How to tame your first agentic AI use case

Business

99x.io

Contents

The team at 99x has had the opportunity to engage in several agentic AI projects implemented using Agentri, its agent orchestration platform. One common challenge many customers faced as they started the agentic journey was selecting their first use case. Selecting the first use case is critical as the outcomes will determine whether an organization builds confidence in adopting AI agents or becomes skeptical of the benefits. This article explores factors to consider when deciding on the first use case and what to watch out for. To provide context, it first starts with a brief introduction to agentic platforms and some recent implementations.

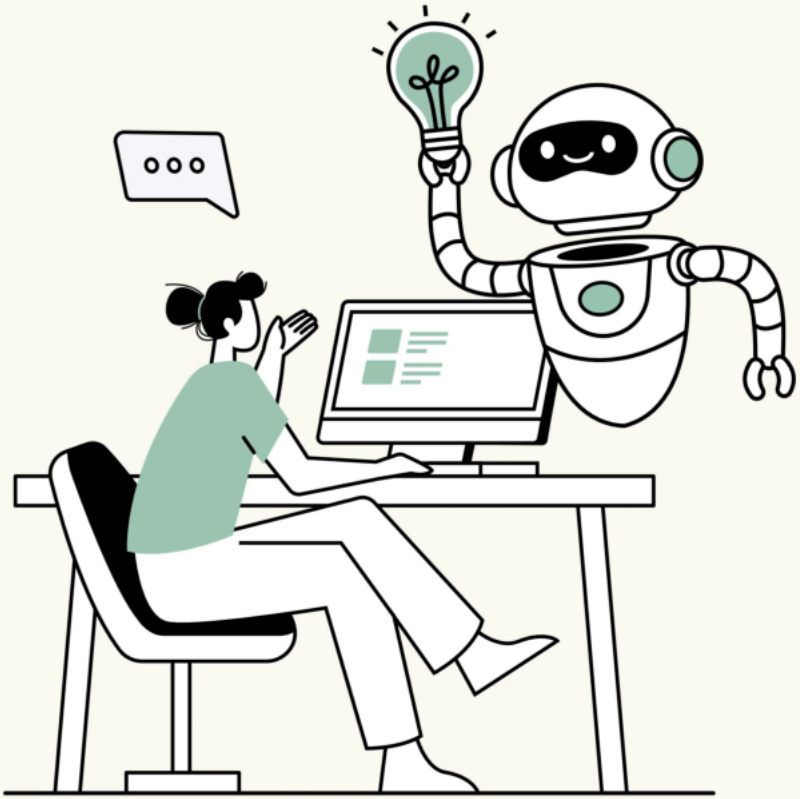

Moving from chatbot co-pilots to agentic AI autopilot

Chatbots co-pilots

- **Passive**
- **Assistive**
- **Limited autonomy**

Agentic AI autopilot

- **Independent lifecycle**
- **Assistive + Proactive**
- **High autonomy**

So, what has changed? On the surface, the user interaction elements of a conventional chatbot and an agentic AI application may look similar. Both accept text inputs, do some processing and then share outputs. While chatbots are passive, assistive and operate with limited autonomy, AI agents go beyond. They can listen and sense the environment, respond to trigger events, have an independent lifecycle even without human interaction, and operate with greater autonomy.

The result is agents able to work alongside humans, like any work colleague, able to sense events, reason implications, perform actions, communicate and ask for clarifications while working towards an outcome. This concept can be extended to have multiple AI agents who work with each other, deciding when to engage with a human, either to ask for assistance or provide an update.

Recent agentic AI implementations

The engagements at 99x span industries including financial services, legal, insurance, automotive and accounting. In customers' business interests, each of the use cases has been simplified to be more generic.

Percy the prospector: The first agent, Percy, enhances lead generation and market research internally at 99x. Percy continuously scans online news platforms, IT magazines and job portals for signals from the Norwegian tech landscape. It then evaluates companies against predefined qualification criteria to see if they are prospects for 99x. If so, Percy creates a comprehensive PDF report of the prospect, drawing from several online sources, and shares the report with the business development team for action.

Sales meeting preparation agent: Built for sales representatives in the insurance industry, this agent is triggered when a salesperson creates a calendar invitation involving a customer or prospect. The meeting request is scanned to identify attendee details. If it is a customer, the agent proceeds to extract the existing policy details and status, fetching this data from a CRM. It then creates a PDF with a consolidated view of the customer and shares it with the salesperson ahead of the meeting. If it is a prospect, the PDF is created using online sources.

Legal self-service agent: This agent empowers anyone to create a comprehensive legal document with the same confidence as working alongside a human legal advisor. For example, when building a software services agreement, the agent creates an initial draft and then walks the user through adding the next level of details and the most relevant legal clauses in this domain. Quality and checking for completeness is built into the entire cycle.

Dispute-resolution support agent: This agent works alongside customer service representatives or those handling back-office functions to reach greater levels of productivity. When a ticket is raised, the agent helps the customer service representative see a full view of the customer's issue, extracting data from multiple sources and building a case history. This empowers the human agent to close the support ticket faster.

Next, the article looks at some pointers on how organizations can decide on their first agentic use case.

Selecting your first agentic use case

Think of job roles and business processes instead of tasks

First, organizations should avoid the trap of building an assistant or copilot that assists humans only when triggered by requests or prompts. Instead, they should look at how a business process, a job role, or an activity performed within a job role can be done using an agent. For example, preparing for a sales meeting is a key part of a salesperson's job role. How can the time taken for preparation be reduced or optimized by using an agent? Are there any job roles that can be fully replaced by an agent?

One 99x customer identified the use of agentic AI in identifying and bidding for tenders. The first step was to identify when a tender announcement was made, then summarize the details to decide if it warranted a bid. The steps of identification and summarization were easily done by an agent and, on completion, the agent prompted the human-in-the-loop to decide on the next steps. In this instance, the agent functioned like a colleague who did the groundwork (or grunt work), increasing both job satisfaction and productivity. This can be extended across several scenarios.

Some examples:

▶ A sales assistant supporting a senior salesperson. Can an agent replace the 'sales assistant' role completely?

▶ Can an agent act as a legal advisor to guide the user to draft a legal document step by step, just as a human legal advisor would do? Once drafted, can the agent then warn the user of any clauses skipped or unspecified and the resulting implications?

▶ In the travel domain, can a set of agents help a user plan an itinerary where one agent manages preferences, another handles third-party bookings, and another takes over payments to complete the transaction?

▶ Can an agent perform complete competition monitoring and escalate when an important development is noticed?

Agents should therefore be thought of in terms of job roles and business processes and not in terms of assisting task execution.

Balance technical feasibility, value creation, and organizational readiness

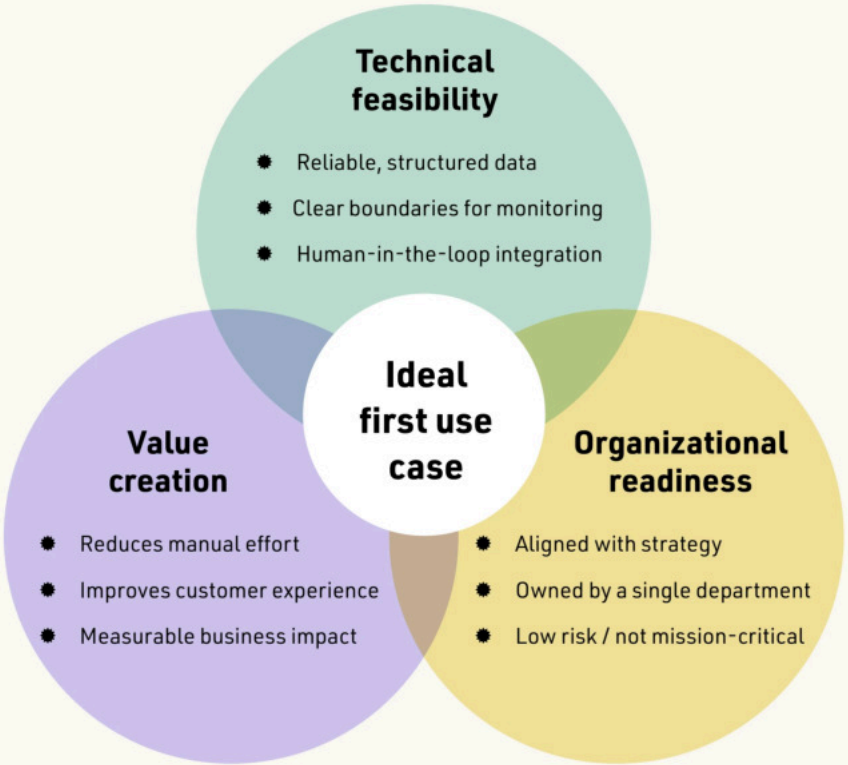

The first use case should solve a problem that matters. Organizations should look for processes or pain points with high visibility and measurable business outcomes. For example, reducing manual effort in customer support ticket triage. The use case must align with broader business strategy, whether that is improving efficiency, driving innovation, or enhancing customer experience.

Agentic AI excels in handling complex, multi-step tasks. However, as a first use case, it is important not to overreach. The problem chosen should be non-trivial but not mission critical. Ideally, it should be owned and executed in a single department instead of spanning multiple parts of the organization. A high-risk deployment can overwhelm stakeholders if things go wrong and be harder to troubleshoot. It is best to start with a problem that can be clearly bounded, where the agent's choices and outputs can be monitored and corrected.

Next, data availability and quality must be considered. Agentic systems thrive on context. The data required to fuel the use case must be accessible, structured, and reliable. Poor-quality data can derail even the most promising ideas. Use cases dependent on fragmented or confidential datasets that are difficult to access and integrate should be avoided.

Finally, the human-in-the-loop must be factored in, as agentic AI does not mean removing humans entirely from workflows. One of the strengths of early implementations is designing collaborative systems where agents take on repetitive tasks in a process while humans provide oversight.

As progress is made, organizations should look beyond the first implementation. A team of agents can be engaged as an agentic studio. The ideal first use case serves as a template for scaling agentic AI across the organization. For example, if an agent can successfully automate one type of compliance review, the same framework can likely be extended to others.

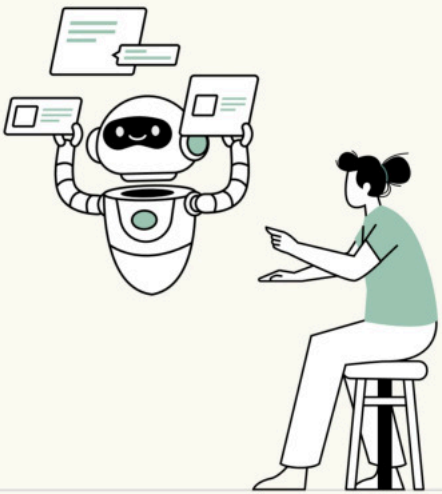

Defining success measures

Measuring success in agentic AI implementations is not just about technical performance, it is about proving business value, building trust and creating a platform to roll out new agents even faster. At a tangible level, success can be measured by improvements in efficiency and productivity. Key questions include:

- Are tasks completed faster?
- How many hours of human effort are eliminated?
- Does using agents allow some parts of the process to be done in parallel?

For example, qualifying and preparing a summary of each prospect previously took 60–90 minutes. The task was repetitive and could become monotonous after a few days. Percy the Prospector now performs this work 24x7 without complaint. The hours saved can be easily calculated. In practice, 99x was able to eliminate the need to hire a new salesperson for prospecting and market research.

By clearly defining success measures upfront, vague claims of efficiency improvements are avoided, and concrete evidence of value is provided. Percy the Prospector, for example, delivered 2x value in improving the number of leads logged into the sales pipeline.

Time (Efficiency & productivity)

- Tasks completed faster
- Parallel task execution
- Hours saved per workflow

(Example: 60–90 mins saved per prospect summary with Percy)

Money (Cost savings & ROI)

- Reduced hiring needs
- Clear financial impact
- Lower operational overhead

(Example: Avoided new sales hire)

People (Scalability & experience)

- Repetitive tasks automated
- Human effort redirected to higher-value work
- Builds trust in AI adoption

(Example: Percy running 24x7, freeing sales team capacity)

What to watch out for as your journey begins

While agentic AI offers tremendous promise, organizations should proceed with a measured approach. The proof of concept that impressed initially will need significant work before it performs reliably in a production environment. As such, the first engagement should not be overhyped but instead treated as an opportunity to prove value. The first use case should be ambitious enough to matter and be visible in the organization but safe enough to manage and tune for success.

Overhyping the PoC

- Proof-of-concept production readiness
- Stakeholder disappointment if outcomes fall short
- Risk of losing trust early

Scope & Safety Issues

- Use case too ambitious higher risk of failure
- Mission-critical processes shouldn't be the first test
- Lack of boundaries makes troubleshooting harder

Operational & Organizational Risks

- Missing data or poor integration
- No clear ownership (too many teams involved)
- Unrealistic timelines or expectations

Conclusion

Agentic AI is not another incremental step in technology. It is a fundamental shift in how organizations can operate right now, by empowering autonomous agents to take meaningful actions driving productivity and customer success. Yet the journey begins with one carefully chosen use case. Selecting wisely is about more than technical feasibility, it is about aligning with strategy, demonstrating measurable value, building trust, and creating a foundation for the future. Done right, it creates a launchpad to derive the most from AI-driven transformation. The key is to start small but with a meaningful engagement, learn fast, and expand with confidence.

About us

99x is a global product engineering company with decades of experience co-creating technology solutions for customers across Europe and beyond. With deep expertise in AI engineering, product development, and scalable delivery, 99x helps organizations turn bold ideas into market-ready products.

If you're considering your first agentic AI use case, the experts at 99x can guide you through the journey, from concept to implementation.

Lets talk Agentic AI or reach us at productengineering@99x.io