Capgemini

Agentic hyper-personalization at scale

The new standard for insurance RFPs

In partnership with MongoDB

MongoDB

Generic proposals are losing deals

Insurance RFP responses are starting to feel like they've been photocopied over and over. Brokers and clients today are no longer just flipping through proposals hoping to find a winner—they're expecting them to speak directly to their unique needs. The days when you could get away with templated, one-size-fits-all responses are behind us. In insurance, trust is built on understanding, and understanding is signaled through specificity.

In fact, many proposals don't even get past the first skim because they sound like they were written for any client, not this client. The root issue is that generic responses signal a lack of investment in the relationship. Insurers risk losing out on high-value deals, wasting time and resources crafting responses that don't convert. As our work with numerous global insurers has revealed, many of these generic documents—especially cover letters and executive summaries—were not even being read by brokers due to their lack of relevance.

Generative AI for hyper-personalization in insurance

Now, let's imagine a private, enterprise-trained generative AI assistant that doesn't just regurgitate past language, but crafts messages so tailored they make your clients feel like VIPs. That's the magic of a custom, private GenAI assistant.

This assistant is no off-the-shelf chatbot. It's trained on your historical RFP data, your previous client interactions, your industry nuances, and even your internal product literature. It understands how you communicate and what your clients care about. More importantly, it learns and evolves. With the help of Agentic AI, a modular framework powered by specialized AI agents, this assistant goes far past simple auto-fill. It reads the RFP, summarizes the client ask, constructs the top winning themes, and proactively drafts personalized responses, summaries, and even intelligent suggestions for improvement.

This is where hyper-personalization becomes real. By utilizing structured and unstructured data alike, the Gen AI assistant pulls out the most relevant insights and shapes them into messaging that resonates. It compiles data from its entire knowledgebase to craft a tailored solution to the client's problem. It's not guessing, it's contextualizing. That means proposals land stronger, faster, and with far better chances of hitting the mark.

Private Gen AI assistant Agentic Workflow

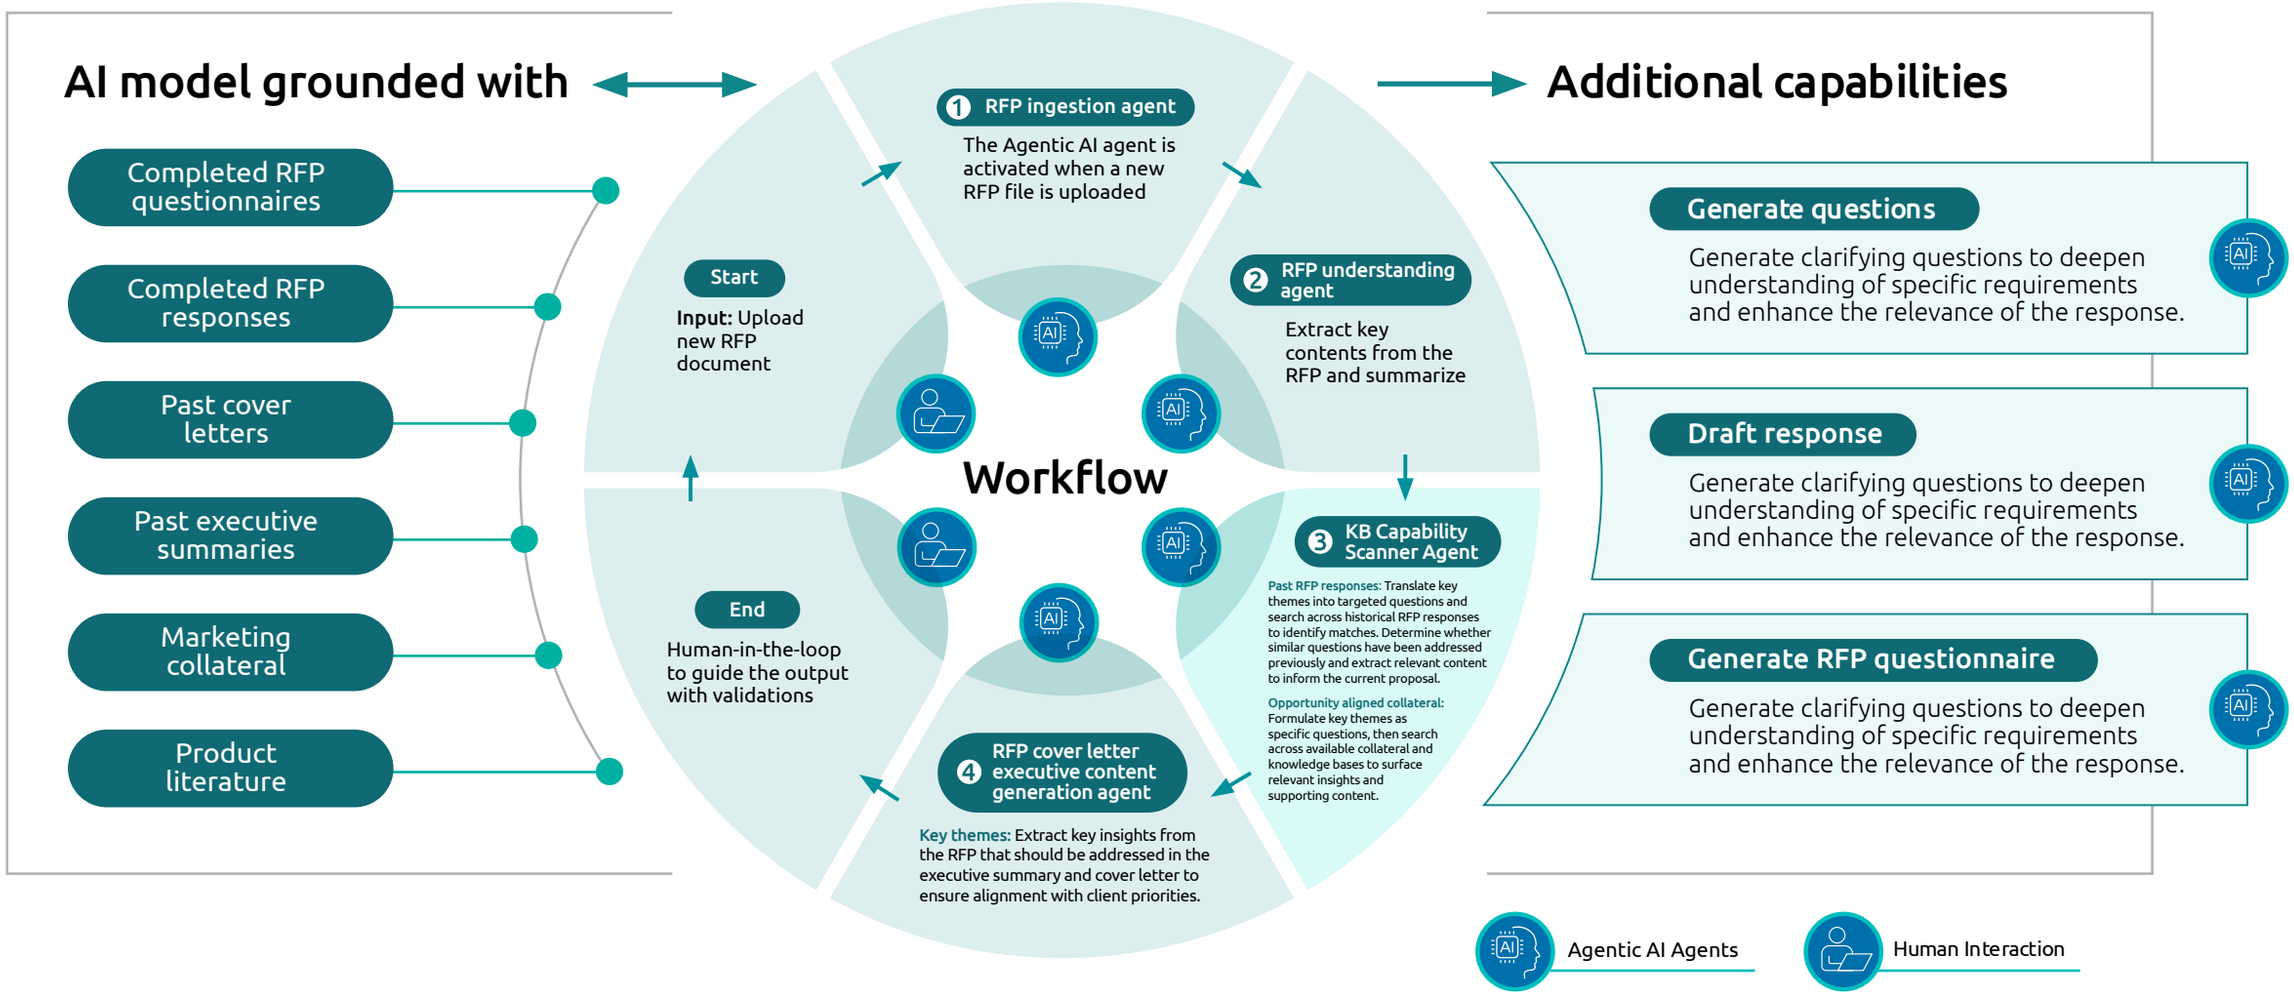

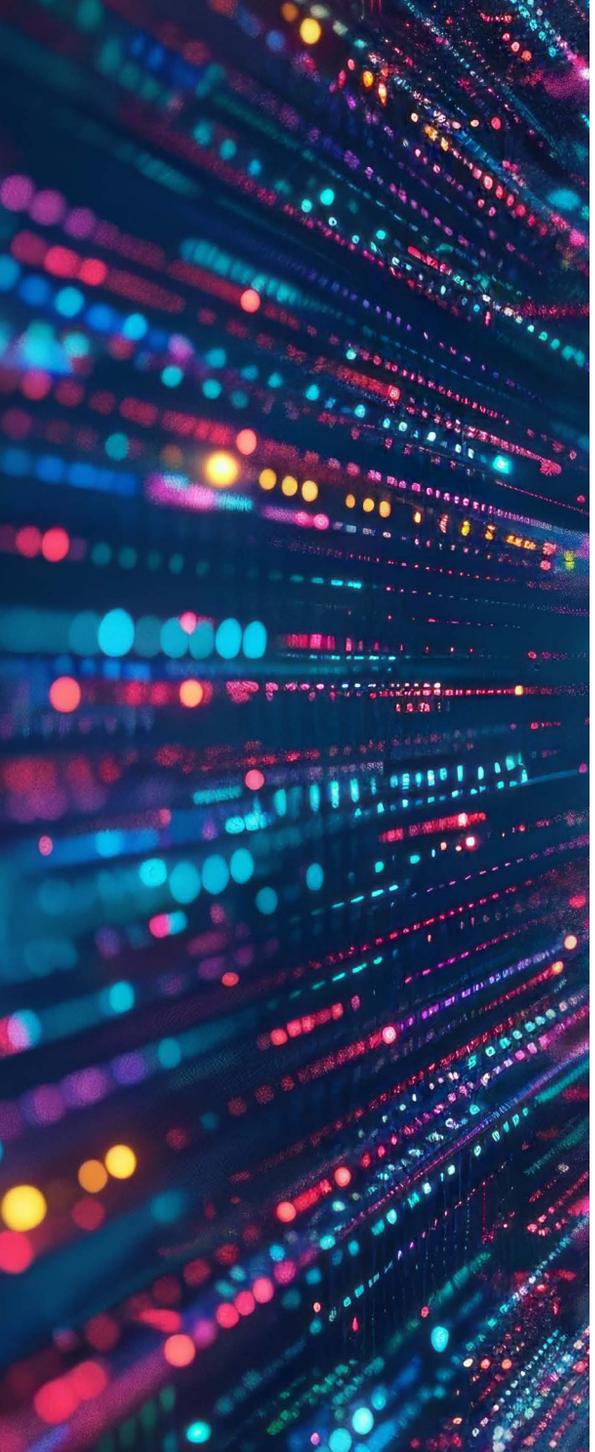

MongoDB: The motor powering AI-driven personalization

Behind the scenes, MongoDB plays a crucial role in making all this magic possible.

Their flexible document model allows for rapid ingestion of diverse data types including past RFPs, client correspondence, marketing decks, and everything else imaginable. This structure is perfect for insurers juggling massive volumes of semi-structured and unstructured data.

MongoDB Atlas Vector Search is particularly crucial here. It enables the Gen AI assistant to rapidly identify, rank, and re-rank the most relevant information based on contextual relevance, delivering responses that are both timely and precise.

Its globally distributed architecture—available across AWS, Azure, and GCP in over 115+ regions—makes it an ideal foundation for building large-scale, enterprise-grade Gen AI applications. By embedding Vector Search directly into the core database, MongoDB eliminates the need to sync data between separate operational and vector databases. This simplification reduces complexity, minimizes the risk of errors, and significantly shortens response times.

Keeping both operational and vector data in a single system also improves performance through reduced latency and advanced indexing capabilities. For organizations building out agentic Gen AI capabilities, MongoDB further supports Graph RAG (Retrieval Augmented Generation) architectures, enhancing contextual accuracy and scalability across use cases.

However, insurance is a heavily regulated industry and data security is critical. MongoDB also offers enterprise-grade encryption, access controls, and supports compliance with key data privacy regulations.

Case study

Less robotic, more calibrated and compelling RFPs at a global insurer

A recent standout example of our custom, private GenAI assistant in action comes from a global insurer who started with a modest request: Can we hyper-personalize our RFP cover letters better? The ask was simple and they were merely looking for a few bullet points to make things feel less robotic.

What we were able to create for them was a revolution in how they respond to RFPs. In just five weeks, our team implemented our custom, private GenAI assistant that not only delivered personalized bullet points but also crafted full executive summaries and tailored cover letters. These were not piecemeal templates—they were coherent, compelling, and calibrated to the specific opportunity at hand.

The feedback we received was immediate and enthusiastic. The Chief Innovation Officer and the Sales leadership team pushed for scaling the solution to other areas. It wasn't just a productivity gain, it was a reputation builder. Brokers began to take notice. The insurer wasn't just responding faster; they were responding smarter.

Business impact, check! Strategic outcomes, check!

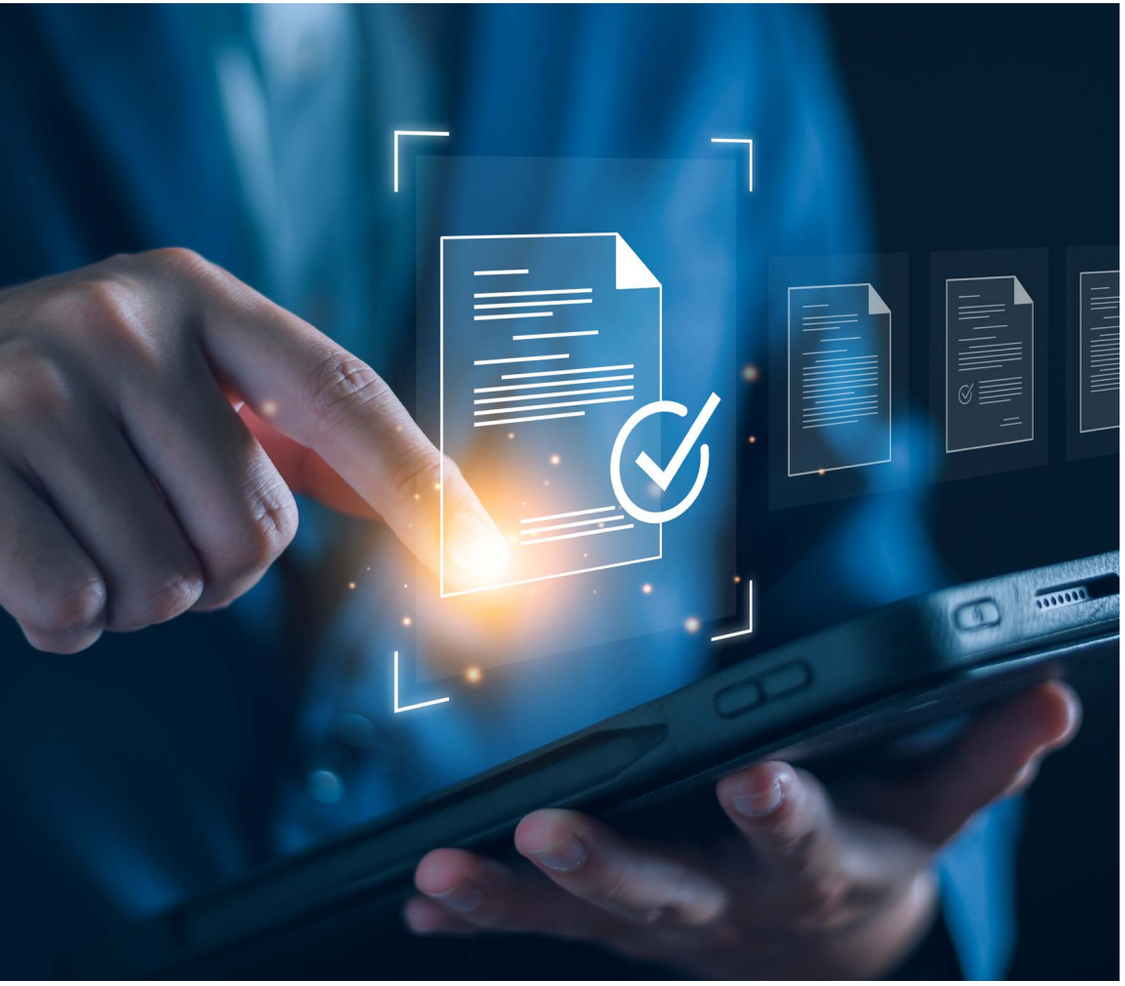

By implementing a custom, private GenAI assistant, insurers gain access to a scalable, cloud-native platform that integrates easily with existing systems—whether it's a CRM, document management platform, or internal knowledge base. Beyond the technical flexibility, the real impact lies in how this approach transforms stagnant, siloed data into living insights that power tailored client engagement.

The platform supports more consistent and efficient proposal development by reducing manual effort, accelerating turnaround times, and improving the quality and relevance of responses. Teams can focus less on reformatting and more on building client relationships. Meanwhile, the built-in security and governance measures ensure that every interaction meets enterprise compliance standards, protecting both client data and institutional knowledge.

Insurers using this model report stronger broker engagement, better win rates, and faster RFP response times. Operational costs drop due to reduced manual formatting and response drafting. From a technical perspective, RAG-enhanced GenAI can offload up to 35% of compute cost compared to full LLM inference on raw content, thanks to targeted document retrieval and short-form reasoning tasks.

As organizations use this solution over time, feedback loops from won/lost deals can be fed back into the model for retraining, improving response quality and alignment. As the assistant matures, it can serve as a strategic enabler across adjacent workflows—claims review, renewal briefs, or even sales coaching.

The future of insurance RFPs

Custom private GenAI assistants represent a rare intersection of technical maturity and business impact. When combined with MongoDB's robust data orchestration capabilities and Capgemini's proven technology blueprint, this solution becomes more than a digital enhancement—it becomes a strategic advantage.

Organizations that embrace this model transition from reactive, templated proposal development to proactive, context-rich client engagement. With the ability to generate intelligent, personalized content at scale, they not only improve operational efficiency but also strengthen their competitive position in a high-stakes market.

This isn't just about responding faster—it's about responding better. As expectations around relevance, precision, and value continue to rise, the future of insurance RFPs will belong to those who invest in intelligent automation and meaningful personalization.

The path forward isn't generic. It's personal, scalable, and ready to deliver lasting impact.

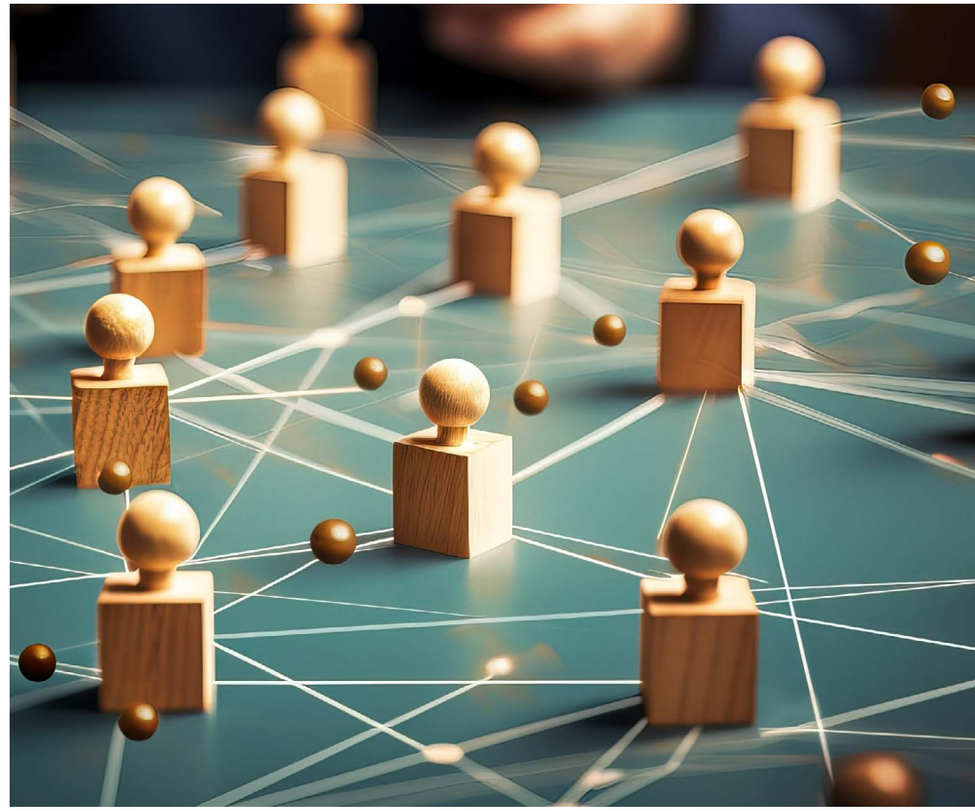

Meet our experts

Pinaki Bhagat
AI & Generative AI Solution Leader
Financial Services Insights & Data
Capgemini

Shounak Acharya
Senior Partner Solutions Architect and PFA
MongoDB

About Capgemini

Capgemini is a global business and technology transformation partner, helping organizations to accelerate their dual transition to a digital and sustainable world, while creating tangible impact for enterprises and society. It is a responsible and diverse group of 340,000 team members in more than 50 countries. With its strong over 55-year heritage, Capgemini is trusted by its clients to unlock the value of technology to address the entire breadth of their business needs. It delivers end-to-end services and solutions leveraging strengths from strategy and design to engineering, all fueled by its market leading capabilities in AI, cloud and data, combined with its deep industry expertise and partner ecosystem. The Group reported 2024 global revenues of €22.5 billion.

Get the Future You Want | www.capgemini.com

About MongoDB

MongoDB's developer data platform offers significant architectural advantages by enabling organizations to securely unify application data (both structured and unstructured) with AI-related data (vectors). This capability allows institutions to build rich, real-time AI applications. At the core of MongoDB's developer data platform is MongoDB Atlas, the most advanced multi-cloud database on the market. Atlas provides unmatched data distribution and cloud mobility, built-in automation for resource and workload optimization, and a flexible document model, among other features. MongoDB also offers the flexibility to deploy applications on-premises, on a single public cloud, or across multiple clouds simultaneously, ensuring resilience, scalability, and the highest levels of data privacy and security.

For more information, please visit www.mongodb.com

In partnership with MongoDB