Kognitos

WHITEPAPER

How to Trust and Govern AI Agents

Table of contents

THE PROMISE OF AGENTIC AI

Agentic AI is new on the scene, but already it's sweeping both nascent AI tech and legacy automation communities.

Based on a recent survey by Capgemini, **82% of executives** at large enterprises plan to integrate AI agents into their operations within the next three years.

Read More

In the context of enterprise process automation, agentic's quick-to-implement and adaptive attributes shine in comparison to traditional solutions like Robotic Process Automation (RPA), Intelligent Document Processing (IDP), and integrated platform as a service (IPaaS) tools.

These AI agents promise a lot, making them attractive to organizations looking to speed up their automation adoption timeline, but with so many tech companies making bold agentic claims or agent-washing their existing solutions, it can be challenging to cut through the noise to determine what is real and what is hype.

Moreover, some agents may also be too risky or inaccurate for a full-scale enterprise implementation. This eBook explores the brilliant potential of agentic AI while providing enterprise CIOs with a rubric to assess agentic vendors.

Understanding Agentic AI

WHAT IS AGENTIC AI?

Agentic AI in general refers to an emerging class of artificial intelligence systems characterized by their enhanced autonomy, decision-making capabilities, and adaptability. In short, AI agents are goal-oriented and make plans and take actions to achieve the goal. But, the definition of agentic AI has been fluid in recent months, as the hype sweeps through the community and marketers jockey for position. There are agents geared toward consumer tasks and agentic solutions focused on the enterprise, all with varying capabilities that are changing week-over-week.

In business processes, agents are capable of making human-grade decisions and performing tasks with little to no human intervention. These systems are designed to learn and adapt, making them suitable for dynamic environments like business processes with high variability (read “lots of exceptions”).

Perhaps the best attribute of agentic AI is that agents can be created rapidly and often without any conventional software programming. Organizations are looking to leverage agentic AI to enhance or replace traditional automation solutions for the benefits listed, and may also realize the additional benefit of a drastically lowered total cost of ownership (TCO).

PROGRESS AND ADOPTION

Today, only 10% of organizations are using AI agents, however, more than half of executives surveyed report that they will adopt agentic AI within the next year. This number rises to 82% over the next three years. This rapid acceleration in agentic AI implementation is due to several reasons:

Low Barrier to Entry: Agentic AI offers an accessible entry point into automation, reducing the headaches setup and maintenance costs associated with RPA or IDP and other tools

Self-Healing Capabilities: These systems can adapt to changes and recover from errors autonomously, reducing downtime and IT resource constraints

Supplementing Traditional Automation: Integrating agentic AI with the previous generation of automation solutions gives organizations greater flexibility and efficiency at lower costs

AGENTIC AI ADOPTION

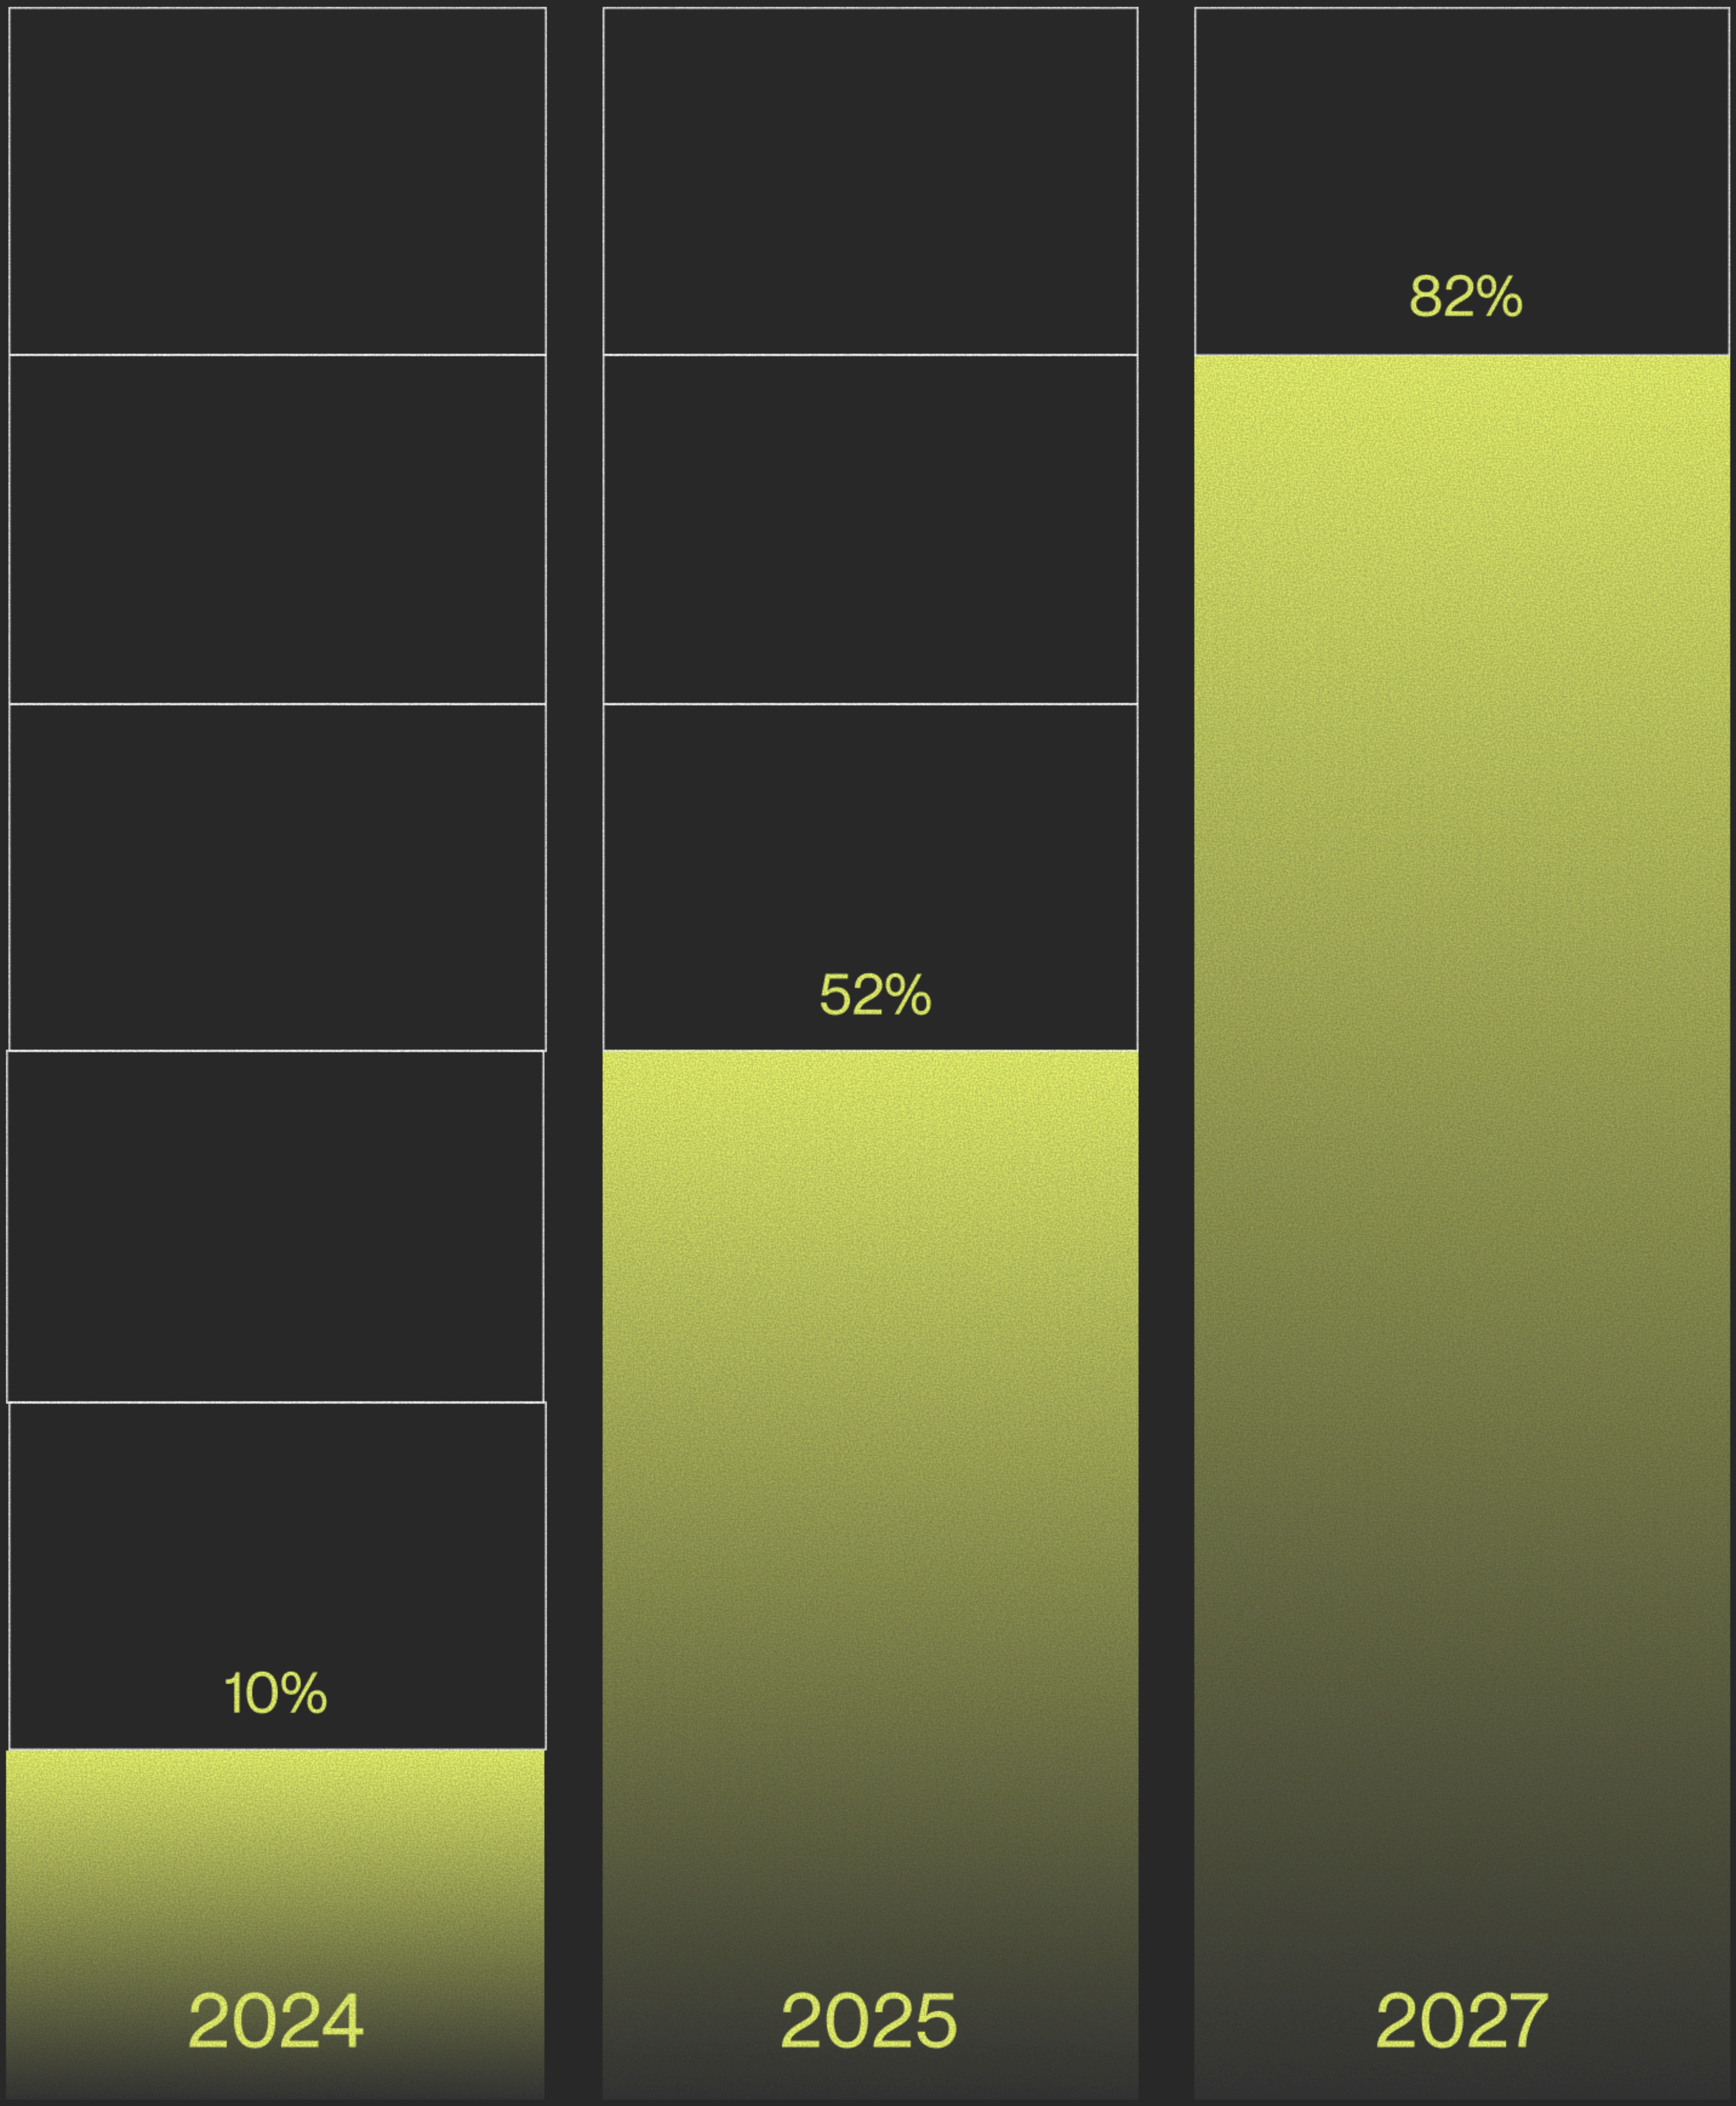

Source: Capgemini Research Institute, Generative AI executive survey, May–June 2024, N = 981 organizations who are at least exploring generative AI capabilities, excluding the public sector.

*The chart excludes 1% that answered unsure/don't know

Rubric for Assessing AI Agents

When evaluating potential vendors for Agentic AI solutions, CIOs should ask potential vendors questions related to each of the following key areas of consideration:

01 Control Mechanisms

There are effectively two types of AI agents: those that can act autonomously toward a goal, and those that achieve these same outcomes by following a process reviewed and approved by people. Defining the process controls the agents in and of itself, eliminating concerns about rogue agents with their own mind acting of their own free will. Much as businesses have enforced human processes for centuries, they should look to control agents via process, as well.

Despite this, McKinsey reports that 91% of leaders don't feel "very prepared" to adopt AI safely and responsibly. Ensuring human oversight and the ability to modify AI actions is crucial for maintaining control over automated processes. Consider these questions:

How does the solution ensure human oversight in decision-making processes?

What controls are in place to modify and review the AI's actions post-deployment?

02 Maintenance & Adaptability

Long-term success and adaptability depend on the ability to quickly and easily maintain and update AI systems. For example, RPA revolutionized automation when it was first introduced, but it requires a team of specialized developers and system upgrades to keep it running effectively as both the technology and the organization evolve. AI agents offer the antithesis to these challenges, but make sure to ask:

What strategies does the vendor offer for maintaining and updating the AI system?

How does the solution handle changes in business processes without major disruptions?

ARTICLE

Implementing generative AI with speed and safety

Read more

Rubric for Assessing AI Agents

03

Reliability & accuracy

Ensuring consistent performance and minimizing errors is critical, especially in high-stakes environments. One of the early concerns about agentic AI is that if given a command to complete a process 10,000 times, it may execute it slightly differently each time. That might not be a concern for creative use cases, but ask these questions in instances with real business implications:

What is the expected error rate for critical processes?

How are errors managed and recorded?

How does the system ensure accuracy in high-stakes environments like finance or healthcare?

04

Governance & Compliance

Enterprise-wide citizen development is a CIO's nightmare. When adopting an agentic solution, control and compliance are key considerations. Vendors should be able to provide compelling answers to these questions:

How does the solution prevent unauthorized creation of AI agents by employees?

What measures are in place to ensure compliance with industry regulations?

05

Learning & Evolution

One of the key differentiators of agentic AI solutions over legacy technology is the ability to adapt and improve over time. Questions to consider:

How does the AI system adapt to new information or changing business needs?

What is the vendor's approach to lifecycle management of AI agents?

A Trusted Path to AI Innovation

While agentic AI presents exciting opportunities for enterprise automation, it can be challenging to evaluate solutions and properly understand the risks associated with full-scale adoption.

Kognitos addresses these challenges through its neurosymbolic AI platform that combines symbolic logic with modern AI capabilities for AI you can trust, every time. The platform combines speed, adaptability, and comprehensive oversight, ensuring that enterprises can harness the power of AI without compromising control or reliability.

Legacy automation solutions are treating agentic AI as a buzzword and jumping to align themselves to the latest technology. In contrast, Kognitos uses modern AI, but with guardrails set in place by human stakeholders.

For forward-thinking CIOs looking to revolutionize their automation strategies, Kognitos provides a secure and trustworthy path forward. Embrace this opportunity to lead your organization into the future of AI automation.

Get in touch

Telephone	(800) 963-8112
Email	SALES@KOGNITOS.COM
Website	KOGNITOS.COM

Kognitos

Agentic AI Assessment Rubric

1 Control Mechanisms

How does the solution ensure human oversight in decision-making processes?

What controls are in place to modify and review the AI's actions post-deployment?

2 Maintenance & Adaptability

What strategies does the vendor offer for maintaining and updating the AI system?

How does the solution handle changes in business processes without major disruptions?

3 Reliability & Accuracy

What is the expected error rate for critical processes?

How does the system ensure accuracy in high-stakes environments like finance or healthcare?

How are errors managed and recorded?

Agentic AI Assessment Rubric

4 Governance & Compliance

How does the solution prevent unauthorized creation of AI agents by employees?

What measures are in place to ensure compliance with industry regulations?

5 Learning & Evolution

How does the AI system adapt to new information or changing business needs?

What is the vendor's approach to lifecycle management of AI agents?