Agentic AI in insurance: From theory to practice

nashtechglobal.com

Delivering technology **excellence**

Insurtech UK

Nash Tech.


FUJITSU


GenAirate
TECHNOLOGIES


Capital
Law

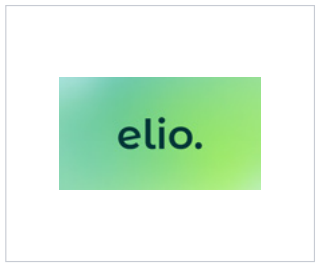
elio.

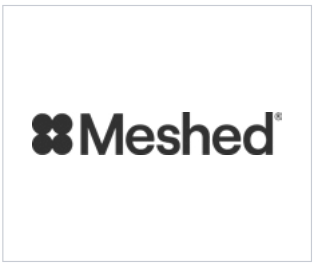
Meshed


mosaic
INSURANCE


Sønr


WPA
Health Insurance

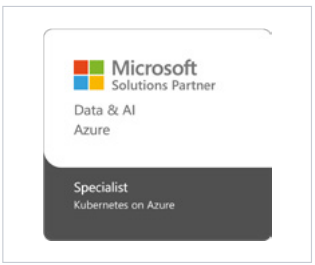
Microsoft
Solutions Partner
Data & AI
Azure
Specialist
Kubernetes on Azure


yu life


DICEUS


Hastings
DIRECT


eip


SAGEIFY


ManyPets


APOLLO
A SKYWARD GROUP COMPANY


sprout.ai

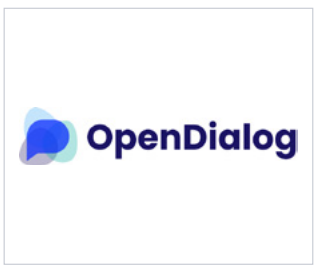
OpenDialog


quivaworks

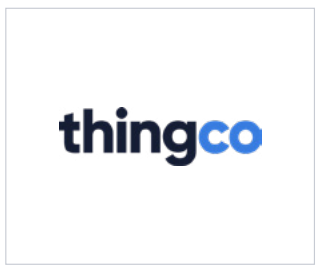
thingco

Our contributors' views on agentic AI in insurance

"It's so early in the journey... experts aren't that expert. And no one's as expert at our customer base as we are. The technology is learnable... the earlier you start the AI journey, the more you learn, and that's how you create an advantage in the marketplace."

David Mitchell-Dawson
Associate Director of Product & Innovations, EIP Limited

"If you've got a complex, clunky manual process, don't AI it. Go back to the drawing board and rethink the process first."

Mike Downing
CIO, WPA Health Insurance

"I don't think it's a technology failing. I think the stumbling block is governance: getting the due diligence and adoption frameworks right."

Hugh Coughlan
CTO Data & Applied Intelligence, Fujitsu

"As we've gone through this process, it's clear that everything is still being worked out in the agentic AI space, there really is no standard yet, and integration remains a challenge."

Dan Gordon
CTO, Sonr

"One of the first agents we're working on is an agent focused on data quality as part of a 'self-improving' system... we're using AI to do the AI... you cut diamonds with diamonds."

Usha Badrinath
Chief Data Officer, Mosaic Insurance

"The AI is the glue between all the points that currently require two or three different SaaS products."

James Fogarty
Co-founder, Elio

"Agentic AI will be our best friend if we have the appropriate understanding and controls that don't allow it to become our worst enemy."

Simon Liste
Head of Core Technology, Apollo

"The early adopters are not necessarily going to win the race. The leaders will be the ones who understand when to adopt new technology without losing focus."

Angelos Charitidis
Data Science Senior Manager, Hastings Direct

NashTech x Insurtech UK agentic AI working group

This white paper is the output of NashTech x Insurtech UK's agentic AI working group. The working group was established as a dedicated forum for senior insurance executives to explore practical applications of agentic AI. Its purpose is to move beyond conceptual discussions and focus on developing safe, effective, and value-driven use cases. This is not a venue for product demonstrations or feature promotion; rather, it serves as a collaborative space to validate assumptions, exchange insights, and share evidence of what is delivering results and what is not.

Through a range of roundtables, webinars, and individual interviews, we engaged senior leaders from across the insurance value chain, including insurers, MGAs, brokers, insurtechs, and technology partners. Discussions covered critical functions including underwriting, claims, fraud, finance, data management, and innovation. Contributors include leading organisations such as Mosaic, Hastings Direct, YuLife, Laka, WPA, among others. We also incorporated practical insights from the broader partner ecosystem, including Microsoft.

Throughout the working group, a clear pattern emerged: while interest in agentic AI is widespread, confidence in scaling solutions within a regulated, customer-trust-driven environment remains limited. Most initiatives are still confined to controlled pilots, often with an internal focus, because the most significant challenges are not technical. They centre on accountability, governance, explainability, and operational risk.

This white paper consolidates the insights gathered, the debates held, and the lessons learned. It is designed for decision-makers who require a clear and actionable plan.

Introduction: Context of AI in insurance

The insurance industry is no stranger to artificial intelligence. For years, machine learning models have been foundational tools for risk assessment, claims processing, and operational efficiency. Yet the recent emergence of agentic AI represents a significant new frontier, a paradigm shift that is creating as much uncertainty as excitement for industry leaders. This is a fundamental evolution in how we conceive automation and decision-making.

Recent data reveals that insurance companies are now operating at adoption levels comparable to technology, media, and telecommunications firms, a remarkable achievement for a traditionally conservative sector.

This acceleration masks a critical paradox; **while AI pilot programmes are universal:**

74% **of insurance companies reporting active initiatives**

Only 7% **have successfully scaled these innovations across key functions**

This **'pilot purgatory'** represents perhaps **the sector's greatest strategic challenge in 2026.**

The NashTech x Insurtech UK working group on agentic AI reinforces these findings through qualitative insights that illuminate the underlying causes. Many insurers within the working group are running internal pilot programmes with AI agents, yet virtually none have let these agents interact directly with customers. This cautious approach is not due to a lack of interest or capability; it stems from the absence of clear regulatory guidance and the high stakes of customer trust.

The absence of clear governance frameworks for agentic AI systems creates institutional paralysis. Senior executives repeatedly cited uncertainty around accountability, explainability requirements, and regulatory expectations as primary barriers to scaling beyond the pilot stage, transforming what should be technical deployment challenges into existential questions of organisational risk tolerance.

> "There is a lack of robust frameworks for tracking AI decisions, which is critical for regulatory compliance. Insurance has a bad reputation. We are perceived as the worst sector in the country. For example, with banking, if you are rejected for a loan, you wouldn't tweet about it or put posts out that you've had a bad experience with your bank, but you do with insurance."
>
> **Jonathan Valentine**
> *Co-founder and CIO, Thingco*

Insurers have been cautious, taking advantage of AI where it's reliable but holding back where outcomes are harder to explain or justify.

All participants in the working group observed that while insurance remains a 'human business' and speed and efficiency are vital, empathy ultimately outweighs the cost in customer interactions.

Before exploring the applications and navigating the risks, however, it is crucial to establish a clear and shared understanding of what agentic AI truly is and how it differs from the technologies that came before it.

Agentic AI – evolution and adoption in business processes

Strategic clarity begins with a shared vocabulary. In today's landscape, terms like AI, generative AI, and agentic AI are often used interchangeably, creating a fog of confusion that hinders effective planning and investment. Establishing a consensus-based definition is the first step toward building a coherent strategy.

Based on insights from the workshops, the recent history of AI can be understood as a clear technological progression:

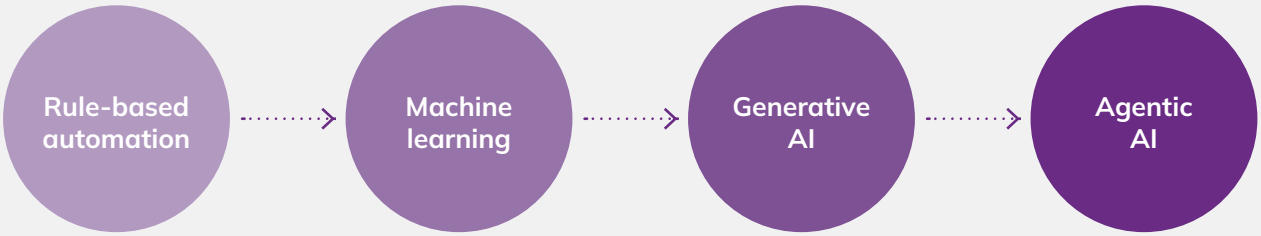

While each step builds on the last, agentic AI introduces a distinct new capability: autonomy in service of a goal. To clarify this distinction, it is helpful to differentiate between individual tools and the overarching system.

Aspect	AI Agent (The tool)	Agentic AI (The system)
Orientation	Reactive and task-oriented: Responds to specific, well-defined commands	Proactive and goal-oriented: Aims to achieve a broader objective
Execution	Single-step execution: Typically performs a single, discrete task	Multi-agent orchestration: Can coordinate multiple specialised agents
Autonomy	Low autonomy: Follows instructions precisely	High autonomy: Can revise its own instructions

Dan Henry, Strategic Account CTO, Microsoft UK, further contextualised this by describing an 'agentic spectrum', outlining a practical, three-phase path to adoption that organisations can follow.

1 **Agent as an assistant**

"The journey begins with individual productivity tools like Microsoft Copilot, which augment the capabilities of a single user.

---

2 **Human-led agents**

The next phase involves delegating specific, discrete tasks within a larger workflow to an AI agent, which then reports back to a human operator.

---

3 **Human-led, agent-operated**

The most advanced phase envisions an integrated system where automated, agent-driven flows are the default."

Why now?

Several trends have converged to make agentic AI feasible and attractive:

The explosion of digital data and APIs in insurance means many tasks such as quotes, claims and service requests, can theoretically be handled end-to-end by software.

Advances in AI models, especially LLMs, enable understanding of unstructured text and reasoning in ways that earlier rule-based systems could not.

The pressing business needs to 'do more with less'. Talent is hard to find, and operational costs are under pressure; this drives interest in AI that can automate complex work, not just trivial tasks.

The ease of availability of supercomputer level compute power and memory provided by modern computer architectures, GPUs, and enterprise grade networking enabling the processing of very large data sets.

Practical considerations for implementing generative and agentic AI in your business

Moving from a conceptual understanding of agentic AI to real-world implementation requires a structured and deliberate strategy. A robust framework is not just about choosing the right technology; it demands that organisational, legal, and risk dimensions are considered from the very outset.

Robert Stenzel, Senior Business and Technology Manager, NashTech, presented a five-phase framework designed to guide organisations through this process in a controlled and methodical way:

1 Prepare

This foundational phase involves establishing the problem context by identifying the business challenge, understanding/mapping the existing process and identifying constraints/challenges as well as success criteria. Crucially, it must include an early and thorough review with legal, compliance, IT security and risk teams to set guardrails, clarify non-negotiables, and avoid pursuing use cases that are fundamentally not viable.

2 Enable

Build a shared baseline understanding across business, technology, and risk stakeholders. This includes educating teams on autonomous workflows, tool capabilities and limitations, governance implications, and typical failure modes. Enablement is not treated as a one-off event but as an ongoing thread that continues through experimentation and delivery.

3 Evaluate

Identify, assess, and prioritise candidate use cases based on strategic relevance, risk exposure, and a clear value hypothesis. At this stage, an initial business case is formed, including expected benefits, cost drivers, and operational implications, sufficient to justify moving into experimentation.

4 Engineer

This phase involves designing and building a focused Proof of Concept (POC) to validate the technical feasibility, boundaries for autonomy, integration complexities, and control mechanisms. In this stage, we are testing value assumptions, technical feasibility, and risk hypotheses. Legal, compliance and risk functions are continuously involved as the AI system's behaviour becomes observable.

5 Decide and mobilise

Based on the results of the POC and a refined, comprehensive business case, the organisation can make a final, evidence-based go or no-go decision. A concrete mobilisation plan is required, covering deployment, scaling, operating model changes, governance, monitoring, and continuous improvement, to ensure the capability can be safely and sustainably operated in production.

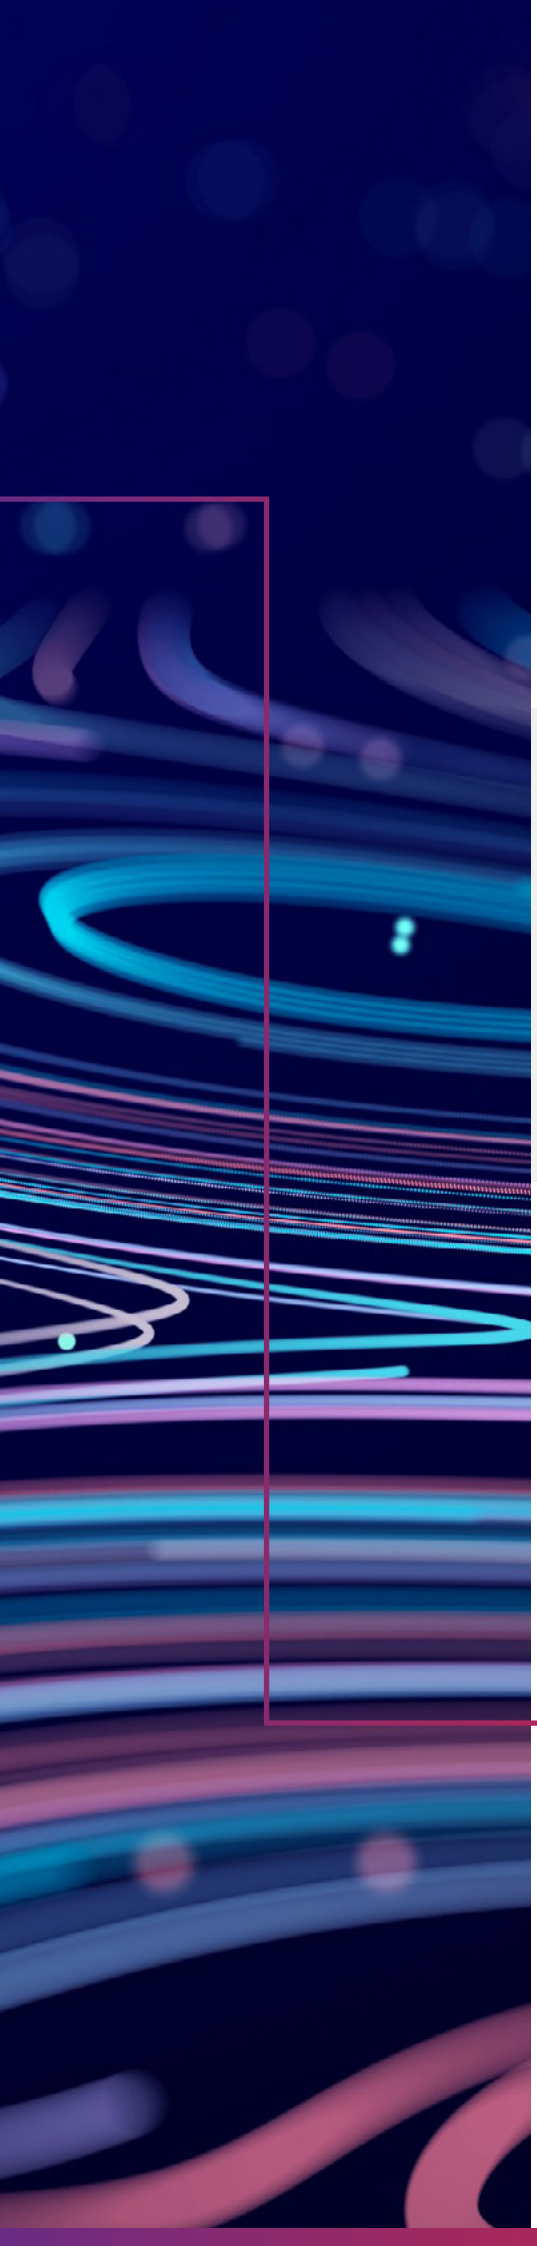

Alongside this process framework, teams must also navigate key technology choices. Thomas Pointer, Senior Technology Consultant, NashTech, explained how each approach involves a distinct trade-off between speed of response, accuracy, cost and the ability to update with recent data.

Prompt engineering

The simplest and fastest approach involving carefully crafted instructions to a base model using publicly available data.

Retrieval-augmented generation (RAG)

This technique allows a model to access and read proprietary data sources, leading to higher accuracy using generic model knowledge on your data sources. There are additional costs each time you run the model in terms of compute (price) and speed of response, but new data can be fed into the solution, ensuring up to date data.

Fine-tuning

The most complex and costly option involves retraining a base model on a specific dataset to change the models underlying behaviour. This has the advantages of speed of response when running the model, but it's difficult and costly to retrain the model, and you cannot easily update the model with the latest data.

Lessons learned from practical experience in pilot projects

The most valuable insights into agentic AI adoption come not from theoretical frameworks, but from the candid experiences of those on the front lines of implementation. The shared challenges from leaders at firms such as Legal & General and Hastings Direct reveal a set of recurring, non-technical hurdles.

The data hurdle

Before an agent can act, it must learn. The foundational problem for many organisations is the state of their data, it is often fragmented, of inconsistent quality, and locked away in siloed systems.

Simon Clayden, Principal Business Architect, articulated this challenge:

"One of our first challenges, though, was where is it? Where is all that information... and it is everywhere."

Without a coherent data strategy, any AI initiative will struggle to get off the ground.

The governance and trust hurdle

In a highly regulated industry like insurance, autonomy must be paired with accountability. Early adopters stress the critical need for systems to be explainable, auditable, and equipped with safe 'off-ramps'. When an AI agent makes a mistake or gets stuck, there must be a clear and reliable process for escalating to a human handler.

The internal bottleneck hurdle

Often, the biggest obstacle is not the technology itself, but organisational inertia. One member of the working group shared how his team, eager to test a simple use case, was told by internal security and network teams that they lacked the bandwidth to support it. This serves as a reminder that even with executive support, internal resource constraints can bring projects to a halt.

The skills gap hurdle

Successfully implementing agentic AI requires a new blend of talent. The demand goes beyond AI engineers to include roles like prompt engineers and business architects. This specialist talent is scarce, creating a significant resourcing challenge.

Risks and blockers: Regulation and adoption inertia

Proactively identifying and mitigating risk is not an obstacle to AI adoption; it is the prerequisite for successful and responsible implementation. Building trust with regulators, customers, and internal stakeholders require a clear-eyed assessment of the potential pitfalls from day one.

Legal and regulatory landscape

The regulatory environment for AI is complex and evolving rapidly. An **OECD September 2024 paper surveyed** 49 jurisdictions to map AI regulation in finance, revealing a landscape characterised by varied approaches. Regulators employ both outcomes-based approaches, (such as the UK's principles-based model) and rules-based approaches, (such as the EU's AI Act).

From a supervisory perspective, regulatory expectations remain consistent regardless of the definitional debate. What matters most is that financial institutions understand the risks and have proper governance, risk management, and control tools in place. The **FCA's September 2025** approach is explicitly principles-based and outcomes-focused, allowing flexibility to adapt to technological change. Critically, the FCA does not plan to introduce specific new regulations for AI; instead, existing frameworks will be applied.

Legal and regularity landscape

The legal and regulatory landscape emerged as one of the most debated topics within the working group. Executives consistently highlighted that the pace of technological innovation is outstripping the evolution of governance frameworks, creating uncertainty for organisations deploying AI. Many discussions centred around how compliance obligations, liability concerns, and ethical considerations were becoming critical factors in shaping adoption strategies.

Regulatory and legal accountability

The question of liability is paramount.

"Ultimate responsibility must remain with a human. This is not optional; it is a foundational principle that regulators expect. When an AI agent decides, accountability must be traceable to a human decision-maker. Organisations must ensure that AI agents can be overridden, that decisions can be audited, and that escalation to human handlers is built into the system from the outset."

Charlotte Gregory
Partner, Capital Law

The 'black box' problem and explainability

For many AI models, the internal logic behind a decision can be opaque. This 'black box' nature presents a significant challenge in regulated processes like fraud detection and financial crime prevention, where every decision must be transparent, explainable, and auditable. The regulatory expectation is straightforward; if an AI system makes a significant decision affecting an individual, that person must be able to understand why, contest the decision, and seek human review.

Data, copyright, and intellectual property considerations

The use of data to train and operate agentic systems introduces multiple legal dimensions. The Data Use and Access Act 2025 (DUAA) and Information Commissioner's Office (ICO) guidance on automated decision-making will shape what organisations can and cannot do. One critical emerging issue is copyright. Computer-generated works are currently under review following the UK Government's Copyright and AI consultation, with precedent being set in cases such as Getty Images versus Stability AI, where it was found that GenAI models do not store or reproduce images, and therefore do not infringe copyright. Organisations must clarify whether copyrighted materials are being used in training data, what IP exists in any outputs, and what confidentiality obligations apply to proprietary data.

Disputes and litigation risks

Early case law is establishing precedent that will shape liability. In **Walters v Open AI LLC**, defamation was examined through the lens of AI-generated content. More broadly, what an AI system says to customers becomes part of public information and can create enforceable liability, even if made in error **(Moffatt v Air Canada)**. Professional negligence claims involving AI decisions are emerging in case law, such as **Ayinde v London Borough of Haringey**, where fake (potentially hallucinated) case references were submitted to the court.

Employment law and ethical use

Organisations deploying AI must develop robust policies around ethical use in the workplace. This extends beyond recruitment, where algorithmic bias can constitute unlawful discrimination, to encompass duty of care obligations to employees. The expectation is that organisations will work within regulatory confines on compliance and reputation management.

Adoption inertia and internal politics

Organisational resistance often proves to be a greater challenge than the technology. Overcoming this requires securing internal champions, educating stakeholders, and reframing the initiative around business process transformation.

Data protection and security

Giving an AI agent access to sensitive data introduces new security vectors. Robust data protection protocols and security reviews must be implemented before a proof of concept begins, not as an afterthought.

Use cases identified through the insurtech workshops

The true value of agentic AI becomes clear when theory gives way to practical application. Four high-impact use cases consistently emerged during the workshop sessions.

Streamlining underwriting data triage

Problem: Underwriters are inundated with high volumes of unstructured submission data from brokers, arriving in inconsistent formats. Average lead time to detection (hours/days)

Solution: An agentic system that ingests broker emails and attachments, automatically structures the data, and uses historical conversion data to prioritise opportunities.

Impact: Could reduce manual processing time by 70%, cutting average quote turnaround from ten days to three days.

> "Customers don't see our AI, they just feel it. There's no chatbot to speak to on our website. You upload your documents and our agents extract the data behind the scenes. Because AI handles the background work, we don't need the same commission structure as traditional brokers, and customers see that reflected in lower costs."
>
> **Jake Wells,** *Co-owner, Meshed*

Accelerating fraud escalation to MGAs

Problem: The manual process of assembling a fraud report for escalation is slow.

Solution: An agent that automates fraud report generation, assembling all relevant data into a structured narrative with human handler review.

Impact: Could significantly reduce claims settlement times, which can be extended by over 100 days during underwriter escalation.

Optimising intermediary and broker operations

Problem: Significant time is wasted on "admin ping-pong", repetitive back-and-forth communication to gather information.

Solution: An agent acting as a 'trust but verify' checkpoint for broker submissions before they reach the underwriter.

Impact: Fewer hand-offs, cleaner submissions, faster approvals, and a more efficient relationship between brokers and underwriters.

Enhancing customer service with real-time knowledge

Problem: Junior agents often need to place customers on hold whilst they seek answers from senior staff.

Solution: A customer service agent powered by Retrieval-Augmented Generation (RAG) connected to verified internal knowledge bases.

Impact: Dramatically improves response times, increases first-call resolution, and empowers agents with real-time expertise.

Elements to consider for the investment case in agentic AI

A successful investment case must articulate value far beyond simple cost savings. It needs to align with broader strategic goals and address compelling business reasons for change.

As one workshop participant asked, "Just because an agentic agent could do it, does it actually mean that you should do it?" The answer must be grounded in a robust business case.

1 Align with strategic vision

The most persuasive proposals frame agentic AI as not a technology experiment, but as a critical enabler of the company's core strategy. **Coco Co, Head of Finance Transformation, Hastings Direct**, provided a powerful example where the finance department's AI project was positioned as essential to achieving the company's vision to "become the biggest digital insurer out there." This elevated the conversation from a departmental efficiency project to a strategic imperative.

2 Focus on scalability

A core argument for agentic AI is its ability to break the linear relationship between growth and headcount. The business case should demonstrate how agentic systems allow revenue and transaction volumes to increase without proportional increases in operational costs and full-time employees (FTEs).

3 Quantify with key metrics

Whilst strategic alignment is crucial, the business case must be supported by tangible, quantifiable metrics:

- Book size managed per agent
- Renewal rates
- Quote-to-bid ratio
- Claims processing time and settlement costs
- First-call resolution rate

4 Define the 'cost of inaction'

Perhaps the most powerful argument (shared by **Dan Henry, Strategic Account CTO, Microsoft UK**) is the one that addresses the 'cost of inaction'. Leadership approval is often unlocked not by the promise of gain, but by the fear of being left behind. A strong business case articulates the competitive risk of failing to invest and build capabilities in an area where rivals are already moving forward.

Recommendations for pilot projects and investigations

The journey into agentic AI can seem daunting, but the collective advice from the workshop series provides a clear, actionable roadmap for launching a successful first pilot.

1 Start with a predictable, high-value problem

Resist the temptation to tackle the most complex, high-profile challenge first. The consensus advice is to select a back-office use case with clear, verifiable outcomes. Functions like 'finance and controls' were cited as popular and safe starting points because the rules are well-defined and results are easily audited.

2 Build a cross-functional 'Think Tank'

Do not let this be a siloed IT project. Create a small, highly motivated team that includes business, technology, legal, and risk stakeholders from day one. This ensures all dimensions of the project are considered from the outset and builds broad organisational buy-in.

3 Frame governance as an enabler, not a blocker

Treat governance not as a hurdle to be overcome, but as a valuable 'good blocker'. These teams can help stress-test the project's design, identify unseen risks, and build guardrails that protect the business.

4 Embed human oversight by design

Human-in-the-loop is non-negotiable, especially in early projects. Establish clear operating protocols for error handling and escalations to a human agent before the system is scaled, ensuring there are always safe off-ramps.

5 Lead with the business process, not the technology

Frame the initiative as a 'business process transformation' project. This shifts the narrative from complex technology to solving a tangible business problem. Agentic AI becomes the tool, not the end goal, making the value proposition easier for non-technical stakeholders to understand.

Summary:
From theory to transformation

The insights from the NashTech x Insurtech UK working group paint a clear picture: **agentic AI is not a distant, futuristic concept but an achievable and increasingly necessary evolution for the insurance industry.**

However, its successful adoption demands a pragmatic, collaborative, and governance-first approach. The journey from hype to implementation is not about a single technological breakthrough, but a methodical process of organisational learning, strategic alignment, and disciplined execution.

1. **Start small with focused use cases** that have measurable outcomes. Pursue early wins in back-office functions to build momentum and learn critical lessons in a controlled environment.

2. **Prioritise legal, risk, and compliance guardrails** from the very beginning. Treat governance as a foundational element of project design that enables responsible innovation.

3. **Build a business case that focuses on strategic value and scalability,** not just efficiency. Align projects with the organisation's core vision and articulate how agentic AI can enable profitable growth and responsible innovation.

4. **Foster a community of practice** to share playbooks. The challenges are universal. Organisations that collaborate will learn faster and avoid "reinventing the wheel in silos."

This white paper is a product of that collaborative spirit. It is intended not as a final word, but as a starting point. The time for passive learning is over. Your task now is to assemble your pilot team, select your first use case from the high-impact examples provided, and begin the work.

The future of the industry will be shaped not by those who watch, but by those who build.

We are experts in technology, delivering smart solutions that solve business challenges and create value. Our award-winning teams apply deep expertise and passion to deliver complex IT projects globally.

Nash Tech.

For more details, please send your enquiry to **info@nashtechglobal.com** or visit our website **www.nashtechglobal.com**

Who are Insurtech UK

Insurtech UK are the voice of the UK's insurtech ecosystem-a thriving community that fosters innovation, builds connections and creates opportunities. Their mission is to transform the insurance industry through technology and position the UK as the global leader in insurance innovation.

Visit **www.insurtechuk.org** for more information