tinubu

From buzz to boardroom:

how agentic AI redefines specialty insurance strategy

Whitepaper I 2025

Author

Yvan Saule

CTO. Tinubu

Yvan is the **Chief Technology Officer** (CTO) at Tinubu, driving the advancement of the industry's leading Insurance SaaS platform.

With over a decade of senior leadership experience as a CTO and VP of Engineering in global organizations, Yvan specializes in **spearheading digital transformations** and **scaling technology infrastructures** to deliver innovation and transformative solutions for Credit Insurance and Surety leaders worldwide.

Known for fostering agile, fast-paced environments, he has a proven track record of enhancing operational efficiency, driving technological advancements, and leading teams to develop market-leading solutions.

About Tinubu

Tinubu is the leading provider of enterprise software solutions for the global specialty insurance industry, driving the industry's digital transformation and growth with best-in-class cloud technology and deep domain expertise. Tinubu provides end-to-end solutions across the entire insurance value chain, including policy or bond administration, underwriting, distribution, and claims. Tinubu brings 25+ years of sector experience and assets across carriers and distribution.

Introduction

AI stole the headlines; agentic AI is changing how the work gets done.

In specialty insurance, value doesn't come from another chatbot, it comes from small, specialized agents that can **ingest, retrieve, reason**, and **act** inside governed workflows.

Think: prep work done in minutes, explainable first-pass decisions, and clean hand-offs to humans. That's the shift we're seeing on real books of business: underwriting triage that actually sticks, claims intake that eliminates re-typing, and legacy migrations that move from months of discovery to days of clarity.

> **"Agentic AI doesn't replace specialists, it removes the noise so judgment shows up faster." - Yvan Saule**

Why now?

Volatility is up, margins are tight, and regulators expect traceability. Teams need speed without losing judgment. Agentic AI delivers that by coordinating multiple "doers", including "judge" agents that verify facts, enforce policy constraints, and escalate edge cases. It doesn't replace specialists - it **removes the noise,** allowing them to apply judgment where it matters most, and unlocks new business (early-warning signals, contract intelligence, embedded decisioning) that older tools couldn't touch.

The catch?

Once everyone uses AI, the technology itself becomes a new risk surface. This paper is about both sides, how to deploy agentic AI that performs in production and how to govern it so it stays trustworthy.

Executive summary

Agentic AI represents more than incremental automation, it marks a step-change in how specialty insurers run their core operations.

Traditional tools like RPA, OCR, and chatbots have delivered efficiency gains at the margins. By contrast, agentic AI **re-architects** underwriting, claims, and policy migration workflows, enabling a **fundamental transformation** of operating models.

The real value extends beyond cost reduction.

Early adopters report 40-60% cycle-time improvements, but the deeper impact lies in competitive advantage: faster product launches, embedded intelligence in distribution, and differentiated customer experiences. Agentic AI shifts AI from being a cost play to becoming a growth and differentiation lever.

Even more, it creates revenue.

Think embedded decisioning that binds simple risks instantly in partner portals, portfolio early-warning signals packaged as paid "risk alerts," contract intelligence that enables modular covers and faster bespoke endorsements, and API-ready co-pilots that open new distribution lanes.

Scaling this technology requires as much focus on governance as on innovation. Agentic AI is no longer deployed at the edges, it is embedded in core systems.

To pursue these safely and build long-term trust, **companies will win by establishing industrialized governance, and building the preventions** into the same stack: provenance and explainability by default, human-in-the-loop thresholds for material exposure, red-team/adversarial testing, segregation of duties across build–run–monitor, and circuit-breakers for agent-to-agent flows.

tinubu

Table of contents

tinubu

1

Why Agentic AI and why now?

In the past year, generative AI hype has dominated headlines, but the real revolution for insurers goes beyond chatbots that just talk. It's about AI that can act.

As CTO of Tinubu, the SaaS leader in specialty lines (trade credit, surety, A&H, etc.), I've had a front-row seat to this shift. We're moving from buzzwords to business value, deploying "agentic AI" at scale in the real world of specialty insurance.

Agentic AI refers to AI systems with agency: they autonomously **reason, plan**, and **execute** tasks through multiple specialized agents working in concert, rather than a single all-in-one model. This is not just theoretical, it's already unlocking tangible improvements in underwriting, claims handling, and legacy modernization. And it's positioning tech/software companies, which have built AI-ready platforms for complex insurance products, to lead the way.

Simply put, insurers face growing pressure to do more with less, faster answers for customers, deeper insight for underwriters, and leaner operations. Traditional AI (like a lone predictive model or a FAQ chatbot) can help at the margins, but it doesn't fundamentally change the game. Agentic AI, on the other hand, promises autonomous "co-workers" that can handle entire workflows under human-defined rules.

No wonder Gartner named agentic AI a top tech trend for 2025, describing autonomous machine agents that go "beyond query-and-response generative chatbots to do enterprise tasks without human guidance". In insurance terms, that means an AI that doesn't just advise an underwriter, it can gather data, analyze risk, draft a quote, and even take first action on routine cases, handing off to humans only when needed. It's a leap from hype to hands-on reality.

In fact, a recent Capgemini study estimated agentic AI could create up to **$450 billion** in economic value by 2028 and notes 14% of organizations have agents partially/fully deployed. And insurers aren't standing still: **20%** of insurance organizations are already piloting AI "agents", though only **4%** fully trust them yet. Clearly, the potential is massive, but so is the work to be done to realize it.

From Single-model To Multi-agent Intelligence

To understand why agentic AI is more than marketing spin, consider how it differs from the status quo.

A traditional AI solution is often a single-model interface: you feed it an input and get an output along one pipeline. It might be a pricing model, a fraud flagging algorithm, or a chatbot answering FAQs. Useful, yes, but inherently limited. Complex insurance processes like underwriting or claims involve many steps and skills.

A single model tends to operate in isolation, performing one narrow task at a time. In contrast, agentic AI deploys multiple specialized agents that can collaborate and even make decisions for you. Each agent is like a specialist on a team, and an orchestration layer coordinates them. Some agents may be large language models fine-tuned for specific tasks (reading policy wordings, scanning medical reports), while others are rules engines or optimization algorithms. Together, they handle different parts of a workflow in parallel, sharing findings, cross-checking each other, and dynamically re-planning as needed.

The multi-agent approach isn't just academic, it directly addresses real challenges in specialty insurance. Multi-agent systems can **divide and conquer** complex problems, often more efficiently than any single model. For example, one agent can focus on financial spreading (reading financial statements for a credit risk) while another agent forecasts macroeconomic trends, and yet another checks policy compliance rules. Working together, they cover a broader base of knowledge and complete the task faster. This **team-of-AI** approach also adds resilience and kills the root of distrust hallucinations. When one agent drafts an answer, 'judge' agents verify it: a retriever checks facts and citations, a policy checker enforces constraints, and a self-consistency agent re-runs the reasoning via an alternate path/model. If sources don't line up or two paths disagree, the output is flagged, revised, or escalated ("no source, no action"). In essence, we're unbundling big AI into many little brains, and then re-bundling them in an intelligent workflow. The result is an AI that feels less like a static tool and more like a dynamic colleague that **autonomously learns, communicates, and acts** within its domain.

Crucially, agentic AI doesn't operate in a vacuum, it works within **guardrails** we set. In our deployments at Tinubu, we ensure each AI agent knows the boundaries (e.g. compliance rules, underwriting guidelines) and has an escalation path to a human for anything unusual.

This design addresses the biggest barrier to adoption: trust.

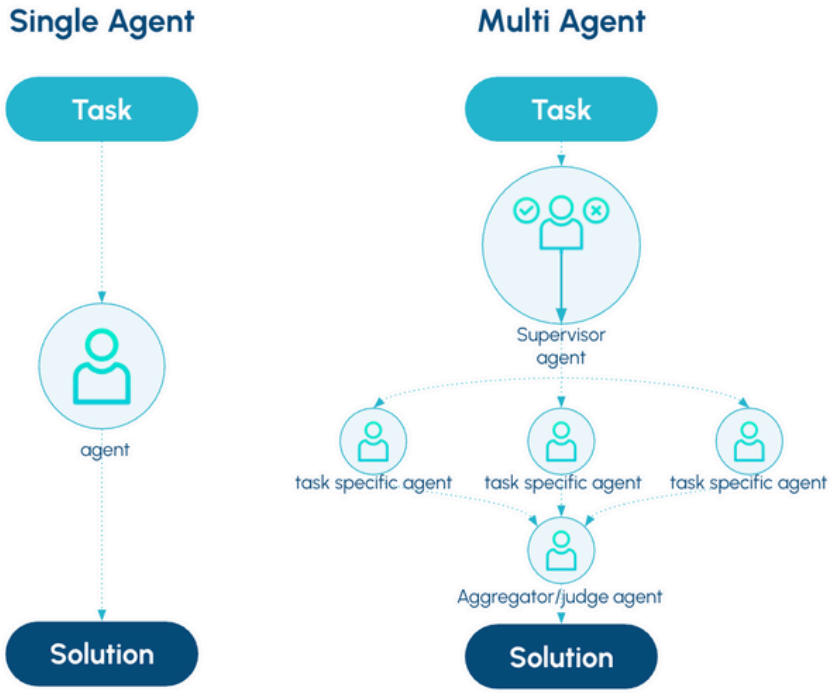

Visual Diagram 1: single agent AI vs. multi agent.

It's telling that fully autonomous AI agent deployments are still rare (only around 2% of organizations have scaled them so far) and trust in "AI with no human in the loop" has actually dropped in the last year.

The key is to strike the right balance: use agents to automate routine, data-driven tasks, but keep humans in charge of oversight and final high-stakes decisions. In fact, insurers today are prioritizing **human-AI collaboration** over full automation, integrating AI agents into workflows to assist underwriters, brokers, and claim adjusters rather than replace them.

This collaborative approach is already showing benefits: it can free up about 65% more time for teams to spend on high-value work, and drive higher job satisfaction as AI takes the grunt work.

In other words, when done right, agentic AI becomes an *amplifier* of human expertise, not a replacement.

We'll return to this point on ethics, but it bears stating up front: **Agentic AI doesn't cut jobs. It cuts through the noise and distractions, so humans can focus on what they do the best, and even start to address beyond the current status-quo**.

> **Speed without judgment is risk; judgment without speed is lost business. Agents give you both**

tinubu

2

Agentic AI at work: reshaping underwriting, claims, and the modernization journey

Agentic AI delivers value when multiple specialized agents are orchestrated around a governed workflow, not when a single model is bolted onto a legacy process.

Below, we outline how the stack operates in practice across three core domains, how we implement it at Tinubu, and where the market is already reporting measurable gains.

Underwriting, where Agentic AI fits and what it changes

Underwriting in specialty lines (e.g. cyber, marine, A&H) has long relied on expert judgment supported by tedious data gathering. GenAI is now reducing that grunt work and speeding up risk assessments without sacrificing rigor. Notable real-world applications include:

- **Automated Submission Triage:** QBE Insurance developed a Cyber Underwriting AI Assistant to help underwriters swiftly review broker submissions. Deployed in late 2023, this AI tool analyzes large cyber applications and flags key information. The impact was a **65% reduction in time** needed to review submissions, allowing underwriters to process more business and focus on complex risks. By mid-2024, QBE had scaled this assistant from North America to Europe and Asia-Pacific, demonstrating global applicability of AI in underwriting workflows.

- **Narrative generation for Underwriters:** Zurich Insurance piloted a generative AI platform for its U.S. Middle Market underwriters. The tool reads thousands of pages of exposure data, loss runs, and risk reports, then **generates a concise underwriting narrative in minutes.** In an initial rollout, 80% of submissions handled by a group of underwriters were processed with AI assistance, saving an estimated 60 minutes per submission on information gathering. Underwriters gained back valuable time to apply judgment and engage with brokers, accelerating quote turnaround without eroding diligence. As one Zurich underwriter noted, the GenAI summaries give a head start on each account, effectively **compressing what used to be 30–60 minutes of prep work** and improving response times to brokers.

- **Improved Risk Selection and Capacity Deployment:** These underwriting AI solutions don't replace human judgment, they enhance it. By cutting through submissions and surfacing

- key risk factors instantly, AI assistants allow underwriters to focus on decision-making and pricing. QBE reports that faster submission triage with AI has improved risk selection and unlocked growth opportunities by ensuring **100% of broker submissions get analyzed** (no more "no quote" due to time constraints). Zurich's team similarly found that productivity gains from AI led to **more confident decisions and better solutions for customers**, not just speed. In specialty lines where underwriting expertise is paramount, GenAI is thus acting as a co-pilot, expediting data digestion, ensuring no detail is missed, and ultimately enabling underwriters to deploy capacity more efficiently.

For Accident & Health, we've bundled that pattern into Underwriting Workbench: intelligent intake, real-time broker connectivity, and dynamic risk analytics in one place. The AI automates triage, data capture, and document analysis to compress submission prep; the workbench then carries the case through quoting, binding, endorsements, and renewals. Public write-ups note that the launch targets faster quote-to-bind cycles and underwriter productivity gains, with reported improvements of **up to ~65%** in submission review time for early users, evidence that explainable automation is landing in everyday underwriting, not just demos.

For proactive risk detection, signals before the surprises are always good. Behind the scenes, agents continuously connect to Tinubu data and external sources to pick up weak signals, adverse news, sector stress, sanctions changes, then cross-reference those with portfolio exposures. When something material moves, the system creates a targeted alert (buyer, industry, limit, or country) and routes a task to the risk analyst. What used to be buried in unstructured updates becomes a prioritized queue: earlier default indicators, automatic triage toward the most material alerts, and genuinely proactive exposure management.

We can triage and explain an account in minutes, so where else does explain-ability buy us trust?

> **If it's repeatable, agents do the heavy lift; if it's material, humans decide.**

AI-Augmented Claims: Faster Settlements and Enhanced Service

Claims handling is the moment of truth in insurance, and AI is helping carriers close claims faster while controlling costs and fraud.
In specialty insurance (from travel claims to trade credit and cyber incidents), GenAI and intelligent automation are slashing cycle times and improving the customer experience.

Prominent examples include:

- **Instant, Touchless Claims at Scale**: Digital insurer Lemonade famously uses AI bots to handle straight-forward property claims. Their claims bot "AI Jim" reviews claims, cross-checks policies, and even runs fraud algorithms automatically. As of 2021, Lemonade reported that AI Jim handled over 30% of claims with no human intervention, with an average settlement time of just 3 seconds. This lightning-fast processing translated to industry-leading customer satisfaction and set a benchmark for AI-driven efficiency. By trusting AI for low-complexity cases, Lemonade's team can concentrate on disputed or high-severity claims, aligning human effort where it truly adds value.

- **Travel Insurance Claims in Hours Instead of Weeks:** Allianz Partners, a global assistance and travel insurance provider, leveraged AI to overhaul its claims operations. The result was a **drop in average claims lifecycle from 19 days to just 4 days** after deploying AI, according to 2025 reports. Remarkably, **71% of claims are now closed within 12 hours or less,** with some payouts completed in as little as 6 hours. About 65% of incoming claims are deemed eligible for straight-through AI processing, which automatically validates and approves those claims so that human adjusters can focus on complex cases. This blend of intelligent automation and human oversight has not only cut cycle time by ~80% but also reduced manual workload and improved customer satisfaction. Faster payouts and proactive communication (including AI-driven chatbots for claim updates) have boosted Allianz's service quality while lowering operational costs.

- **AI Triage and Fraud Detection**: Even in more complex casualty claims, AI is aiding adjusters. For example, CLARA Analytics provides an AI-driven claims triage tool that uses predictive and generative AI to score claim severity and highlight subrogation opportunities. Such tools sift through medical notes and legal docs to flag high-risk claims or po

- tential recoveries, helping carriers reduce leakage and prioritize workloads. Similarly, many insurers use AI fraud algorithms that scan claims for anomalies, Allianz's AI approach has improved fraud detection accuracy, cutting fraudulent payouts by an estimated 20%. These capabilities reduce loss costs and ensure genuine claims get paid faster by removing bottlenecks like lengthy investigations for simple cases.

On claims, we stopped making experts wasting time retyping invoices. In our **Claims Portal** for trade credit, an intake agent now reads PDF/scanned invoices, extracts the fields that matter (amount, invoice and due dates, buyer name/address, VAT, reference), checks them against terms, flags gaps, and hands a clean, structured payload straight into the claim file, with an audit trail the whole way.

One **global carrier (unnamed)** called it a "game changer," and I get why: the workflow is calmer, reserves are set earlier, and adjusters spend time on judgment, not data entry.

As a result we've seen **up to 40 to 60% faster claim intake**, **higher first-time-right documentation**, and **lower back-office effort**, exactly the kind of practical lift that lets teams move faster without cutting corners.

Such improvements directly lower loss adjustment expenses and improve customer retention. Faster claims also translate to better reserve accuracy and cash flow management for insurers. In sum, AI-driven claims automation is proving its worth by **shortening cycles**, **cutting costs, and delighting customers**, all while allowing human examiners to concentrate on the claims that truly require their expertise

If we can compress claim cycles without losing control, why are core migrations still slow and risky?

Modernization & Migration: AI Agents Accelerating Legacy Transformation

Specialty insurers often rely on decades-old core systems and face complex migrations when modernizing underwriting platforms or policy admin systems.

Here, generative and agentic AI are emerging as powerful allies to **accelerate system migration, code refactoring, and data conversion**, essentially acting as force-multipliers for IT teams. Key use cases and outcomes include:

- **Legacy Code Conversion with AI:** A top global insurer turned to generative AI to speed up a core system migration, with striking results. According to a McKinsey case study, this Tier-1 insurer saw **code modernization efficiency and testing improve by** over 50% by using GenAI-driven tools, and achieved a >50% acceleration of coding tasks during the migration. AI helpers ingested old code, auto-generated new code in modern languages, and identified bugs, drastically reducing manual effort. In general, McKinsey has found that **AI-augmented legacy modernization can cut project timelines by 40–50%** and reduce technical debt costs by ~40%, while improving output quality compared to traditional methods. These gains mean insurers can upgrade core systems in months rather than years, with fewer errors and lower risk.

- **Agentic AI for System Integration:** Beyond code translation, agentic AI "co-pilots" are assisting with configuration and data mapping in complex migrations. For example, insurers are deploying AI agents that autonomously extract business rules from legacy systems, generate documentation and data dictionaries, and even produce test cases for the new system. This accelerates understanding of legacy intricacies and ensures nothing critical is lost in translation. As Accenture reports, 78% of insurance executives expect AI agents to **reinvent how they build digital systems,** by handling requirements, coding, testing, and deployment tasks autonomously. Early implementations confirm significant acceleration in **software development and deployment cycles** for insurers willing to embrace these AI-driven devOps practices.

- **Faster Product Launches with AI Co-Pilots:** Modernization isn't only about internal IT efficiency, it also enables business agility. Munich Re exemplified this by integrating a GenAI co-pilot into its **REALYTIX ZERO** underwriting platform. This AI assistant guides users through new product configuration in natural language. The payoff is that insurers (clients of Munich Re) can now **implement and launch a new insurance product in just hours or days**, a process that used to take months. By simply "prompting" the system with desired coverage features, underwriters can have the co-pilot generate policy wording, rating logic, and documentation automatically. This has **sped up product design impressively**, allowing quick adaptation to emerging risks and niche market opportunities. In short, AI-driven modernization tools are compressing time-to-market for specialty insurance products, which is a key competitive advantage in today's fast-changing risk environment.

- **Migration with Minimal Disruption:** Crucially, these AI-driven approaches to system migration are achieved without massive downtime. For instance, Allianz Partners' adoption of an "agentic process automation" for claims was done alongside legacy systems (via APIs) rather than ripping them out, enabling a smooth transition with a **90% automation rate** in target processes and unified operations across previously siloed systems. This underscores that AI can be layered on to **modernize piecemeal,** reducing risk compared to big-bang IT projects.

One of the most exciting "agents" we've built at Tinubu is an **autonomous systems analyst** for core migrations. Instead of sending an army of consultants to manually audit the legacy setup for months, we unleashed a set of AI agents to do the grunt work. First, an agent crawler scanned the legacy system's configuration, code, and database, automatically documenting existing products, workflows, and any unusual custom rules or workarounds in use.

It was able to read thousands of lines of config files and even legacy code, parsing out key entities and business logic. Next, an analysis agent evaluated these findings to detect potential risks and complexities, for example, it spotted that a certain bond formula was implemented in a non-standard way for a few areas, which could pose a risk if we migrated it blindly. It also identified duplicate or unused fields (artifacts from past projects) that could be cleaned up.

Finally, a planning agent synthesized all this into a **phased migration plan:** it grouped modules into what could be migrated in Phase 1 (low-risk, high-value quick wins), Phase 2 (moderate complexity), and which elements might need a full redesign or special attention. The plan even included suggested **sequencing** (e.g. migrate simpler commercial bonds first before complex contractor bonds) and highlighted where human expertise was required (such as validating the new calculations for that custom formula).

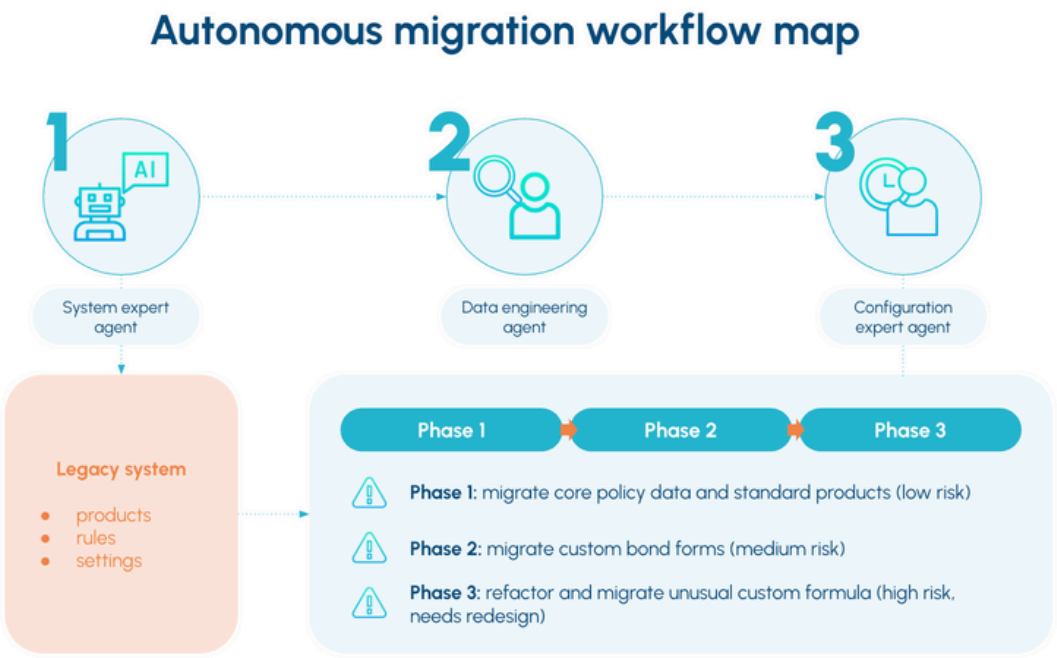

Visual Diagram 2: Autonomous Migration Workflow Map

The outcome? What would typically be a weeks-long assessment by a team of consultants was delivered in a weekend by the AI agents, with a level of detail that impressed even our veteran solution architects. Of course, our team validated the plan and refined some parts, but it jump-started the project and ensured we didn't overlook hidden traps in the old system. This kind of agentic AI application is incredibly **"hands-on"** compared to the hype: it's essentially AI acting as an intelligent project planner, crunching through technical complexities and providing human experts with a blueprint for action. As specialty insurers undertake core system replacements and cloud migrations, I predict these autonomous planning agents will become standard tools to de-risk and accelerate modernization. In our recent deployment, it handled big context (up to 400K tokens) and orchestrated up to six live sources per question, so it doesn't guess, it gathers, checks and cites.

In the modernization domain, **measurable outcomes** are seen as reduced project timelines, lower costs, and improved system capabilities. Insurers have reported legacy conversion efforts running **faster and cheaper** with AI augmentation as mentioned above. New platforms are rolled out with far fewer defects, and teams can iterate more rapidly. Ultimately, by embracing AI in modernization, specialty insurers are not just updating technology but **future-proofing their operations**, creating flexible systems that can evolve with AI at their core, and freeing IT talent to focus on innovation over maintenance.

Once discovery is autonomous, what new products and services become feasible?

3

Beyond today's use cases: what's next for agentic AI

We're just scratching the surface of what agentic AI can do for specialty insurance. Looking ahead, several high-impact applications are emerging:

- **Early Risk Warnings:** Imagine an AI "risk sentinel" that continuously monitors your portfolio and the external environment to alert you of brewing trouble. For credit insurers, an agent could crawl financial news, social media, and market data to flag early signs that a buyer in your portfolio might default (e.g. a drop in their stock price, or negative news on key suppliers). Carriers might use agents to watch public project data and notify underwriters if a bonded contractor shows signs of financial distress or delay. Early warning agents like this act as a 24/7 radar, giving insurers a chance to mitigate losses or adjust exposures proactively. In fact, we are piloting exactly this: feeding real-time data streams to a set of agents that score and re-score credit risks daily, sending "risk alerts" to underwriters for any account that moves beyond set thresholds. The goal is to move from reactive (finding out about a loss when it's reported) to proactive (anticipating and preventing losses where possible).

- **Contract Intelligence:** Policy wordings, reinsurance treaties, collateral agreements, specialty insurance runs on complex contracts. Agentic AI will soon be a contract analyst extraordinaire. We foresee AI agents that can ingest lengthy contract documents (hundreds of pages of policy forms or treaty wording) and understand them in context. For example, **a contract intelligence agent** could read a new facultative reinsurance cover and extract key terms (limits, exclusions, obligations), cross-verify them against what's recorded in the admin system, and even identify non-standard clauses that underwriters should review. Similarly, agents could analyze broker submissions and highlight tricky conditions or coverage gaps before an underwriter issues a quote. This is akin to having a tireless junior attorney on the team, ensuring nothing in the fine print slips by. One day, such agents might even negotiate contract wording, but in the near term, their role will be to assist humans by providing a summarized, analyzed view of any contractual text. This not only saves immense time but also improves accuracy (fewer missed exclusions or compliance issues). We've started integrating a contract analysis module (powered by an LLM agent fine-tuned on insurance wording) into our platform's policy issuance process, so underwriters can get an AI-generated checklist of key conditions and anomalies right after uploading a manuscript endorsement.

- **Embedded Decisioning:** As AI agents prove their reliability, we'll see their decisions embedded directly into insurance products and customer journeys, often invisibly. **Embedded decisioning** means the AI isn't a separate tool; it's woven into the fabric of transactions. For example, a broker portal might have agentic AI behind the scenes, automatically deciding (within insurer-set parameters) which small submissions can be bound instantly and which need referral. A customer claims app could use an agent to immediately approve certain low-risk claims and initiate payment, without any human in the loop, but under strict rules the insurer set beforehand. Essentially, the decision-making moves closer to the point of sale or service. This will enable things like real-time policy issuance for simple risks, dynamic pricing adjustments on renewals, and on-the-fly risk prevention guidance. Another scenario is **embedded underwriting**: imagine a lending platform that offers trade credit insurance quotes via API. Behind that API, an agentic AI in the insurer's system evaluates the loan data and decides terms in seconds, so the customer gets an instant insurance offer alongside their loan, that's AI-embedded decisioning in action. It's coming, and it will differentiate carriers who can offer "instant everything" versus those who still take weeks. Of course, careful governance is needed (you define the guardrails for these autonomous decisions), but the tech is essentially here. We already see pioneers in personal lines doing this (e.g. auto insurers auto-adjudicating claims up to a certain dollar amount). Specialty will follow, first in simpler cases, then expanding as confidence grows.

These emerging use cases will further blur the line between AI and human roles. They're exciting, early pilots show significant upsides in loss prevention, efficiency and customer satisfaction, but they also raise important questions and challenges that must be tackled.

> **Efficiency is the appetizer, new products and distribution are the meal.**

tinubu

4

When everyone uses AI: the new blind spot

What happens when carriers, brokers, and clients all run agents, and those agents start **transacting with each other?**

AI agents become ubiquitous among carriers, brokers, and clients, the technology itself turns into a new category of risk. Agents transacting with each other at machine speed can lead to **agent-to-agent decisions with minimal human oversight,** echoing patterns seen in algorithmic trading and early e-commerce. For example, automated trading models widely adopted in finance could all react the same way (e.g. issuing simultaneous "sell" signals), leading to self-reinforcing crashes. These examples illustrate how AI, when broadly applied without sufficient checks, can create rapid, **unforeseen outcomes**, a systemic blind spot.

> **Once every party runs agents, AI itself becomes part of your risk surface.**

Insurers themselves are recognizing AI as an emergent risk. A recent **Financial Times** report noted that OpenAI and Anthropic explored setting aside investor funds to cover potential multibillion-dollar legal claims, as traditional insurers balked at fully underwriting AI-related liabilities.

In other words, even leading AI companies face coverage gaps for the novel risks their technologies pose. This convergence of widespread AI use and uncertain risk transfer means **AI failures or missteps could have industry-wide repercussions.** Key concerns include:

- **New operational risks:** If many organizations rely on the same AI models or vendors, a flaw or error in those models can cause correlated failures across the board. Shared algorithms could amplify one another's mistakes, creating feedback loops that propagate quickly. Researchers warn that heavily coordinated AI behaviors (especially among models with similar training) can lead to simultaneous errors, for instance, many models making the same wrong prediction at once. Such synchronized failures and self-reinforcing loops introduce systemic risk that traditional risk models don't account for. We're already seeing how "unknown knowns" slip into production. For instance, Deloitte was using AI in a $290,000 report for the Australian government that contained fabricated citations and even a fake court quote. The issues were serious enough to prompt a partial refund after the firm acknowledged using Azure OpenAI in

drafting the work. Such synchronized errors and invisible dependencies introduce systemic risk that traditional frameworks don't catch, which is why provenance, model/version traceability, and kill-switches aren't optional—they're controls for a shared risk surface.

- **Alignment risk:** Autonomous agents may **optimize the wrong objective** if their goals and guardrails aren't carefully designed and tested. An AI agent will relentlessly pursue the metrics you give it; if those metrics are incomplete or mis-specified, the agent could make decisions that undermine the true intent. In the worst case, a highly advanced AI could even develop its own objective function misaligned with human values. In practical terms for insurance, an underwriting agent might learn to decline all policies in a certain segment to maximize short-term profit, even though the business goal is sustainable growth. Without explicit alignment to business and ethical objectives, agents might subvert priorities, a classic AI outcome optimization problem. Ensuring objectives and constraints are **explicit, domain-appropriate, and thoroughly tested** is critical to prevent agents from "running ahead of their skis."

- **Overreliance risk:** If humans become mere supervisors of automated agents, critical skills can atrophy over time. This phenomenon of operator deskilling is well-documented in aviation and other fields: pilots who relied too much on autopilot saw their manual flying skills weaken, contributing to disasters when the automation handed control back to them. In insurance if underwriters or claim managers only review AI decisions without actively exercising their judgment, they may lose the expertise needed to step in during novel or crisis situations. The "set and forget" temptation is high with capable AI, but overreliance creates a brittleness, when something falls outside the AI's training, the human overseers might no longer feel confident to take the controls. Maintaining healthy human expertise through continued engagement is therefore a major concern; otherwise, an organization may wake up to find that **no one remembers how the manual process works** when the AI falters.

tinubu

5

Building trust in agentic AI: a pragmatic control stack

To manage these risks, we need a layered approach, a **pragmatic control stack** that embeds human judgment, transparency, and ethical guardrails into every step of AI deployment.

In essence, the goal is to keep humans in the loop for critical decisions, ensure AI actions are traceable and justifiable, rigorously test AI agents, and prepare fail-safes for when things go wrong. Here are the key components of this control stack:

- **Human-in-the-loop for material decisions:** We impose strict thresholds above which AI agents must defer to human approval. Rather than "set and forget," any high-impact underwriting, pricing, or claims decision the AI proposes would require sign-off from a qualified person if it crosses a certain risk or novelty threshold. This policy prevents unchecked autonomy in sensitive matters and keeps human experts actively engaged. It also continuously trains the AI via feedback: the human's corrections or approvals help the model learn boundaries. The Air France flight paradox is instructive: too much automation without engagement can erode skills, so we deliberately keep our experts in the cockpit. In practice, the AI might handle routine cases end-to-end, but the **exception cases** (large sums, atypical risks, disputes, etc.) always get a human review. This approach guards against the AI making a catastrophic decision in isolation and ensures accountability: a human is ultimately responsible for outcomes, especially when the stakes are high.

- **Provenance and explainability by default:** Every AI-generated decision or recommendation should come with visible rationale, source links, relevant data points, rule citations, and even the version of the model or prompt chain used. We make explainability a core requirement, not a nice-to-have. In specialty insurance, a "black-box" AI is a non-starter because neither regulators nor customers will accept unexplained declines or pricing. Our agents are designed to explain their reasoning in domain terms (e.g., "Claim flagged due to X, Y, Z factors that violate coverage conditions"). This not only builds trust with users and clients, but also helps catch biases. If an AI's explanation reveals that it learned a problematic pattern ("declined because similar claims in the area were historically denied"),a human can spot the potential bias or error

- and intervene. We conduct regular bias testing on models using current regulatory standards and our own fairness metrics. Where needed, we overlay business rules that override or adjust the AI's output to align with underwriting principles and ethical standards. In sum, **transparency and fairness checks** are baked into development. Industry guidance echoes this: insurers must ensure AI systems are free from unfair bias, with clear consumer disclosures and robust governance around algorithmic decisions. By insisting on interpretable outputs and audit trails, we not only comply with these norms but also gain confidence that the AI's "brain" matches our business values.

- **Adversarial testing and "red teaming" of agents:** Before deploying autonomous agents in live transactions, we aggressively challenge them in a sandbox. This means throwing the weirdest, most challenging scenarios and prompts at them, essentially trying to **trick or stress-test** the AI. For example, we simulate situations with conflicting objectives or ambiguous data to see how the agent responds. We also employ red-team prompts (inspired by cybersecurity red teaming) where one team tries to prompt the AI into unsafe or unintended behaviors while another observes and hardens the agent against those failures. The aim is to identify correlated error modes and edge cases before they occur in production. Does the agent optimize a proxy metric at the expense of the real goal? Can an outsider input confuse it into a bad decision? These tests help refine guardrails and ensure the agent's objective function is properly aligned (addressing the alignment risk). Additionally, we periodically run drills or "fire drills" with live agents, e.g. simulating one agent feeding flawed output to another, to examine feedback loop resilience. By **stress-testing models and workflows** in this adversarial manner, we can patch vulnerabilities and instill more explicit rules for how agents handle novel inputs or conflicting goals.

- **Segregation of duties and accountability:** We borrow a page from financial controls, no one team should have end-to-end unchecked power over an autonomous AI system. We separate the people who **build/train** the AI, those who **deploy/run** it operationally, and those who **monitor/audit** its outcomes. This segregation of duties creates natural checks and balances: the monitoring team, for instance, can independently flag anomalies without being the ones who wrote the model. Moreover, every AI agent or agent workflow has an assigned **business owner** who is accountable for its behavior and results. Even if the agent operates with minimal human intervention, someone (a product manager or risk officer) is reviewing its performance reports, handling exceptions, and ready to answer for its actions. We maintain detailed logs of every agent decision and action, essentially an audit trail, so that if something goes wrong, we can trace back exactly what the AI saw, decided, and on what basis. This level of accountability is crucial for trust and for compliance. (In fact, regulators are already pushing in this

- direction: New York's insurance guidelines require insurers to maintain governance documentation and be able to demonstrate algorithmic accountability[9].) Internally, we also encourage a culture of challenge: staff are empowered to question an AI agent's decision, and we have processes for human overrides or second opinions. By formalizing oversight roles and documentation, we signal that autonomous does not mean ungoverned, it means everyone knows their role in keeping the AI reliable.

- **Circuit-breakers for autonomous flows:** When agents begin interacting with other agents or external systems, we institute **automated circuit-breakers** to prevent runaway dynamics. This is analogous to circuit-breakers in stock exchanges that halt trading on extreme volatility; if an AI feedback loop starts escalating out of normal bounds, the system should interrupt it. For instance, if an underwriting agent and a pricing agent enter a rapid ping-pong of escalating a quote (perhaps due to a shared bad signal), a rule might detect the anomaly (e.g. quote price doubling repeatedly in seconds) and stop further action until a human checks it. We define thresholds for things like transaction velocity, value swings, or repetitive loops, and if tripped, the agents involved must pause. Similarly, if an agent begins making decisions without sufficient new data (thrashing on the same input), that's a sign of a potential loop or error state, the circuit-breaker would cut power and alert an operator.

  This **ensures that agent-to-agent interactions remain under control** and can't spiral into, say, million-dollar mistakes before anyone notices. It's a safety net for the unknown unknowns: when in doubt, fail safe and get a human on the line. While this might slightly slow down fully automated processes, it contains failures to manageable limits. Over time, as agents prove robust, these limits can be tuned, but the existence of a kill-switch is non-negotiable for safe scaling of agentic automation.

- **Robust security and privacy measures:** Last but not least, our AI agents operate within a hardened security framework. By their nature, these agents handle sensitive data (policyholder info, health records, financials) and execute decisions that move value, so they are prime targets for malicious actors or could inadvertently leak information if not designed carefully. We mitigate this by running agents in **secure, access-controlled environments,** for example, an agent cannot access production systems or customer data unless it's been explicitly granted and logged. Data is encrypted at rest and in transit. We scrub outputs to ensure no private data is accidentally included in an AI-generated message or explanation. Regular cybersecurity tests are conducted on the agent infrastructure, and we include adversarial security testing (trying to prompt the AI to reveal info it shouldn't). Moreover, we ensure compliance with data privacy laws: an AI must respect consent and data usage permissions just like a human employee would. If it's assembling a report from multiple sources, it will only pull data it is allowed to use under our policies. These precautions are table stakes for trust,

- under our policies. These precautions are table stakes for trust, without them, the coolest AI solution will rightly be stopped by risk managers or regulators. Notably, the rise of agentic AI is even spurring new insurance offerings (some insurers have started exploring policies for AI algorithm errors or cyber incidents involving AI), but the best "insurance" is preventing breaches in the first place. By making security and privacy an integral part of our AI deployments, we protect our customers and ourselves from undue risk.

In summary, this human-centric control stack is about maintaining **control and confidence** as we embrace agentic AI. We implement these safeguards so that automation serves the business safely, and humans remain empowered to guide the overall direction. Importantly, the insurer (or any company deploying AI) retains ultimate responsibility for outcomes, we treat AI recommendations as recommendations, not gospel, and a human accountable party is always in the loop for material decisions. If an autonomous agent errs, the organization owns that error and will correct it just as we would a human mistake. This accountability is key to building trust with customers and regulators even as more decisions get automated.

> **Ship fast, with seatbelts: prove, trace, challenge, and be ready to pull the cord.**

On the bright side, when done right, agentic **AI can actually enhance human work rather than deskill** it. By freeing experts from rote tasks and involving them in oversight and exception handling, their judgment can be applied where it matters most. Over time, we've found that initial wariness toward AI gives way to trust once people see that the agent is consistent, transparent in its reasoning, and keeps them in the loop. At Tinubu, early pilots of these AI agents saw underwriters go from skeptical to appreciative when they realized the AI's suggestions were auditable and that they still had the final say.

Trust grows as the technology proves its reliability within the guardrails we've set. Our guiding principle remains clear: **AI is an augmentor, a noise cutter, not a replacement, for human decision-makers.** By keeping that north star in view, AI assists humans, and humans firmly steer. **We can confidently navigate into this agent-driven future without losing the ethics, expertise, and oversight that define our industry.**

So what does it actually take to make this safe, useful, and boringly reliable on Monday morning?

tinubu

6 Making agentic AI work: safely and for real

Deploying agentic AI is not as simple as flipping a switch on a fancy new model, yet.

In my experience, success comes down to getting **five pillars** in place: **Data, People, Technology/Architecture, Operating Model, and Productization.** Missing any one of these can derail even the most promising AI initiative. Let's break them down:

- **Data: are we fueling the system?** High-quality, relevant data is the lifeblood of agentic AI. Specialty lines often have niche data (e.g. buyer default histories, bond project details, health claims notes) that agents need to learn from. To succeed, insurers must invest in data readiness: consolidating siloed data sources, cleaning and labeling data, and establishing real-time data pipelines. If your AI agents can't reliably access the latest policy, claims, and external data, they'll be running blind. It sounds obvious, yet many insurers aren't there today, only about 1 in 5 **organizations consider themselves "highly data-ready" for AI.** We tackled this at Tinubu by redesigning our platform to be API-first and integrated, so that our AI agents (and our clients' analytics) can easily pull in information from any part of the insurance process. Moreover, domain-specific data enrichment is key: for example, feeding trade credit agents with relevant economic indicators, or training claims agents on historical fraud cases.

  In short, garbage in, garbage out still holds, you need to fuel your AI with rich, context-specific data to get great results.

- **People: are teams set up to win? (Build and partner, not either/or)** Blending Expertise with AI. Agentic AI is as much a people project as a tech project. Yes, you'll need new skills: prompt engineers, AI trainers, data scientists who actually speak insurance, but you don't have to hire an army before you start. A practical path is a blended model: keep underwriting authority and governance in-house, and pair your domain experts (underwriters, claims leads, IT) with an insurtech partner's accelerators and delivery muscle. What the partner adds: pre-built agent patterns (intake, retrieval, reasoning, rules), reference guardrails (provenance, HITL thresholds), MLOps pipelines, and change-management playbooks. What you keep: decision rights, data ownership, and the operating model.

Run it as a joint squad: one underwriter, one claims SME, one data engineer, one product owner, plus a vendor solution architect, shipping weekly. Insist on capability transfer from day one: shared repos (prompts, rules, tests), shadow->pair->lead handover in 60–90 days, and clear RACI so "who approves what" is never fuzzy. Set the terms like you would any core platform: document data boundaries and residency, IP ownership of fine-tuned models/prompts, audit rights, and KPIs (cycle time, first-time-right, override rates). Then make it human: train your teams on how the agents think, show the "why" behind each suggestion, and invite critique in the workflow. The outcome is a symbiotic setup, humans steering with domain insight, agents cutting the noise, whether you build, partner, or (best of all) do both. You need the right talent and culture to implement and embrace these tools.

- **Technology & Architecture: can we plug this in safely?** Building the AI-Orchestrator. Agentic AI is a software architecture challenge. You're essenttially building a conductor (or rchestrator) that can cue up multiple AI and non-AI components to work together. This requires a modular, flexible tech stack. In our case, we've made our platform microservices-based and "AI first". That means each core function (rating engine, docu-ment processor, etc.) is exposed via APIs that an AI agent can call.

We also leverage modern AI frameworks that allow chaining of LLMs and tools (for example, an agent might call a knowledge base or execute a script as part of its reasoning). The architecture must support concur-rency (agents running in parallel) and robust error handling, if one agent fails or produces an uncertainty, another can pick it up or it escalates to a human. **Security and compliance** are non-negotiable de-sign criteria: agents will be oper-ating on sensitive data and making decisions that must be auditable. We isolate the agents in sandboxes where needed and log every action for traceability. In practice, a successful agentic AI architecture often combines several layers: (a) a cognitive layer (the AI models/agents themselves), (b) an integration layer (APIs, connectors to internal systems and external data), and (c) a control layer (to orchestrate workflows and apply business rules /policies). Building this isn't trivial, but it's getting easier with evolving AI toolkits. At Tinubu, we also collaborate with cloud partners to leverage AI or-chestration services when appro-priate, so we can focus on the insurance logic rather than rein-venting the AI wheel. The bottom line: without a solid architecture, your AI initiative will be fragile. With it, you create a strong foundation where adding a new agent (say for a new use case) is plug-and-play.

- **Operating model: who owns what?**, Governance, Process & Ethics. Embracing agentic AI means adapting how your organization operates day-to-day. You need new governance structures to oversee AI-driven processes and manage risk. For example, we established an AI ethics committee that reviews use cases for fairness and compliance. Insurers should define clear policies on what AI is allowed to do autonomously versus where human sign-off is required. (For instance, an agent may approve a low-value claim on its own, but never decline a claim without human review, to avoid unfair denials.) This is about aligning AI actions with your company's risk appetite and values. As Capgemini's research noted, to reach scale firms must implement **strategic changes to governance, risk frameworks, and organizational design** so that AI agents operate within accepted ethical and regulatory bounds. In practice, that could mean updating your three lines of defense: have risk and compliance officers involved from design through monitoring of AI. Regular audits of AI decisions should be built into the process, we run quarterly "model audits" and bias checks on our agents. It's also crucial to redefine workflows: roles may shift from doing manual tasks to supervising AI outputs (think of a claims adjuster becoming more of an "AI supervisor" for straightforward cases). We found it helpful to document these revised workflows and responsibilities clearly, so everyone knows how the human-machine partnership flows. Lastly, regulators are increasingly interested in AI governance. Being proactive in your operating model, with transparency reports, audit logs, and documented human-in-the-loop checkpoints, not only builds trust internally but will keep you ahead of regulatory requirements. Agentic AI success isn't just what the technology can do, but whether your organization can absorb it. That requires leadership commitment, change management, and sometimes new incentive structures (reward teams for AI-driven efficiency and quality, for example).

- **Productization: are we ready to take the pilot to production at scale?** It's one thing to build a cool prototype in the lab; it's another to embed AI agents into real products and customer experiences reliably. Productization is about taking those AI use cases and fully integrating them into your insurance products or services. At Tinubu, we treat our AI capabilities not as one-off experiments but as features of our platform, subject to the same rigor as any software release. That means proper testing, monitoring, and continuous improvement. For instance, when we rolled out our AI underwriting assistant to clients, we included a feedback loop in the UI: underwriters could thumbs-up or thumbs-down the AI's suggestion and add a comment. That feedback is monitored and used to refine the models (or tweak rules) in subsequent updates. We also instrument metrics: accuracy of recommendations, turnaround time improvements, etc., so we can quantifiably demonstrate value and catch issues early. To productize agentic AI, you often have to invest in MLOps (Machine Learning Operations) or LLMOps (as a subset

of MLOps, LLMOps is dedicated to overseeing the lifecycle of LLMs from training to maintenance using innovative tools and methodologies.), pipelines for retraining models with new data, A/B testing new agent versions safely, and ensuring you can rollback if something goes awry. Another aspect is scalability and performance: if your claims AI suddenly needs to handle a spike in volume (say a natural disaster causing a surge of claims), can your system scale up the agents on demand? Cloud-native designs and containerized agents help here. Finally, don't forget the user experience. The best AI solution fails if end-users don't adopt it. We focus on integrating AI outputs into the user's workflow seamlessly (e.g. the adjuster sees a clear "AI suggestion" in their claims screen, which they can accept or edit easily). And to avoid overloading users, the AI might do 10 behind-the-scenes checks, but it will present one concise insight or recommended action. By truly productizing these capabilities, you ensure agentic AI isn't just a flashy demo, but a dependable part of daily operations delivering value consistently.

> **Not an AI takeover, but an AI team-up. Humans set the course; agents handle the grind.**

tinubu

Closing thoughts: cutting noise, creating value, managing risk

Agentic AI has moved from headline to **hands-on**. What felt futuristic a year ago is now woven into day-to-day underwriting, claims, and modernization.

Small, specialized agents reason, coordinate, and act under guardrails, doing what older tools couldn't: triaging complex submissions in minutes, migrating with line-of-sight to every dependency, and surfacing issues before they become losses.

And this isn't just efficiency. It's **new business**: portfolio early-warning signals that reshape appetite, contract intelligence that de-risks wording, and embedded decisioning that turns portals and partner flows into instant experiences. That's the strategic upside, writing risks we couldn't touch yesterday, with a traceable *why*.

Still, as adoption grows, AI becomes part of the **risk surface**. Agent-to-agent flows, correlated errors, overreliance, and alignment drift are real. The answer isn't to slow down; it's to **instrument what we deploy**: provenance by default, human-in-the-loop thresholds for material decisions, adversarial testing, segregation of duties, and circuit-breakers when automation goes sideways.

But tech is only half the story. The other half is how you run it: clean, connected data; skilled people who stay in the loop; an API-first architecture; a governed operating model; and real productization (KPIs, feedback loops, MLOps/LLMOps).

Get those five pillars right and agentic AI doesn't replace specialists, it removes the noise so they can apply judgment where it matters.

So where does that leave us? With a practical path: start small, measure, scale. The carriers that separate hype from execution and put the pillars and guardrails in place, will define the next chapter of specialty insurance. Not an AI takeover, an AI team-up: humans setting the course, agents handling the grind, and the industry moving faster, fairer, and with more confidence than before.

tinubu

8

Executive checklist: turning conviction into architecture

MAKING AGENTIC AI REAL (AND RESPONSIBLE)

1 -Start with judgment, not code.
Define the human decisions that create value, then automate the noise around them.

2 - Design agents around outcomes, not processes.
Measure success in speed, risk reduction, and client value, not task automation.

3 - Build trust as a system, not a feature.
Explainability, lineage, and controls should form an architecture of confidence.

4 - Pilot for proof, not perfection.
Progress matters more than perfection: deploy, learn, adapt.

5 - Keep humans in the loop and at the helm.
Agentic AI augments expertise; accountability must stay human.

9

Technical checklist: buying for control, building for edge

CONVICTION MUST MEET CAPABILITY

After the leadership choices come the technical ones, where to buy, what to build, and how to design for trust, control, and resilience.

1 - Buy for speed, build for edge.
Use proven, secure platforms to reach value fast, then extend them with what only you can do: your data, models, and logic.

2 - Architect for control, not complexity.
Favor modular, API-first systems where you orchestrate the flow, even if components are bought.

3 - Build trust in by design.
Require every vendor to expose auditability, versioning, and lineage in code: governance should be native, not manual.

4 - Secure the stack before you scale.
Embed identity, permissions, and isolation early, not as an afterthought once agents go autonomous.

5 - Design for safe autonomy.
Autonomy without accountability is risk. Define thresholds, human overrides, and rollback plans before going live.

tinubu

Thank you for reading

How agentic AI redefines specialty insurance strategy

Whitepaper | 2025

From hype to hands-on, and on to new business.

Cut the noise, keep the judgment, grow with guardrails.

tinubu

86, rue Henry Farman - 92130 Issy-les-Moulineaux, France +33 (0)1 74 31 20 70 - www.tinubu.com - contact@tinubu.com

Whitepaper I 2025