

Agentic AI in Action

An Executive Perspective on Building a Zero-Friction SDLC

In an era where applications are no longer just built, but co-created with autonomous AI agents, software development is entering a transformative new phase. In this exclusive interview, [Deepak Mittal](#), Founder & CEO of [NextGen Invent](#), and [Michael Kaminaka](#), Chief Growth Officer at NextGen Invent, unpack how Agentic AI is redefining how products are conceived, engineered, tested, deployed, and evolved, paving the way for a zero-friction SDLC. They explore how intelligent AI agents streamline every stage of development, from early concept exploration and architecture planning to coding, testing, integration, and ongoing optimization, reducing manual effort, accelerating delivery, and improving reliability.

Deepak emphasized that the success of AI adoption hinges on an organization's AI maturity, leadership's willingness to take calculated risks, and ability to navigate industry-specific nuances. He also highlights how AI is elevating the fundamentals of Software Development Life Cycle (SDLC), Product Management, enhancing traditional frameworks such as 3C (Customer, Company, Competitor), STP (Segmentation, Targeting, Positioning), and 4P (Product, Pricing, Promotion, Placement) by helping teams save significant time in research and enabling stronger product-market alignment.

"To work effectively with AI, software professionals need a combination of technical, analytical, and human-centric skills. They must understand AI's capabilities and limitations, evaluate outputs critically, and collaborate across disciplines. Awareness of privacy, security, bias, and accountability is essential, along with the ability to apply creativity rather than just speed in development." - Deepak Mittal, Founder & CEO, NextGen Invent

Enabling the Zero-Friction SDLC with Generative AI: Fastest ROI, Ethical Imperatives & Risk Mitigation

Michael Kaminaka: Across the Software Development Life Cycle, where do you see Generative AI delivering the fastest ROI, during requirements gathering, coding, testing, or ongoing maintenance stages?

Deepak Mittal: In our experience, we've seen AI deliver the fastest ROI in a few clear areas.

- *First, documentation, including requirements, test cases, and code comments. That's the low-hanging fruit, and the returns are both immediate and measurable.*
- *Second, unit testing and code reviews, where automation enhances speed, quality, and consistency across teams.*
- *Third, coding itself, but here, the ROI varies significantly. Refactoring isolated utility code is quite different from developing highly integrated business logic. That spectrum makes the returns more difficult to quantify.*

Most people focus only on **short-term ROI**, but I'd recommend looking deeper. If better testing and documentation reduce bugs and rework, your **total cost of ownership** decreases substantially.

That long-term value often outweighs the short-term gains of faster sprints or quicker releases to production.

Michael Kaminaka: What do you see as the biggest risks and inherent limitations of relying on AI in software development, and how can organizations proactively address and mitigate them?

Deepak Mittal: When it comes to limitations, several challenges stand out:

- **Domain-specific knowledge gaps:** Much business logic, such as risk modeling for P&C, is not publicly available, making it difficult for AI to learn effectively.
- **New technologies and platforms:** AI tools take time to adapt. For example, asking AI to code in Palantir or other niche platforms often fails initially.
- **Legacy code integration:** Incorporating AI into existing systems is challenging, particularly when the model is not trained in your organization's coding style.
- **Data limitations:** Building AI models using AI are constrained by limited proprietary data; public datasets are often insufficient.
- **Tool fragmentation:** One AI tool rarely covers everything. Using multiple tools together introduces complexity and inefficiency.
- **Strategic decision-making gaps:** AI cannot yet decide which technology stack, cloud provider, or third-party vendor is most suitable.
- **Opportunity recognition:** AI may miss opportunities that rely on human judgment or deep contextual understanding. For instance, platforms like Mayo's data ecosystem require years of collaborative work; AI cannot replicate that without human guidance.

Regarding risks, organizations must be cautious about:

- **Security and privacy exposure:** Risks include prompt injection, data leakage, data poisoning, and potential jailbreaks.
- **Knowledge gaps:** Many organizations lack the expertise to evaluate the risks of generative AI in software development effectively. Relying solely on AI to assess its own safety is not sufficient.
- **Reputation risk:** Using AI tools without proper governance can lead to missteps. Organizations must be deliberate in adoption to avoid damage to credibility or trust.

Mitigating these challenges requires a combination of **strong oversight, human validation, governance frameworks, and continuous skill-building** to ensure AI complements development rather than introducing unforeseen risks.

Michael Kaminaka: *What essential skills should today's software professionals develop to collaborate effectively with AI systems while maintaining creativity, accountability, and technical excellence?*

Deepak Mittal: To work effectively with AI, software professionals need a combination of technical, analytical, and human-centric skills. They must understand AI's capabilities and limitations, evaluate outputs critically, and collaborate across disciplines. Awareness of privacy, security, bias, and accountability is essential, along with the ability to apply creativity rather than just speed in development.

Key skills include:

- **Prompt engineering:** Developing effective prompts to guide AI behavior accurately.
- **Data literacy:** Understanding and validating AI outputs using strong mathematical and analytical skills.

- **Developer creativity:** Solving novel problems in innovative ways, beyond mere productivity.
- **Cross-disciplinary knowledge:** Combining math, product thinking, customer-centricity, and collaboration skills.
- **Critical evaluation:** Reviewing AI-generated code and maintaining healthy skepticism of outputs.
- **Business and Industry knowledge:** Understanding the broader business context and industry dynamics to align AI solutions with organizational goals.

Professionals equipped with these skills can harness AI as a powerful collaborator, enhancing innovation while ensuring software quality, ethical standards, and long-term strategic impact.

Michael Kaminaka: How are AI coding assistants changing the developer experience today? Are they primarily improving productivity, or fundamentally shifting the definition of what “good engineering” means?

Deepak Mittal: In my view, AI pair programming plays out very differently depending on experience. Junior developers often over-trust the AI's output, which can be detrimental if they lack the foundation or domain expertise to catch mistakes.

Senior developers, on the other hand, know how to validate results and refine prompts, so they see real productivity gains. In short, AI is boosting senior talent but can set back juniors if not managed carefully.

I'd also add that what I call “zero-friction SDLC” has two pathways: one driven by productivity and the other by creativity.

AI pair programming today is tilted toward the productivity side, speeding up coding, testing, and documentation, but it's not yet enabling the creative breakthroughs that emerge from design thinking and deep problem-solving.

Michael Kaminaka: In your view, what aspects of Generative AI in software engineering are currently overhyped, and which areas do you believe hold the most underrated, high-impact opportunities?

Deepak Mittal: When evaluating the hype versus reality, I prefer to look at it through two distinct lenses, **business and technology**.

From the business perspective, the hype sounds incredible: “Use this tool and you’ll get 50, 70, even 90% productivity gains. Just type a few prompts, and the code will be ready to deploy within an hour. Everyone becomes a full-stack developer, with zero defects and bulletproof security.” That’s the story we often hear.

On the other hand, several **underrated opportunities** are quietly emerging. Technologies such as the **Model Context Protocol (MCP)** and **virtual assistants**, when used beyond simple chat functions to address **dashboard fatigue** and improve workflow efficiency, are true game changers. These innovations bring practical, measurable value to teams, enhancing focus, reducing context-switching, and enabling more intelligent, frictionless collaboration.

Michael Kaminaka: What key ethical and security considerations should organizations prioritize when integrating AI into their software development processes to ensure responsible and transparent adoption?

Deepak Mittal: Ethics and security in AI are **dynamic and context-dependent**. What is considered ethical today may not be tomorrow. For example, in healthcare, is it ethical to automatically decline claims in the first review? AI systems will learn from historical patterns, but organizations must carefully consider local and global ethical standards; what’s acceptable in the USA may not be acceptable in the United Kingdom.

To address these concerns, organizations need a strong AI governance framework and AI model observability.

This includes testing for diverse and edge cases, implementing guardrails, and monitoring prompt injections, data leakage, and model drift.

Other critical considerations include:

- **IP and licensing risks:** Ensuring proper rights to training data and avoiding code-license contamination.
- **Human oversight:** Keeping humans in the loop where decisions have a significant impact.
- **AI-on-AI development:** If an AI model is used to create another AI model, observability and testing requirements increase.
- **Documentation and testing rigor:** Focus on detailed testing and comprehensive documentation becomes even more critical.
- **Tools for observability:** Platforms like AI Foundry will become essential for monitoring and managing AI systems effectively.

By integrating governance, observability, and human oversight, organizations can mitigate risks while responsibly leveraging AI to enhance development processes. Knowing where the source of information came from, how the end user leveraged the information, and having a secure logging of events is crucial for maintaining the integrity of the body of knowledge.

Wrapping Up

The interview reinforces that Agentic AI is not just enhancing development workflows; it is reshaping how software is conceived, engineered, tested, deployed, and evolved, laying the foundation for a truly zero-friction SDLC.

From autonomous code generation and intelligent testing to continuous optimization and automated delivery pipelines, Agentic AI is enabling a development model where bottlenecks are reduced, iteration cycles are accelerated, and quality becomes more consistent. Yet, the discussion makes it clear that these gains materialize only when organizations combine advanced AI tools with strong human oversight, governance, and ethical guardrails.

Leaders must evaluate their organization's AI maturity, cultural readiness, and risk appetite, aligning AI adoption with broader business goals and industry-specific requirements. Teams across product and engineering must deepen skills in data literacy, critical evaluation, prompt and agent orchestration, and collaborative decision-making, while continuing to rely on human judgment, creativity, and strategic thinking. Ethical and security considerations, ranging from data privacy and IP protection to bias prevention, model drift management, and system observability, remain central to responsible scaling.

The future of software development will rely on a synergistic partnership between humans and AI, where [agentic AI enabled generative AI development services](#) empower teams to automate repetitive and data-intensive tasks, while enabling developers to focus on creative problem-solving and strategic innovation. Organizations that embrace this approach will not only accelerate delivery and improve quality but also establish a foundation for sustainable, long-term growth in an AI-driven world.



Deepak Mittal

Founder & CEO, NextGen Invent

Deepak Mittal is an accomplished serial entrepreneur with a proven track record in product, AI, and data strategy. He is experienced in transforming an organization into an AI-enabled organization. Currently, Deepak holds the Founder & CEO position of NextGen Invent Corporation, a leading global professional services company that has always been at the forefront of providing solutions using disruptive technologies such as Agentic AI, Computer Vision, and Intelligent Automation. During his career, Deepak has served as a member of the Board of Directors of many companies, including D4DT, Optym, CMR Institute, CBS Alum Club of New York, and Launch Right Now. He is a strategic advisor to various organizations that grew from small start-ups to unicorns or had successful exits. Deepak holds expertise in evaluating new product ideas, devising an appropriate product and AI strategy, transforming organizations into insight-driven organizations, and structuring strategic partnerships. Deepak has an MBA from Columbia Business School and an M.S. in Industrial Management Engineering from the Indian Institute of Technology.



Michael Kaminaka

Chief Growth Officer, NextGen Invent

Michael Kaminaka is a seasoned healthcare executive with over 30 years of experience spanning inpatient and outpatient care, value-based care, health insurance, life sciences, and employer health solutions. Known for his relentless curiosity and hands-on approach, Mike has dedicated his career to helping organizations reduce friction, enhance workflows, and accelerate their data transformation journeys. Guided by his core belief to “meet the customer where they are and help them reach their desired destination,” Mike combines deep industry insight with a collaborative mindset to drive meaningful, technology-enabled change across the healthcare ecosystem.



GET IN TOUCH WITH US



FOLLOW US ON



SUBSCRIBE NOW

