



Q3 2025

TECH TRENDS REPORT

dexian®



Unlocking the Power of AI Agents While Protecting Your Assets

Earlier this year, OpenAI captivated tech journalists and enthusiasts alike when they previewed [Operator](#), an AI agent that can type, click, and scroll a web page. This “computer-using agent” captures what makes the next step in artificial intelligence so compelling. Operator acts autonomously, using the context and the tools available, to execute a task using its own reasoning.

Since then, AI agents have become the new bauble everyone wants to leverage::

- [PwC](#) found that 79% of senior executives already adopted AI agents
- [IBM](#) discovered that 99% of developers were exploring or developing AI agents.

AI agents are clearly shifting work dynamics, but we still aren't convinced this is the inflection point that puts humans out of work (besides, we don't even believe that's a good idea). An AI agent is just a tool. An advanced, sophisticated tool; but a tool all the same.

This report explores:

- **Intro to AI Agents: Key Applications**
- **How to Deliver Value with AI Agents**
- **The Intersection of AI Agents and Human Beings**

Let's explore the world of AI agents





Unlocking the Power of AI Agents While Protecting Your Assets

What is an AI agent? Simply, it is a program that can act autonomously to grasp context and interface with other tools to accomplish a user's requested task. AI agents incorporate function calling and multistep reasoning to large language models, unlocking a new level of independence when solving problems. In practice, they may ultimately serve as digital coworkers for people, capable of executing tasks with minimal human intervention, which is a major upgrade from generative AI.

Many AI agents within a business context fit into one of following categories, which we've outlined with some examples:

DELIBERATIVE AGENTS

Instead of kneejerk reactions to stimuli, deliberative AI agents are more methodical in their actions. They use internal models of the world, a defining characteristic of deliberative agents, to reason through options and pick a better outcome. They're more complex and a little slower, but they can gather data from environmental inputs to deliberate, make informed decisions using cognitive computing features, and learn over time.

Examples:

- Autonomous vehicles that plan routes, anticipate traffic, and make navigation decisions in real time.
- Supply chain systems that optimize scheduling, resource allocation, and logistics across networks.
- Healthcare decision tools that analyze patient data to recommend accurate diagnoses and treatments.

UTILITY-BASED AGENTS

These AI agents are focused not only on achieving a goal but doing so by following utilitarian theory and picking the best possible outcome. Utility-based agents evaluate multiple potential actions and choose the option that maximizes their overall utility.

Examples:

- Ride-sharing apps that assign drivers and routes to minimize fuel consumption and reduce emissions
- Smart grid systems that balance electricity demand and supply to minimize costs and outages.
- Hospital scheduling systems that balance staff availability, patient urgency, and resource constraints to improve care efficiency



LEARNING AI AGENTS

When we talk about dynamic tools that can interact with and adapt their behavior with each performance, these are the AI agents we have in mind. Learning AI agents diligently gather information, learn the correct path of choices to take, and refine their strategy based on evolving environments.

Examples:

- Energy management systems that optimize renewable power distribution by learning consumption patterns as they predict weather forecasts
- eCommerce platforms that learn through user behavior, engagement, and preferences to optimize pricing strategies, offers, and marketing to individual consumers
- Financial trading bots that learn how to evaluate market data, spot trends, and forecast future events before making real-time trading recommendations

MULTI-AGENT SYSTEMS

Just like the name implies, these systems involve multiple AI agents working together to solve a problem. They can collaborate, share information, or even test run competing strategies in complex situations. This system often benefits from AI orchestration, which a centralized agent oversees a squadron of AI agents with different domains or specialization.

Examples:

Here's how AI agents can help with freight and passenger train coordination:

- Scheduling agents can coordinate timetables to avoid conflicts across shared tracks
- Routing agents can adjust train paths in response to congestion or track maintenance
- Optimization agents can manage acceleration and braking patterns to reduce energy use
- Passenger service agents adapt carriage assignments and seat allocations to balance load across the network

Major AI companies are also exploring these tools within their own context. [Workday](#) has created a centralized management platform to orchestrate their own AI agents and third-party creations. [Salesforce](#) has released their own New Agentforce Command Center to optimize and scale the AI agents to enhance human productivity. [OpenAI's ChatGPT Agent](#) is orchestrating tools like Operator and Deep Research to creatively satisfy user requests.

The Dexian team is also experimenting with our own AI agents. One of the agents we've built can listen to a live interview between managers and candidates, evaluating the conversation and offering recommendations. We've built explainability into the model to ensure all actions and recommendations can be verified for their objectivity.

We've also built an AI agent that implement and configure ServiceNow for our customers. Our digital ServiceNow worker can gather requirements, generate detailed instructions, and configure systems and user access according to those requirements. This has resulted in a process that is 70% to 80% faster than human-driven processes. That's incredible progress, but requires a smart implementation framework to achieve.



How to Deliver Value with AI Agents

Though there is great potential for AI agents, their ability to generate value isn't foolproof. [Gartner](#) predicts that 40% of AI agent projects will be canceled by the end of 2027 because of unclear business value, insufficient risk control, and rising costs. Organizations that overcome that ambiguity and risk do so by following defined processes and measuring the performance for continuous improvements.

If your organization doesn't already have a clear roadmap or benchmarks in place, you can protect your investment with Dexian's methodology. First, let's look at the value framework necessary to achieve success.

Dexian's 5-Step Value-Led Approach to AI Agents

START WITH VALUE

First, prioritize your value dimension. How do you want your AI agent to positively impact your organization?



ENTRY POINT & USE CASE SELECTION

Even after you pinpoint your desired value, you need to select an entry point that can deliver quick wins and build organizational confidence. When possible, identify high-visibility pilot programs for your first use case.



PICK YOUR AI AGENT MODEL

You have your desired value and use case, but how will you access agents that align with your needs.



BUILD

Best when your processes are highly specialized and require custom-designed AI agents.



BOOST

Ideal when you want to extend the functionality of existing platforms without replacing them.



BUY

Most effective when speed to value is critical and proven solutions already fit your needs.

IMPLEMENTATION

As you're preparing to implement your AI agent, you need to cultivate a comprehensive view and understanding of wider implementation considerations.

DATA AND INFRASTRUCTURE

Garbage in, garbage out still applies to AI agents. Make sure you build secure, well-governed pipelines and infrastructure to support data quality and compliance.

APPLICATIONS

Your agent needs to integrate seamlessly into existing workflows if it's going to deliver clear business value. Verify that APIs, data sources, and other essential tools are accessible to function calling features.

APP FRAMEWORK AND AI MODELS

If you are looking for consistent and reliable outputs from AI agents, you need to pick the right framework. Review AutoGen, CrewAI, LangChain, and other frameworks to find which serves your complexity, usability, performance, and scalability needs.

FINE-TUNING AND PRE-TRAINING

Even though the goal of AI agents is to allow them to autonomously problem-solve, you'll get your desired results faster by giving it specialized training in your desired field. To do this, you can either license a pre-trained model or build your own. You can even use tools like Salesforce's [CRMarena-Pro](#) to create digital twins of your AI agents you can treat as a flight simulator.

RESPONSIBLE AI

More than creating AI agents that work, it's important to prepare them to make a positive impact. Incorporate the input of SMEs and everyday employees to promote ethical standards, transparency, and compliance from the start.



What about Measuring Results?

Measuring the ROI of AI agents means looking at both cost savings and strategic impact. That said, not all these benefits are immediately obvious and measurable? Freeing up people for high-value work is hard to measure, but it's definitely part of the ROI equation.

- Efficiency gains are often easiest to capture, as automation reduces processing time, frees employees for higher-value work, and minimizes errors that drive up costs.
- Beyond savings, agents can enable revenue by improving decision-making, optimizing supply chains, or enhancing customer interactions. Just know these benefits often emerge over a longer horizon.
- ROI also compounds with scalability, as single-use applications evolve into networks of agents that deliver enterprise-wide agility and faster time-to-market.

Taken together, these KPIs can help determine when your AI agents are on track to deliver value and when they need to be recalibrated before they waste too much of your budget and trust.





The Intersection of AI Agents and Human Beings

Where do humans fit in the equation? That's the billion-dollar question ([literally](#)). It's easy to imagine a world where digital workers will dominate the workforce, but plenty of real-world examples show how humans need to offer clear guidance, guardrails, reinforcement, and quality control to AI agents.

In July, an [AI coding agent from Replit](#) deleted a live database after running unauthorized commands and disregarding explicit instructions. Hallucinations and rogue policy decisions from Cursor's customer support AI tool compelled numerous customers to leave the platform. Even real AI agents with high quality and governance standards hallucinate and make mistakes.

That doesn't even account for the threat to cybersecurity as mismanaged AI agents gain or are given access to sensitive data. A [survey](#) earlier this year found that 63% of security leaders believed AI agents were one of their biggest cybersecurity threats.

Beyond the ethical implications, organizations are on the hook for the actions their AI agents take. When [Air Canada's chatbot](#) promised a passenger discounts on their flight that weren't available, the airline was held responsible in court for that commitment. AI agents will fall under that same expectation, requiring human oversight during crucial stages.

So, where should you have a human in the loop to interject? There are three key areas:

- Whenever an action for an important task cannot be easily undone (e.g. flight purchases, contract agreements)
- Whenever legal and regulatory compliance are involved (e.g. financial reporting, healthcare data privacy, safety standards)
- Whenever an AI agent needs human supervision to reinforce learning (e.g., reviewing flagged transactions, validating medical diagnoses, approving high-risk decisions)



Even when AI agents are capable of handling a task without much supervision, human users need to put them on the right footing. For example, prompt engineering is still important for AI agents. Communication is inherently contextual, which means organizations need to clearly communicate best practices for developing prompts for agentic AI. This includes keeping prompts concise, grounding them in domain-specific context, and testing variations to ensure consistent, reliable outputs.

Another example is organizational adoption: even the most advanced AI agents won't deliver value if employees don't trust or understand them. Companies need to invest in ongoing training (or even formalized upskilling programs) to establish clear roles for how humans and agents collaborate. Plus, creating feedback channels and communicating where users can share their challenges and thoughts on improvements can refine these tools in real-time.

At the end of the day, implementing clear change management for humans and AI will determine whether you unlock an AI agent's full value together or waste money on your investment.

Key Takeaways for Q3 2025

- AI agents are autonomous programs that use context, tool integration, and multistep reasoning to execute task, ultimately serving as digital coworkers that go beyond traditional generative AI.
- The major types of AI agents include deliberative AI agents, utility-based agents, learning agents, and multi-agent systems
- 40% of AI agent projects will be canceled by the end of 2027 because of unclear business value, insufficient risk control, and rising costs. A clear implementation roadmap can mitigate those risks.
- Organizations need to keep humans in the loop with AI agents' decisions to offer clear guidance, guardrails, reinforcement, and quality control to AI agents.

Learn how to responsibly harness AI agents. **Partner with Dexian Today.**

[EXPLORE OUR ARTIFICIAL INTELLIGENCE SOLUTIONS](#)