Multi-agent AI – 21st century automation revolution

Capgemini

Automation has been the fundamental technology underpinning economic transformation for centuries.

Whether it was Britain's late 18th and 19th century industrial revolution, the United States' post-World War II boom, or South Korea's industrialization starting in the 1960s, all benefited from automation to increase productivity, efficiency, and profitability. The effect was radical transformation of economy and society.[1] Now, in the 21st century, a new wave of automation is underway as agentic AI use widens across economies.

An *artificial intelligence agent* is a software program that can interact with its environment, collect data, and use this to autonomously perform tasks to meet predetermined goals. As an evolution from technologies like robotic process automation (RPA) and machine learning (ML), AI agents can, perceive, reason, and act in changing environments to achieve their goals. How they reach them is largely left to them to decide.

AI agents employ a range of advanced technologies to interact with users and perform tasks autonomously and effectively. Large language models (LLMs) are often the primary interface between AI agents and users. They are a type of foundation model trained on vast datasets, primarily text. LLMs encode knowledge by recognizing patterns, enabling AI agents to reason, inform their decision-making, and communicate.

An agent can understand and generate human-like text or verbal responses using natural language processing (NLP), making human-AI interactions more natural and efficient.

[1] D. Acemoglu and P. Restrepo, Journal of Economic Perspectives, Vol. 33, Automation and New Tasks: How Technology Displaces and Reinstates Labor, pp. 3–30

Agentic AI doing what it does best

Since generative AI took off as a popular phenomenon, companies have rushed to create their own versions. They are used for advanced search, analysis and interaction with documents, particularly in the legal, HR, and technology fields. While these improve on previous systems, to stop at this point underestimates the full potential. LLMs are improving iteratively – where two years ago an LLM with retrieval-augmented generation (RAG) could produce synopses – current LLMs using enhanced retrieval methods can generate more sophisticated output.

The fundamental difference between using a standalone LLM and employing a multi-agent system is that with the latter, individual agents are created to specialize in specific tasks – often not limited to language – and can collaborate with each other. They can execute more complex tasks and integrate with external tools such as web searches, APIs, and dedicated databases.

Tangible and intangible *value*

Why should enterprises go the agentic AI route? In short, because it will soon be in every business function where it's feasible.

82% of organizations plan to integrate AI agents by 2027.[2]

There are more reasons to move to agentic AI than cost-saving alone.

According to a UK Customer Satisfaction survey, a defining feature of high-performing organizations is *"an appropriate balance of people and technology, combining speed, efficiency, and personal care."*

In other words, the ideal retail customer service experience is a blend of human empathy and technological efficiency.[3] Businesses can therefore use agentic AI to improve and differentiate their offer to customers ahead of competitors, e.g., by adding communication channels and styles that appeal to specific customer bases.

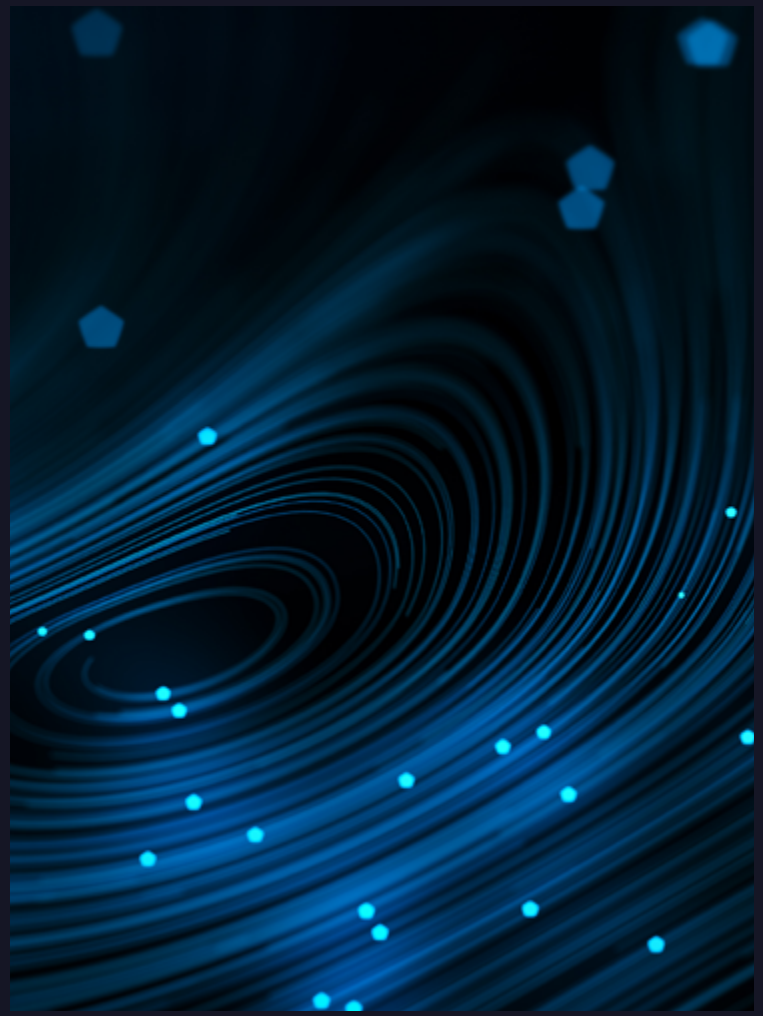

Companies are re-organizing their customer and IT support services to raise quality standards by augmenting them with AI agents in hybrid models. For example, an AI agent can automatically draft responses to customer queries based on historical customer interaction data.

More ambitiously, agents can take ownership of a client issue. Previously, if a customer contacted a chatbot to request a refund for a product and the request was non-standard, the chatbot would escalate the issue to a human customer-service representative. Now, an AI agent can request more information from a customer, for example, proof of purchase or a photograph of an item, analyze the enquiry and offer a possible solution. It can decide alone to override standard procedure, if circumstances justify making an exception. It is highly likely to resolve the matter without human input.

[2] Capgemini Research Institute, Harnessing the value of generative AI: 2nd edition, 2024

[3] UK Index (UKCSI), January 2025

Design for *maximum return*

An AI agentic system increases overall productivity, service quality, customer satisfaction, and loyalty. The design of agent-to-agent and agent-to-human exchanges according to frameworks, rules, risks, and protocols is crucial for this.

Unlike predecessors, AI agents are:

- Autonomous
- Goal-oriented
- Context-aware, using relevant data to make decisions
- Adaptive, adjusting behavior and responses as data or interactions change
- Proactive, initiating action independently without user prompts
- Language-aware, interpreting and responding in human language

Agent **creation**

To build AI agents requires:

1. Defining their roles,
2. Identifying and locating the data they will use,
3. Defining which tasks or goals they execute,
4. Setting boundaries with guardrails.

With multiple agents, each has its own specialized role, while it cooperates with others in a decentralized structure. They can solve more complex tasks collaboratively, such as processing insurance claims: one agent verifies documentation, another evaluates policy criteria, and a third processes payments, completing the task jointly.

Typical goals for agents:

- Customer service
- Content recommendation
- Predictive analytics
- Fraud detection
- Real-time monitoring
- Inventory management
- Scheduling and task management
- Orchestration in manufacturing

The *agentic* system

The visual representation of an agentic process can closely resemble traditional organizational process diagrams. As organizations transition toward agentic systems, their business experts should collaborate closely with AI specialists to effectively design and streamline these processes. However, integrating AI agents into existing systems can be complex, disruptive, and destabilizing to operations if not managed carefully.

Building an architecture that accurately reflects real-world activities first requires creating digital descriptions and definitions of business operations. Clearly defined tasks can then be mapped to AI agents as needed. This approach contrasts with traditional methods, which impose a predefined architecture onto a business and archetypes onto its people.

In an agent-first enterprise, modeling agentic interaction involves visualizing communication and transactions between internal and external parties. This includes understanding the underlying motivation, similar to traditional business process modeling. Organizations should map all such interactions to accurately represent the entire value network, not only the value chain.

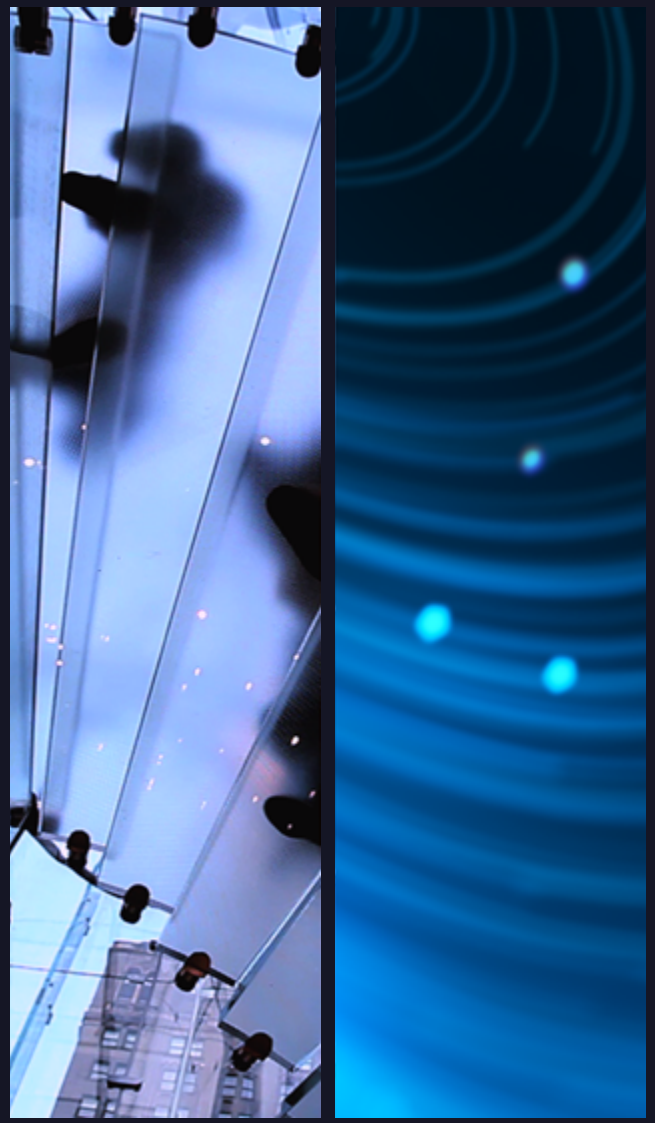

Data as *AI oxygen*

Data requires optimization as the foundation for the entire agentic architecture, since its fragmentation blocks AI agents from working effectively. 80% of organizations store more than 50% of their data in hybrid and multi-cloud infrastructures,[4] which complicates integration, availability, and management. This can have an impact on AI agent deployment and effectiveness.

To optimize data for agentic use means:

- Evaluating its quality
- Establishing governance, management, and security
- Creating a pipeline to ensure real-time or near-real-time availability of high-quality data
- Continuous monitoring
- Ongoing improvement using feedback loops

[4] Tremblay, T., Kohezion, What is data fragmentation?, September 2, 2024

Evolution towards agentic AI

Across sectors and industries, whichever domain decides to implement agentic AI architecture, a similar process applies across specializations to maximize returns on investment.

Leaping into agentic AI immediately without first evaluating simpler automation options is akin to creating a solution in search of a problem: RPA and scripts alone can achieve automation for tasks that have defined stable inputs and treatment. Agentic AI should be adopted specifically for achieving hyper-automation – advanced, next-generation automation for tasks that have variable inputs and variable types of treatment.

1. **Map the processes** – conversion to agentic AI requires preparation for automation. Processes are broken down in sub-processes, macro tasks, and micro tasks.

   When the framework is applied to a financial and accounting digital operation such as a credit-to-cash process, it can be broken down into:

   6 sub-processes | 12 macro tasks | 57 micro tasks

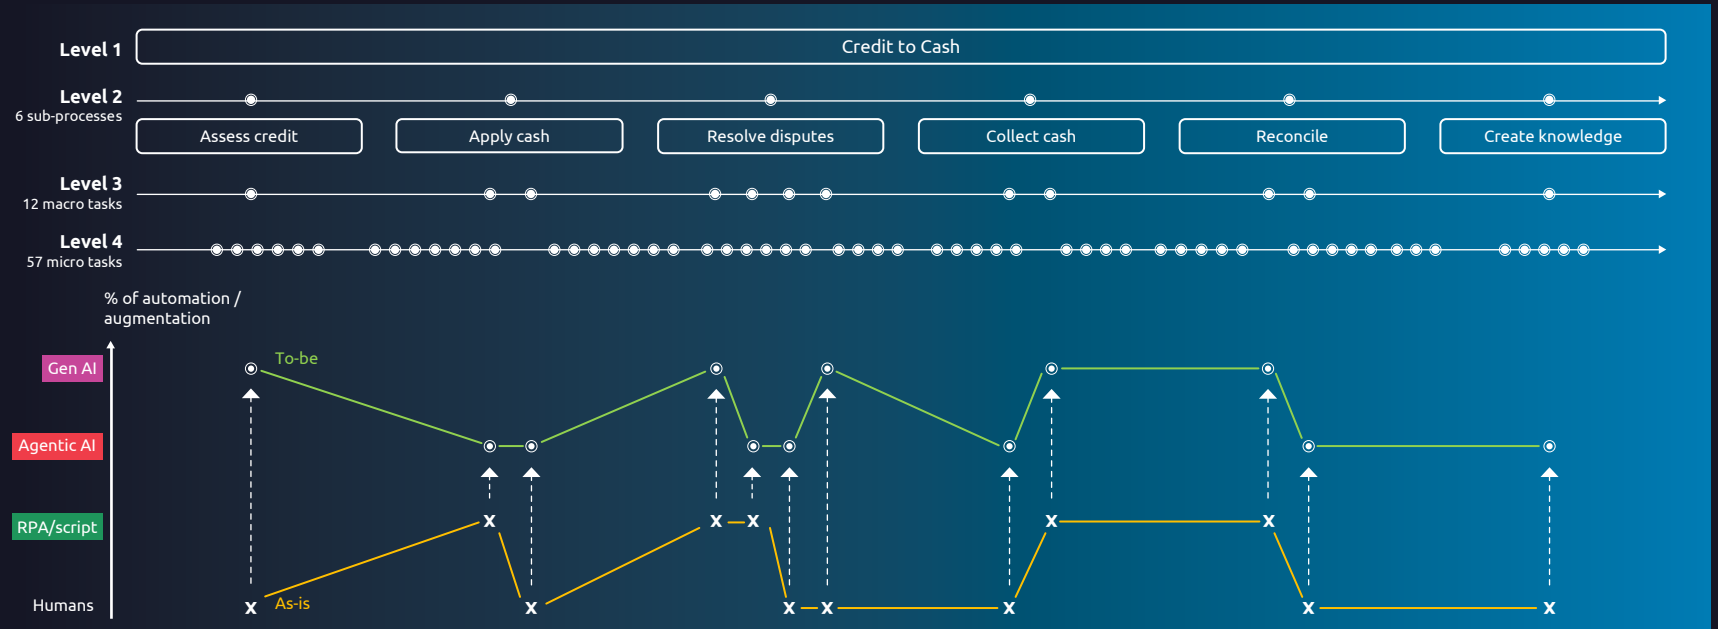

Figure 1 Example of accounting credit-to-cash process, with different levels of automations (RPA/script, Agentic AI, Gen AI)

2. **Convert** existing architecture into automation components starting with RPA/scripts for tasks with defined and stable input, then hyperautomation with Agentic AI for tasks that have variables inputs and variable types of treatment
3. Finally, for tasks that have not been automated, **augment** human workers with AI assistants powered by Generative AI.

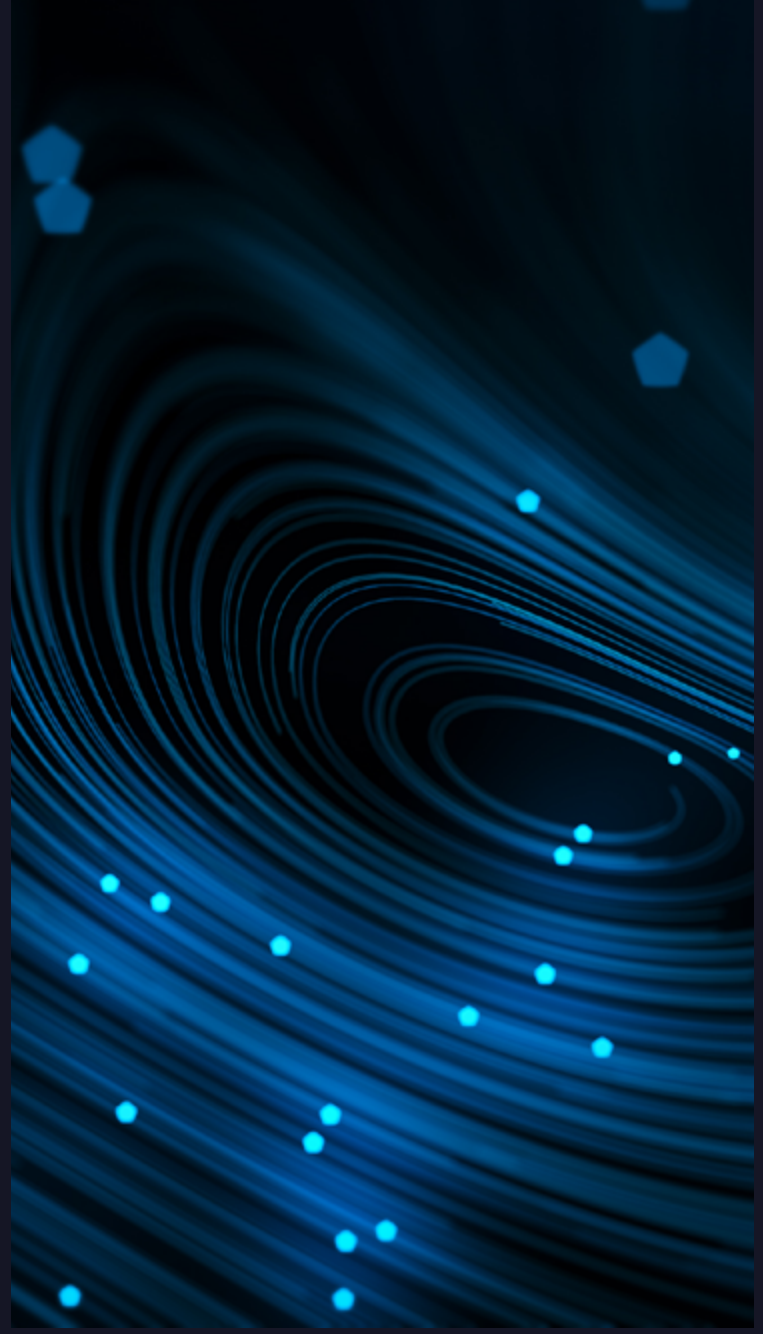

Governance, guardrails, and responsibility

As mentioned earlier, it is a multi-agent system's capability for specialized roles, distributed decision-making, autonomous coordination and collaboration that distinguishes it from a non-agentic Gen AI system. Its greater complexity necessitates specific governance strategies.

In a multi-agent system, agents are monitored while "trusted" to act independently, with selective use of human oversight. AI agents require clearly defined governance frameworks that specify when and in what form human authorization is necessary, for example:

- During initial configuration
- At critical decision points in a workflow
- Before interacting with sensitive data
- In operations involving other agents and humans

Human intervention is a safeguard in case AI agents' decision-making is biased, inaccurate, or breaches company ethics. Such "misbehavior" could damage client and employee trust, negatively impact brand reputation, or result in legal violations.

Testing for compliance and failure, including for bias, fairness, and operational performance, is essential. The aim is to establish points of failure and define boundaries for guardrails. The results should demonstrate agents' compliance with regulations and be the basis for contingency plans in case of failure.

The autonomy of AI agents makes it more complex to trace errors and determine their root causes. When there is a fault, it is not a breakage in classic IT form. Issues are harder to replicate, given the complexity, sophistication and uniqueness of each instance of action. Agent activity must be systematically logged, capturing performed tasks, actions taken, evaluation metrics, and the agent's internal state for effective monitoring and error tracing. These are defined and tested in the proof-of-concept stage and updated as needed during the agent's whole lifecycle.

Orchestration and integration

Processes with potentially thousands of largely autonomous digital workers will need careful orchestration. Success depends on maximizing automation within a well-designed architecture. Without it, there is risk of breakdown or disruption.

The predicted wide use of agentic AI will push IT specialists towards developing skills to support an agentic AI system by onboarding, managing, and training AI agents as “digital workers”.

Managing an agentic AI system requires:

- Configuring agents for specific roles, similar to recruiting and onboarding human employees
- Checking that agents comply with relevant laws, including the EU’s GDPR and AI Act
- Designing human-agent collaboration
- Leading change management with employee training on AI adoption

The sectors and industries already a *step ahead*

The arrival of agentic AI in economies around the world will prompt organizations to review their processes for suitability and potential gains in productivity and cost saving. The following sectors and industries are already on an agentic AI journey, but more will join as expertise in adoption spreads.

Consumer

AI-powered interactive home assistant devices used to oversee the elderly and infirm, locate mislaid items, and monitor home security

Manufacturing

Smart camera-based process monitoring for improved shopfloor performance and safety compliance

Life Sciences

- Drug discovery support to extract actionable insights from drug mechanisms, disease progression and clinical outcomes
- Refining clinical trial design and monitoring real-time data for mid-trial adjustments

Financial Services

- Fraud detection agents identifying suspicious transactions and initiating appropriate responses
- Financial planning and investment management services creating personalized investment strategies and dynamically monitoring client portfolios

Retail and Supply Chain

AI-driven agents monitoring shelves in-store and warehouses and automatically trigger stock replenishment using stock-keeping unit (SKU) codes

Experts to contact

Itziar Goicoechea
AI Agents for Enterprise Offer Leader

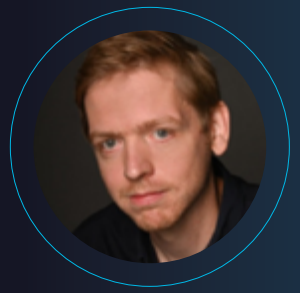

Mark Oost
AI, Analytics, Agents Global Leader

About Capgemini

Capgemini is a global business and technology transformation partner, helping organizations to accelerate their dual transition to a digital and sustainable world, while creating tangible impact for enterprises and society. It is a responsible and diverse group of 340,000 team members in more than 50 countries. With its strong over 55-year heritage, Capgemini is trusted by its clients to unlock the value of technology to address the entire breadth of their business needs. It delivers end-to-end services and solutions leveraging strengths from strategy and design to engineering, all fueled by its market leading capabilities in AI, generative AI, cloud and data, combined with its deep industry expertise and partner ecosystem. The Group reported 2024 global revenues of €22.1 billion.

www.capgemini.com

Capgemini

Get the future you want