



**NOESIS**  
Technology for real growth

 **outsystems**

eBook

# The Impact Of Agentic AI

From agentic AI fundamentals to enterprise production:  
a practical implementation guide for regulated industries

# Contents

<b>The agentic AI opportunity</b>	<b>3</b>
<b>Agentic AI VS Generative AI</b>	<b>4</b>
<b>What defines an AI agent?</b>	<b>5</b>
<b>Single agent vs. multi-agent architectures</b>	<b>6</b>
<b>Agentic AI across industries</b>	<b>7</b>
<b>4 foundational steps to getting started with agentic AI</b>	<b>8</b>
<b>Why Noesis?</b>	<b>11</b>
<b>Why OutSystems?</b>	<b>11</b>

# The agentic AI opportunity

Artificial intelligence, particularly generative AI, has fundamentally reshaped how businesses operate, how software is built, and how companies connect with customers. Recent studies indicate that most organizations have already incorporated AI into their software development lifecycle processes and plan to increase their AI investment over the next 12 months.

**The next evolution is already underway: Agentic AI.** This technology uses autonomous agents to transform work by orchestrating end-to-end workflows, making real-time decisions, and executing complex tasks independently. The transformative impact becomes clear when 2 these systems augment human expertise and deliver true end-to-end automation.

However, going from pilot to full-scale enterprise production remains difficult. Bringing agentic systems to life requires specialized resources and introduces new risks. The market is saturated with thousands of AI tools that promise impact, but few deliver, leading to increased fragmentation and sprawl.

**At Noesis, we've implemented AI-powered solutions particularly in highly regulated sectors like banking, insurance, and healthcare.** Our experience reveals a consistent pattern: organizations struggle not with building individual agents, but with scaling them in environments where GDPR compliance, legacy system integration, and governance frameworks are mandatory rather than optional.

This guide shares what we've learned from moving agentic AI from experimental pilots to production systems that deliver measurable business value. Whether you're just starting to explore agentic AI or evolving your strategy, we'll walk you through what it is, where it's headed, and how to implement it successfully in enterprise environments where governance and compliance cannot be afterthoughts.



# Agentic AI VS Generative AI

To understand agentic AI, it's critical to grasp its relationship to generative AI. Generative AI responds to user prompts and creates content such as text, images, or code based on patterns in data. It enhances productivity by accelerating content creation and code generation, but it requires human direction for each task.

Agentic AI leverages generative AI's ability to understand natural language prompts, but goes further by carrying out sophisticated, multi-step tasks that integrate data, tools, and context to deliver outcomes with minimal human intervention. While generative AI produces outputs, **agentic AI perceives situations, reasons through options, executes actions, and adapts based on results.**

## The four key functions of agentic AI:



### Perceive

Rapidly ingest and aggregate massive volumes of raw data from multiple sources to support actionable insight.



### Reason

Apply advanced logic and analysis to interpret input, uncover patterns, and determine the best course of action based on available information and defined objectives.



### Execute

Make informed decisions in real time and take action to drive outcomes, whether that means updating systems, triggering workflows, or escalating to humans.



### Adapt

Continuously learn from every interaction, refining strategies and responses to perform better with each new data point and scenario encountered.

# What defines an AI agent?

AI agents are AI-driven components built to handle business tasks like processing data, responding to customer inquiries, executing workflows, or solving problems. To be truly defined as an agent, it needs autonomy and decision-making power. More specifically, an AI agent is:



## **Autonomous**

It can reason, plan, act, and evolve without constant human oversight, operating within defined boundaries and escalating only when necessary.



## **Contextually aware**

An AI agent modifies its behavior based on changing data, available tools, and evolving situations rather than following rigid scripts.



## **Purpose-driven**

The agent approaches tasks with clear goals and measurable outcomes, understanding what success looks like for each interaction.



## **Memory-informed**

It learns from past interactions and outcomes to improve performance over time, building institutional knowledge that transcends individual sessions.



## **Collaborative**

It works in tandem with humans and other systems to solve problems, recognizing when human judgment is required and when it can proceed independently.



# Single agent vs. multi-agent architectures

Choosing between single-agent and multi-agent architecture depends on the complexity and specific needs of your use case.

**Single-agent systems** have one AI agent that acts as the central intelligence for a specific task, with access to necessary tools and memory. This approach is generally less complex, easier to develop and manage, and uses fewer computational resources. However, single agents may struggle with highly complex or dynamic tasks and lack the ability to handle scenarios requiring diverse expertise or parallel processing.

**Multi-agent architectures** employ several AI agents working together to solve more complex problems. Often, a coordinator agent manages the overall operation while specialized agents, each with its own memory and tools, handle specific task groups. Multi-agent architectures excel at complex and dynamic tasks, can process information in parallel, and allow for smaller specialized models rather than one massive general-purpose model.

**In our implementations:** Simple and well-defined tasks such as document extraction or basic workflow automation suit single agents. Complex and dynamic scenarios require specialized knowledge and collaboration, such as customer onboarding processes spanning multiple systems or intelligent fraud detection across transaction types, benefit from multi-agent architectures.

For enterprises in regulated industries, the architecture choice also impacts governance complexity. Single agents are easier to audit and explain, which matters for GDPR and sector-specific compliance. Multi-agent systems require more sophisticated observability but deliver greater capability for complex business processes.


# Agentic AI across industries

This section highlights how Agentic AI can help in different industries:



## Banking and financial services

- › Monitoring transactions for fraud
- › Audit prep
- › Regulatory reporting
- › Virtual financial assistants
- › Banking chatbots
- › Personalized product offers



## Manufacturing

- › Monitoring inventory
- › Adjusting to disruptions
- › Placing orders
- › Predictive maintenance
- › Quality control
- › Equipment troubleshooting



## Government

- › Handling forms, applications, and permits
- › Cross-agency data analysis for generating insights
- › Proactive citizen outreach during emergencies



## Healthcare

- › Transcribing patient encounters
- › Virtual health assistants and triage bots
- › Automating billing coding and prior authorization



## Insurance

- › Triaging and automatically approving claims
- › Underwriting agents that evaluate risk and recommend pricing



## Retail

- › Dynamic pricing
- › Automated order fulfillment
- › Supply chain optimization



## Pharmaceutical

- › Adverse event monitoring and reporting
- › Regulatory compliance documentation
- › Supply chain tracking for temperature-sensitive products

# 4 foundational steps to getting started with agentic AI

As you look to harness the full potential of agentic AI, carefully prepare your operational environment in 4 steps.





## 1 | Start with a technical framework

---

### **Unify access to structured and unstructured data**

AI agents are as effective as the data they can access. To ensure effective reasoning and action, establish a foundational data layer so agents work from a single, trusted view of information across your enterprise.

### **Create a scalable cloud-native infrastructure**

Running AI agents on inflexible legacy systems limits their potential and slows innovation. Cloud-native architecture provides scalable capacity for growing workloads, ensuring AI agents can handle complex tasks as your program matures.

### **Build on a platform designed for the full AI agent lifecycle**

Disparate, one-off AI tools create fragmentation that impedes progress, complicates integration, and compromises security. A unified, AI-powered low-code development platform enables existing teams to rapidly build, customize, and scale AI agents and their applications while seamlessly integrating with existing core systems, ERPs, CRMs, and the data that resides in them.



## 2 | Build an AI-ready culture

---

### **Invest in AI fluency across roles, not just in IT**

To maximize AI agent value, promote organization-wide AI literacy. Educating employees helps them recognize how agents fit into workflows and identify high-impact use cases. This broad understanding enables faster opportunity spotting, misuse avoidance, and smarter collaborative decisions across business units.

### **Break down silos and promote collaborative experimentation**

Encourage cross-functional projects to leverage diverse perspectives and data sources for AI agents. This collaboration uncovers use cases and builds trust by demonstrating how combined expertise yields better results. Business units understand workflows and pain points; IT understands technical constraints and integration requirements. Effective agents require both perspectives.

### **Align leadership on a vision of agentic AI success**

Define and agree on measurable goals and success metrics for agentic AI before scaling investment. This top-down clarity fosters accountability and accelerates AI-driven growth. Leadership alignment prevents fragmented initiatives where departments pursue conflicting approaches, creating the agent sprawl that undermines enterprise value.



### 3 | Govern for trust, act with confidence

---

#### **Design ethical guardrails and clear fallback paths**

Successful agentic AI relies on trust. Establish policies defining agent capabilities and clear fallback procedures for uncertainty. This protects your brand, builds stakeholder trust, and creates a safety net for new risks. In regulated European industries, these guardrails must satisfy both internal governance and external regulatory requirements around algorithmic decision-making.

#### **Track and audit AI agent reasoning and behavior**

Transparency is key to understanding and improving agent performance. Implement monitoring tools to see how agents make decisions, what data they use, and potential biases. This strengthens compliance and helps refine agent behavior for better results. For GDPR-regulated environments, audit trails documenting agent reasoning are mandatory, not optional.

#### **Balance control with adaptability through human-on-the-loop practices**

While AI agents offer speed and scale, human oversight remains essential. A human-on-the-loop approach ensures outputs align with context, ethics, and business goals, allowing intervention when needed and enabling confident deployment in high-impact areas. This model scales better than human-in-the-loop approaches that require approval for every action.



### 4 | Start with relevant problems, then evaluate use case complexity vs. value

---

#### **Deliver measurable business impact**

Implement AI agents to solve real problems and deliver measurable results. Define KPIs upfront to track performance, validate value, inform investment decisions, and scale AI across the organization. In our implementations, the highest ROI comes from use cases with quantifiable impact: efficiency gains, revenue improvements, or measurable customer experience enhancements.

#### **Address pain points**

Identify inefficiencies, bottlenecks, or missed opportunities that cost time or money. Prioritizing these pain points secures quick wins, demonstrates value, and builds momentum for broader AI adoption. Start with high-volume, rules-based processes where manual effort is expensive and errors create compliance or customer satisfaction issues.

#### **Evaluate agentic AI solution providers**

Choose partners who align with your long-term AI roadmap, not just immediate needs. Prioritize vendors offering a unified, AI-powered low-code platform that supports the full AI agent lifecycle with enterprise-grade security, scalability, and governance, and seamless integration with diverse data sources and custom AI models. As OutSystems Premier Partners with three decades of experience in European regulated industries, we combine platform capabilities with implementation expertise in banking, insurance, and healthcare where governance is mandatory.

## Why Noesis?

With over 15 years as a trusted OutSystems partner, Noesis helps companies accelerate their digital transformation through enterprise-grade, low-code solutions. With a team of over 600 certified developers and more than 150 successful projects, we combine the AI-powered OutSystems platform with agile methodologies to build complex applications in just 6 weeks, when it traditionally takes 6 months. From strategy to ongoing support, we offer end-to-end delivery focused on generating real value, rapid return on investment, and competitive advantage.

## Why OutSystems?

The Noesis and OutSystems partnership dates back to the very first project developed with the low-code platform. Noesis is one of the first, and more experienced OutSystems partners worldwide, and counts with a dedicated team of more than 150 professionals, with over 230 certifications, to deliver custom made solutions in record-time, with the greatest business impact.



Talk with us and let's make the next success story

## Recognized by OutSystems





**NOESIS**  **outsystems**

[www.noesis-corporation.com](http://www.noesis-corporation.com)



© Noesis. All rights reserved.

eBook