and traceable despite frequent revisions. However, any such revision will still depend on product and pricing actuaries prior sign-off, at least for the foreseeable future.

Multi-Tier Policy Structures & Bundling

Regulatory constraints, such as file-and-use laws requiring regulatory filing and approval before selling insurance products, often limit speed to market and personalization capabilities. A flexible product engine capable of creating comprehensive "omnibus" products covering diverse benefit, liability, and rating combinations, which require only a single filing, provides a solution to achieving AI-driven personalization within permissible boundaries (where regulators allow this). Multi-tier product structures and virtual product bundling enhance flexibility, allowing insurers to effectively assemble technical products into compelling market propositions.

AI-Driven Configurations

To accelerate product development, AI-ready platforms will support AI-driven setup of new product propositions. Product specifications devised by product actuaries, whether human or AI-generated, could be automatically configured within the core platform (e.g., uploaded product specifications get directly translated into configurations).

Clear documentation and AI-compatible protocols are essential for enabling direct AI interactions with the product engine. Automatically testing AI-configured products against predefined specifications represents the natural next evolution, ensuring rapid and accurate product deployment.

4. Dynamic Scalability & Performance

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

Key Capabilities

Low Latency, Real-time Performance

Traditional insurance core platforms often rely on batch processing, such as overnight policy updates. In contrast, agentic AI necessitates real-time responsiveness, typically within seconds or minutes. Platforms must therefore support real-time data synchronization or event streaming.

Systems restricted to daily batch jobs may severely limit AI agent effectiveness or potentially result in incorrect decisions based on outdated information. Although insurance is relatively less time-sensitive than other industries, interactive and time-sensitive AI tasks require prompt responses.

For example, when a customer submits a motor insurance claim with accident photos, traditional overnight batch processing may delay the processing and synchronization to the next day. If real-time processing is enabled, an AI agent may analyze images and provide a preliminary settlement estimate within seconds, directly following up with customers or service partners for additional details, preventing customer frustration and operational risks caused by outdated batch jobs.

High Concurrency & Elastic Scaling

AI agents significantly amplify transaction volume and system event complexity. A single task traditionally executed by a human might now involve multiple specialized AI agents concurrently performing numerous actions, such as checking databases, calling internal and external APIs, and completing complex workflows. Thus, platforms must effectively handle high concurrency and parallel processing.

Adopting cloud-native architecture with container orchestration solutions (such as Kubernetes) or serverless components enables seamless addition of computing power and storage as needed. True microservices and distributed architectures naturally facilitate horizontal scaling.

An example scenario could involve a collaborative AI agent ecosystem concurrently processing a single claim: one parsing the submission, another retrieving external data, another verifying policy eligibility, another screening for fraud or abuse, another exploring cross-selling opportunities, and another drafting customer responses—all in parallel and with the autonomy to dynamically branch workflows and invoke other workflows and agents (where needed).

The platform must robustly manage this parallel activity, potentially requiring advanced architectural solutions like multi-tier caching, database read/write separation, optimized indexing, query optimization, and sharding.

Resilience & Fault Tolerance

At scale, system failures are inevitable. The underlying infrastructure may be subject to maintenance or failure, and external dependencies may be unstable, or an agent might produce an error. The core platform must therefore exhibit resilience, ensuring that isolated failures do not cascade throughout the system. Implementing microservices and distributed architectures provides crucial fault isolation, enabling individual services to fail independently without compromising the entire platform. Employing proper message queues as buffers further enhances resilience.

For example, if a critical third-party dependency within a microservice has a short-term outage, such as a government electronic health record service used during the claim submission and verification process, a claims agent may inform the customer that processing is currently not available and add the claim to the message queue of the required services. Once the third-party service is resumed, the process is continued without repeated submission requirements.

Performance Management & Governance

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. Core platforms should therefore include tools to monitor resource utilization and apply governance policies. This might involve rate limiting the frequency

of expensive AI services or scheduling resource-heavy, non-urgent AI tasks during off-peak periods, ensuring optimal allocation of resources and minimizing contention with interactive workloads.

For example, in the case of claims submission, an agent may complete basic verifications (e.g., data integrity check) and assign resource weights to the pending list (e.g., priority, financial impact). Lower ranked cases can then wait for batch inference for in-depth processing (e.g., FWA checks, agentic assessments), which can be managed through EDA. Batch inferencing allows the processing of large jobs asynchronously, allowing cloud vendors and foundational model providers to optimize resource scheduling and achieve economies of scale. This generally leads to significant discounts offered.[4]

[4] Discount models vary by provider and may include lower per token pricing or lower node-hour pricing, amongst others.

5. 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, a robust security, authority management and access control framework becomes critical. Core platforms must extend and reinforce security and auditing mechanisms to manage AI-driven operations effectively.

Key Capabilities

Fine-Grained Authorization & Identity Propagation

Every AI agent's action must be traceable and confined within defined permissions linked to user identities or service accounts. The core platform must maintain the initiating user's authentication context throughout the entire chain of agent interactions. Administrators should be able to define explicit constraints (e.g., via GUI configuration), including whitelisting or blacklisting certain operations—for example, restricting agents from deleting records or limiting their actions to specific APIs.

Regulatory compliance reinforces this necessity. For instance, regulations might prohibit using protected health information for pricing. Agent design is one way to address this. However, effective access control can fundamentally limit AI agent access to such sensitive data.

Real-Time Monitoring & Logging of AI Activity

Core platforms must treat AI components as first-class citizens within their monitoring infrastructures. This involves capturing detailed metrics—such as CPU and memory usage, decision frequency, and event queue lengths—and distributed tracing across services.

An AI-ready core platform must comprehensively log every AI action and trace transactions across distributed systems. Detailed audit trails become crucial during forensic analyses, helping determine precisely what information the AI possessed and how it acted upon it.

Modern observability solutions aggregate logs, metrics, and traces to deliver continuous insights. For agentic AI, monitoring may track metrics like actions per minute, success or failure rates, and AI response times. Integration with security information and event management (SIEM) systems

enables swift quarantine or deactivation of malfunctioning agents.

Interactive dashboards visualizing AI performance—such as task completion times or pending task volumes—provide critical insights to operations teams. These tools support optimizing performance and cost-efficiency by highlighting expensive API usage or inefficiencies.

Protection Against AI-specific Threats

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.

AI-specific risks, such as prompt injection—where malicious inputs manipulate AI behaviors—necessitate integrating content filtering and rigorous validation processes before AI agents act. Traditional cybersecurity principles, including least privilege, defense-in-depth, and strict input validation, must be rigorously enforced, supplemented by specialized AI security measures.

6. Unified Data Management & Governance

Agentic AI effectiveness hinges on the quality and accessibility of data. Insurance carriers maintain extensive datasets, including customer details, policy records, claims history, and historical underwriting data. The diversity and volume of data are continually expanding with evolving business models such as embedded insurance.

Therefore, 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.

Key Capabilities

High-Quality, Unified Data Access

Insurance data is often fragmented across schemas, data models, and systems. Core platforms should enforce data quality standards and provide unified access to data across functional modules through a central repository or data warehouse. Integration with customer data platforms (CDPs), or even a CDP part of the core platform, helps unify data from core modules with external sources. Standardized data schemas across the organization enable AI agents to interpret and utilize multi-source data effectively.

For instance, an AI agent conducting fraud detection or cross-selling initiatives must access a consolidated, 360-degree customer view. Core platforms must connect and persist data to unified datasets, maintaining consistent identifiers for entities like customers, vehicles, and properties.

Systematic Data Lineage & Data Consistency

The core platform should systematically track data lineage—detailing data origins, modifications, and destinations. When AI agents generate outputs (e.g., underwriting decisions), metadata indicating the AI origin, algorithmic inputs, and transformations must accompany these results, on top of standard metadata such as timestamps. The core platform's data catalog should integrate with the AI pipeline, documenting which tables or streams the AI is allowed to use.

As previously noted, real-time data synchronization is critical for agentic AI effectiveness. Platforms must support data streams and near-instantaneous replication to maintain accurate, current information. Effective concurrency controls are also vital for managing simultaneous read-write activities by multiple AI agents. Transactions initiated by AI should respect the same ACID properties (atomicity, consistency, isolation, and durability) as human-initiated ones.

Strict Data Security & Compliance

Core platforms must limit the risk of AI integration breaching compliance requirements. Naturally, data access by AI agents should respect privacy laws and internal governance rules. If an AI agent is interacting outside of the core platform, this is difficult to enforce.

Core platforms must enforce stringent data access controls (as outlined earlier) and manage partial data exposure via mechanisms like data masking without compromising AI functionality. Given the sensitivity of personal, financial, or medical data—regulated under standards like HIPAA—platforms must rigorously govern data visibility and utilization by both human users and AI agents.

Separately, core platforms need to be able to rapidly adapt to changing privacy and AI regulations from insurance regulators and other governing bodies.

7. Multi-Tenancy & Multi-Country Support

Investments into agentic AI, whether the actual agents or the supporting core platform, can be significant—particularly for insurers already struggling with increasing IT budgets at the rate required to keep pace with technology change and inflation. The best investments are those that can be shared and re-used across entities and borders.

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 train and improve them on larger data sets.

As such, the insurance core platform needs to support multi-carrier, multi-currency, multi-language, multi-time zone, and multi-regulatory requirements. Ideally, the core platform supports multi-tenancy, ensuring that different country entities can be onboarded as country tenants, running on the same version and facilitating the transferability of AI applications.

8. Easy Maintainability & Upgradeability

Last but not least, agentic AI is evolving rapidly. Underlying concepts, models, and technologies frequently iterate and advance, with regulatory frameworks anticipated to change at a somewhat slower pace. This evolving landscape not only opens new opportunities for insurers to transform operations but also introduces new insurable risks.

The rapid pace of change in AI demands agility from insurers and their core platforms. Consequently, AI-first core insurance platforms must be designed for easy maintainability and, most importantly, upgradeability. It is crucial that these platforms be maintained and developed by partners possessing a clear strategy and roadmap to ensure continuous adaptation and leadership in an evolving market.

Peak3 and Graphene: Your Enablers for the Agentic Future

Peak3's AI-ready cloud core platform, Graphene, is purpose-built from the outset on microservices architecture to ensure robust integration and seamless interoperability. Each functionality within Graphene can be easily exposed to external applications, such as AI agents, allowing effortless interaction. Its cloud-native, cloud vendor-agnostic architecture guarantees the dynamic scalability and performance required by modern insurers.

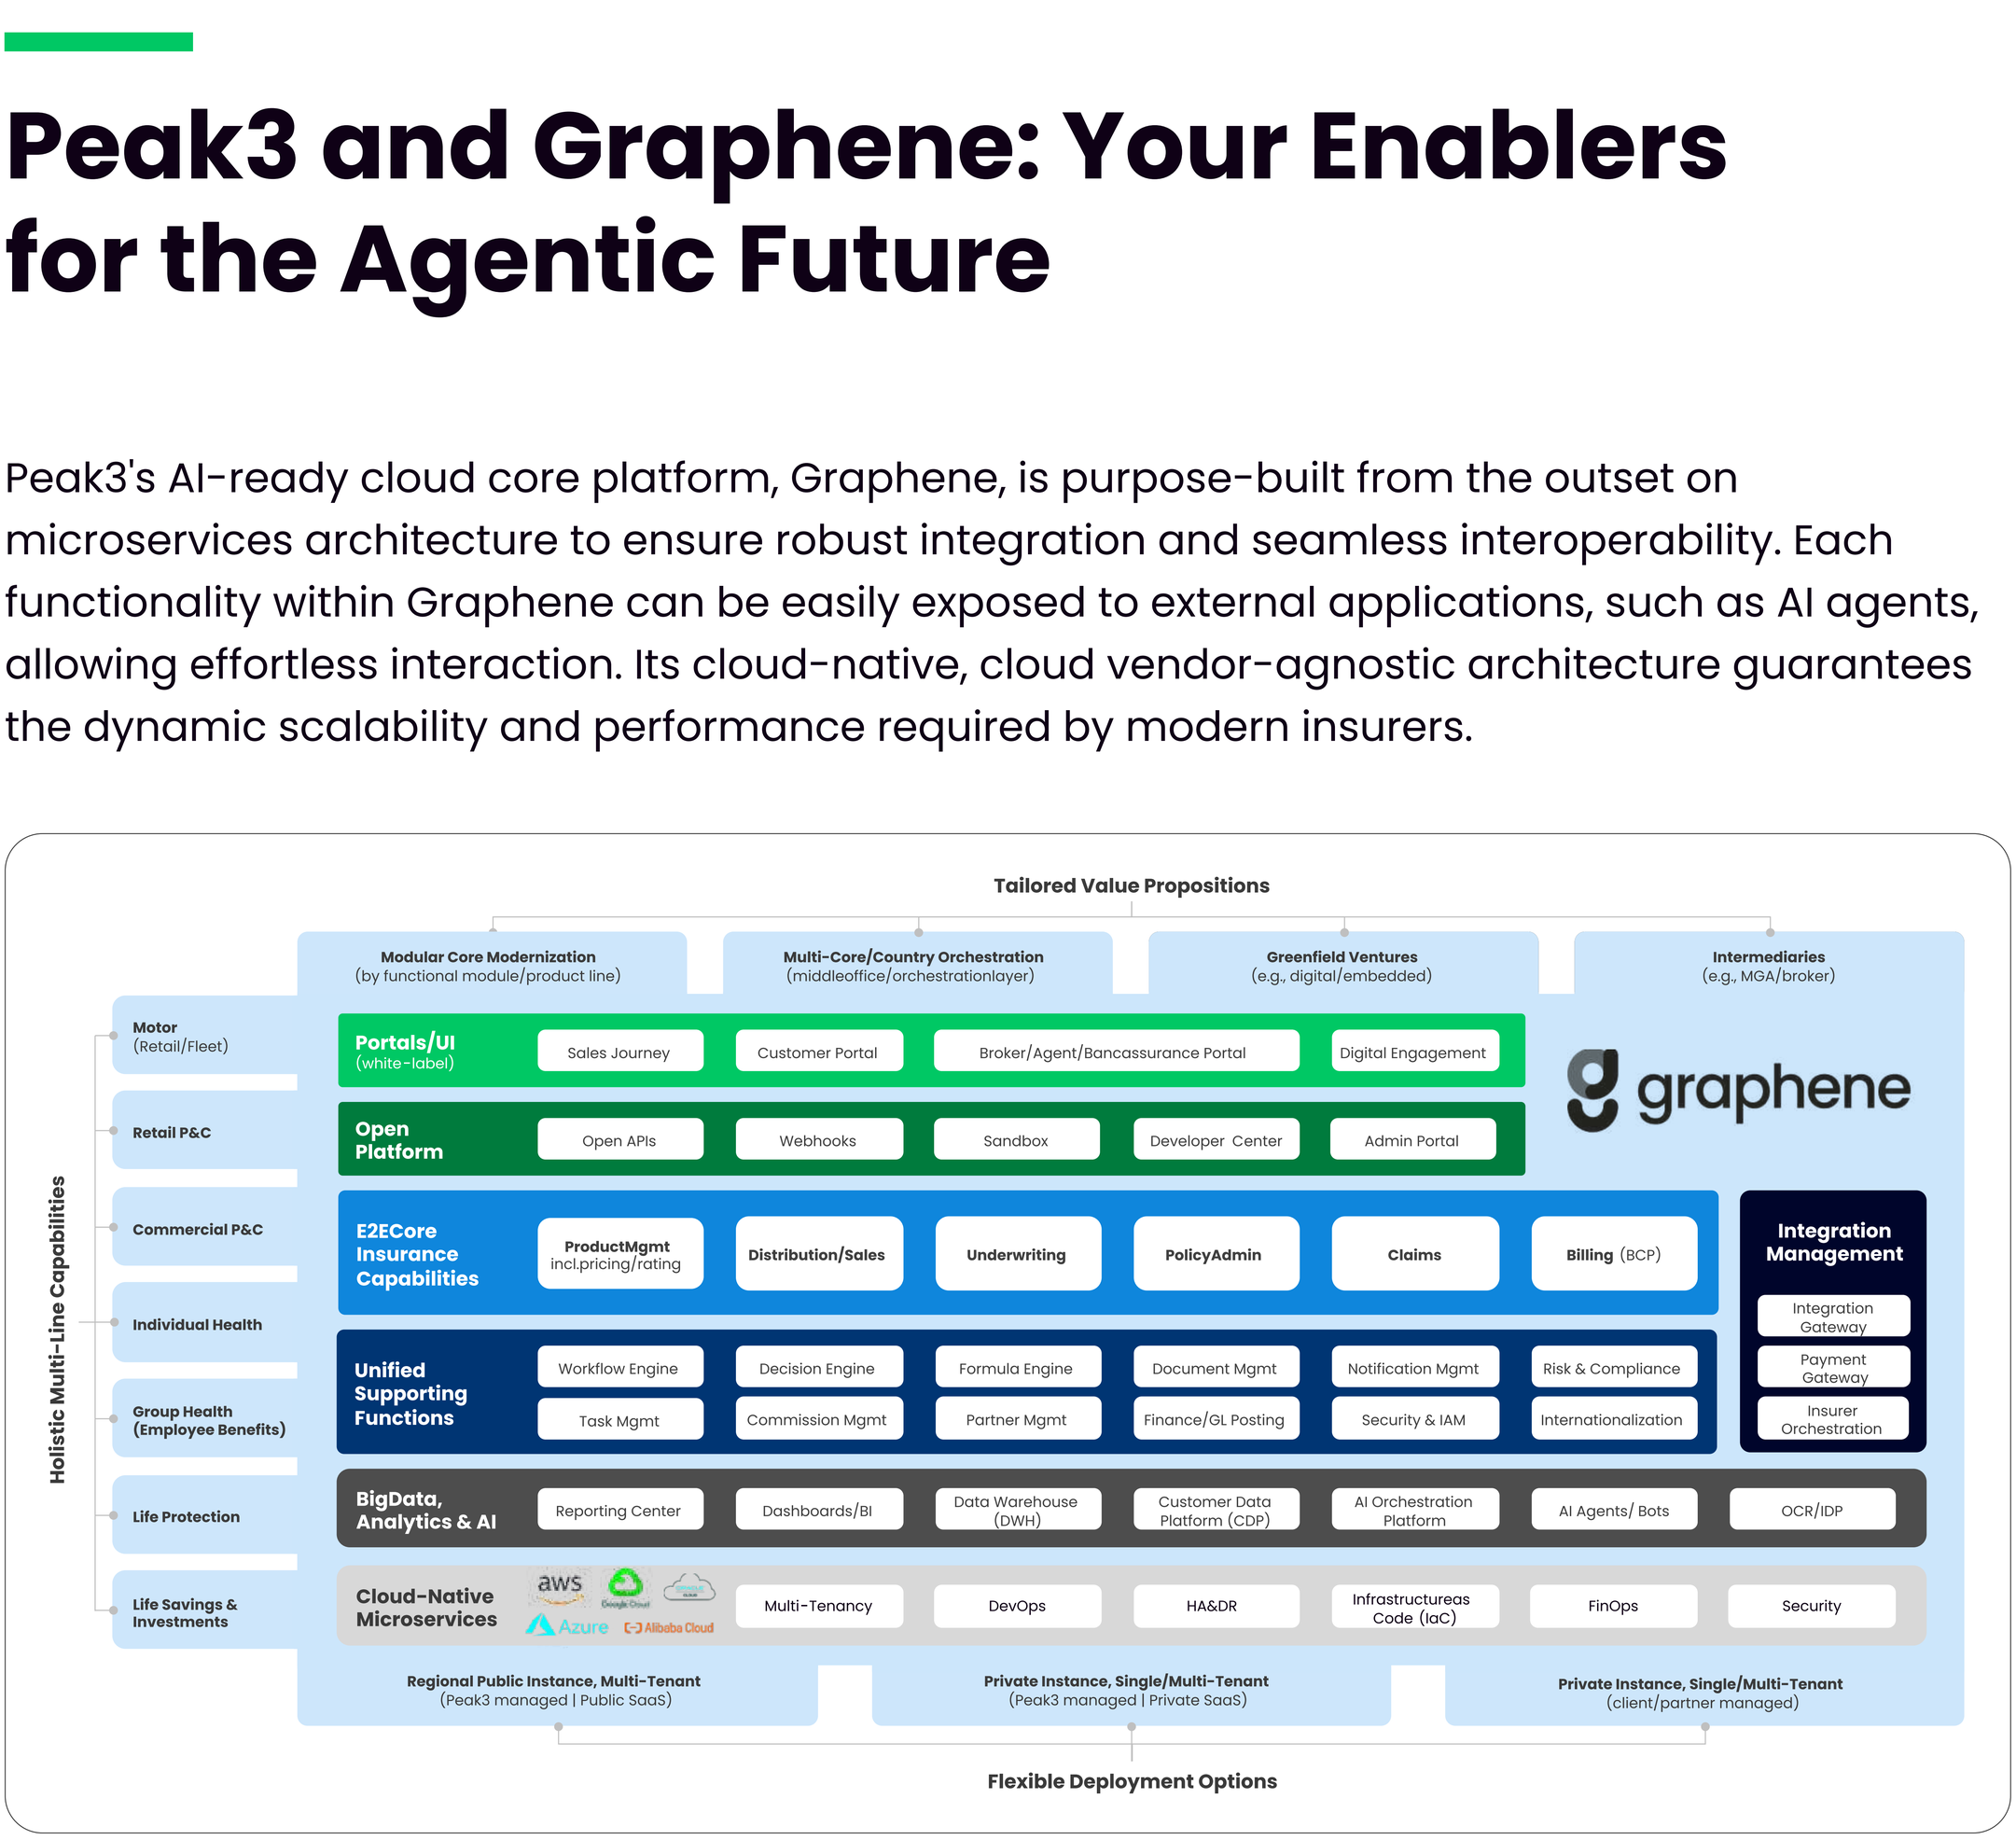

ILLUSTRATION 3: Simplified overview of Graphene by Peak3

Enabling AI-First Operations Across the Insurance Value Chain

Graphene's flexible workflow orchestration capabilities can adeptly manage processes driven solely by AI as well as those requiring human-in-the-loop interventions. Graphene's advanced and flexible product engine facilitates deep personalization not only in customer interactions but also in product development and design.

Security and compliance are embedded at every level, with stringent security protocols and granular authority management ensuring that AI agents can operate within controlled, observable, and compliant environments.

Graphene's built-in big data infrastructure, along with an optional customer data platform, ensures data quality and unified access across systems. Its end-to-end but modular capabilities span the entire insurance value chain—across Property & Casualty, Life, and Health lines—streamlining data governance and standardizing taxonomies without relying on disparate data silos and inconsistent definitions.

Graphene includes a built-in AI agent orchestration platform, allowing insurers to centrally configure and manage AI agents, integrate with foundational models, manage RAG pipelines, and so on. In addition, Graphene provides pre-integrated third-party and Graphene-native AI capabilities such as fraud, waste and abuse detection, intelligent document processing, and intelligent chatbots.

Peak3's ambition is clear: to establish Graphene as the intelligent, interconnected core platform enabling insurers to scale AI-first business models effectively. We continually enhance Graphene to support robust AI-first and agentic operations.

Ready to build your AI-first and agentic insurance operations and business models? Please reach out to us here or at hello@peak3.com.

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.

Peak3

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

To access our latest news and content, please follow us on LinkedIn.