From Proof of Concept to Production : A Guide to De-Risking Generative AI for Real-World Impact

Learn how to tackle technology, security and regulatory risks to implement generative AI prototypes at scale

By **Mohammad Waisam and Nancy Silver** June 26, 2025

Mohammad Wasim
Group Vice President, Global
AWS Alliance Lead

Nancy Silver
Vice President, Business Development
AWS Partnership

Intro: Why do most generative AI prototypes fail?

Building an AI prototype is easy. Getting it into production? That's where things start to get tricky.

Here are the top three reasons why generative AI proofs of concepts (PoC) fail, that we've observed across industries:

- They fail to capitalize on the early mover advantage—and by the time they're perfect, the organization has lost competitive edge
- Organizations don't invest in developing internal talent and AI expertise
- There's no clear framework for measuring success and managing implementation risks

But what about the success stories? The companies that bridge the gap between prototype and production are able to gain a significant edge over competitors.

> "Generative AI experiments are a cost. Generative AI products are cost savings."
>
> **Francesca Sorrentino**
>
> Senior Client Partner, Generative AI Ethics Task Force

The early mover advantage

Many companies hesitate to invest in generative AI because:

- The technology is still evolving
- The regulatory requirements are complex
- ROI is uncertain
- Mistakes can lead to security and reputation risks
- There's no clear vision or unified goal

However, if we observe that generative AI technology is following the same pattern of cloud technology, we can assume that early movers in the space will gain a long-term competitive advantage and larger market share—like Amazon, Microsoft or Google.

This trend is even more true with generative AI technology because of one thing: data.

Think of generative AI as a snowball rolling downhill. The more user data it collects, the better it gets. This creates a feedback loop that's tough for latecomers to replicate.

There’s no shortcut here—real-world adoption is the only way to improve your AI’s performance.

The talent edge

Early adopters aren’t just ahead in data and technology; they’re also ahead in developing AI talent.

Your employees are most likely experimenting with AI on their own, but enterprise-scale solutions require specialized expertise, like data engineers, machine learning experts and product managers familiar with AI. The gap between current skills (prompt engineering on public tools) and future needs (collaborating with agentic AI co-workers) is an opportunity—but only for companies that invest in upskilling their workforce now.

The talent edge is a major opportunity. The key? Hands-on experience. Invest in upskilling your workforce now to lead the race rather than scrambling to catch up later.

> "Ubiquitous use of AI will not equal a level playing field. Your people and your people’s skills will be a huge differentiator in a war for AI talent.”
>
> **Simon James**
>
> Managing Director, Data & AI, Publicis Sapient

What separates success from failure?

Aside from the early mover advantage and the talent edge, there isn't necessarily a perfect formula for generative AI implementation success, or ROI. Any consultancy that tells you so is probably trying to sell you something.

Rather than obsess over the perfect plan and fear failure while standing in place, it's better to take action with a clear understanding of the risks, and how to mitigate them.

Our generative AI ethics and governance task force has identified **five key risk categories**, through researching hundreds of generative AI PoCs internally and externally:

1. **Model and technology risks:** Choosing the right AI architecture for cost, speed and scalability
2. **Customer experience risks:** Ensuring AI-generated content is relevant, clear and useful
3. **Customer safety risks:** Preventing AI from generating harmful or biased outputs
4. **Data security risks:** Protecting proprietary and sensitive information
5. **Legal and regulatory risks:** Staying ahead of evolving AI laws and ethical considerations

This article breaks down each risk area and provides strategies to mitigate them, so that you can move forward to the most important thing: action.

Top model and technology risks

Model and technology risk involves choosing AI tools that balance quality, speed and cost while preparing for future updates and unexpected usage. This often requires special databases and secure environments that traditional IT setups might not have. Here are the key risks to address when moving from proof of concept to production:

- **Beyond accuracy—think cost and speed.** The "best" AI model (which is literally changing every single day) isn't just about accuracy—For enterprises, it's about what's cost effective and easy to implement. Scalability and seamless updates should be your priority.
- **Prevent overuse with rate limits.** Generative AI can be a powerful tool, but like any good thing, moderation is key. Set usage caps to avoid excessive use and ensure smooth operation for everyone.
- **Future-proof your tech stack.** Your traditional systems might not be AI-ready. After you build your PoC, be prepared for roadblocks from existing architecture like slow APIs and on-premises data limits. Consider upgrades, workarounds or phased implementation to ensure seamless integration.
- **Break down generative AI silos:** Ensure technical teams across the organization understand generative AI's strengths and weaknesses to prevent misalignment and investing in LLMs when they're not actually necessary. Collaboration and upskilling are key.

Top customer experience risks

Customer experience risks involve ensuring AI-generated content remains relevant, clear and helpful to users. Poor AI interactions can frustrate customers and damage trust in your brand. Here are key strategies to mitigate these risks:

- **Break down big questions, get better answers.** Don't overwhelm your AI model with complex questions. Oftentimes, splitting queries into smaller, clearer ones can reduce irrelevant responses or fabricated information ("hallucinations")
- **Train the AI to speak your customers' language.** Use prompt engineering to ensure customer-facing AI tools understand key customer queries and tasks. **Trustworthy data = trustworthy answers.** AI is only as good as the data it's trained on—prioritize high-quality, pre-verified datasets.
- **Empower users through UX.** AI adoption improves with intuitive user experiences. Design tools that make AI easy to use.
- **Keep humans in the loop.** Generative AI should enhance customer interactions, not fully replace human oversight.

Top customer safety risks

Customer safety risks involve preventing your AI from generating harmful, biased or misleading content that could hurt users or spread misinformation. Your organization bears ultimate responsibility for AI outputs. Consider these protective measures:

- **Don't rely on LLM safety nets alone.** Prebuilt safeguards (from OpenAI or Meta) are helpful, but they can't fully protect your customers or your brand reputation. At the end of the day, it is your responsibility if your generative AI tool fails – not the company that built the underlying LLM or the customer using the tool.
- **Use banned words lists—but don't stop there.** Legal teams should start by creating keyword filters, but ethical hacking (red teaming) is the next step that is essential to catch deeper security vulnerabilities.
- **Double-check with AI:** Use "constitutional AI"—where a second AI reviews the first model's output—to flag harmful or incorrect responses.
- **Use licensed data, not free web data.** Open web scraping can lead to copyright-infringing outputs—opt for pre-cleared or proprietary data.

Top data security risks

Data security risks involve protecting sensitive business information and customer data when using AI systems. Breaches can lead to regulatory penalties and loss of customer trust. Here's how to safeguard your data:

- **Follow existing privacy laws.** Existing data privacy rules apply to generative AI too, so avoid sneaky data use and begin with ethical guidelines.
- **Avoid personal data in early models.** Start with anonymized data to reduce compliance risks.
- **Secure sensitive data.** If AI must process confidential information, use masking or pseudonymization to protect it.

- **Balance transparency and security.** Let users understand AI decisions without exposing proprietary model details.

Top legal and regulatory risks

Legal and regulatory risks involve navigating the complex and rapidly evolving landscape of AI laws across different regions. Staying compliant requires proactive planning and documentation. Consider these approaches:

- **Avoid high-risk applications.** Using AI for things like medical devices, finance or law enforcement means stricter regulations, so consider alternative applications. However, sometimes AI in these fields (like healthcare) can actually lead to saving lives. High risk = high reward–as long as you stay within the bounds of the law.
- **Document everything.** Keep detailed records of AI training data, model purposes and tool limitations to stay audit ready, even if there isn't a legal requirement to do so.
- **Transparency is key.** Make it clear when users are interacting with AI and disclose its limitations to avoid misleading expectations.

Turning generative AI into a scalable business asset

The transition from AI PoCs to AI products comes with challenges—but companies that tackle these challenges head-on, rather than waiting for others to lead the way, will be the ones that win.

At Publicis Sapient, we specialize in digital business transformation, helping companies curate enterprise data, prioritize AI use cases and build tailored strategies for sustainable success. Whether you're modernizing legacy systems or implementing enterprise AI solutions, our approach ensures your business is equipped to scale AI effectively and ethically.

From concept to enterprise-scale AI

Our Bodhi platform provides an enterprise-ready framework for developing, deploying and scaling generative AI solutions. With a structured approach to technology, operations and ethics, we help businesses move beyond experimentation and into AI-powered transformation.

By taking the right approach—grounded in strategy, security and scalability—organizations can unlock the full potential of Generative AI without unnecessary risk.

AI success isn't accidental—it's engineered. Let's build something transformative together.

The future of AI is being built today. Are you ready to lead the charge?

Related Topics

Artificial Intelligence | Risk Management