Elsewhen Reports

# Building the Agentic Enterprise

## AI Agents & Multi-Agent Systems

by Leon Gauhman & Vahid Panjganj

# In today's enterprise, it's common to have hundreds of software applications, most of which operate in silos, creating a tangled web of inefficiencies that hamper productivity.

AI-powered multi-agent systems (MAS) and AI agents address these inefficiencies by delivering three core advantages to enterprises: increased productivity, cost-efficiency, and adaptability. These capabilities form the foundation of what we call the agentic enterprise — an organisation where AI-driven systems collaborate seamlessly, adapt dynamically, and operate with a degree of independent reasoning.

By enabling real-time data flow and decision-making across previously isolated systems, these networks streamline workflows and reduce manual intervention, freeing up valuable time and resources. Their adaptability allows organisations to build solutions tailored to specific operational needs, reducing reliance on rigid, one-size-fits-all SaaS tools. By moving towards proprietary multi-agent systems and setting the groundwork for the agentic enterprise, businesses can build flexible, customised workflows that better align with their goals.

MAS are composed of AI agents, which range from basic task-specific tools to more advanced autonomous agents capable of agentic reasoning, planning, and adapting to changing conditions. While not all AI agents are necessarily fully autonomous, the integration of even basic agents can streamline operations and improve efficiency across complex workflows.

The adoption of AI agents is happening faster than you may think. According to a recent survey of over 1,300 professionals, 51% of companies already have AI agents in production, with 78% actively planning to implement them soon. Mid-sized companies (100–2000 employees) are leading the way, with 63% already deploying agents. Interestingly, this momentum isn't confined to the tech sector; 90% of non-tech companies are using or planning to use agents, nearly matching the adoption rates of tech companies at 89%.

"The adoption of AI agents is happening faster than you may think."

Elsewhere, with the arrival of multi-agent frameworks like OpenAI's Swarm, LangGraph, and ReAct, enterprises can now bridge disconnected agents through intelligent, collaborative systems, creating a cohesive network that transforms decision-making and operational efficiency. These frameworks are driving multi-agent systems to become increasingly agentic, meaning they are powered by autonomous agents capable of reasoning, planning, learning from past interactions and communicating with each other. This shift allows MAS to go beyond simple task coordination, fundamentally reshaping how software integrates within enterprise environments.

At the same time, Stripe now lets AI agents make controlled purchases using temporary virtual cards — another sign they're moving from concept to reality.

In this white paper, we outline the strategic opportunities unlocked by AI agents and multi-agent systems and explore how they lay the foundation for the agentic enterprise.

— Leon Gauhman

# Contents

# AI Agents

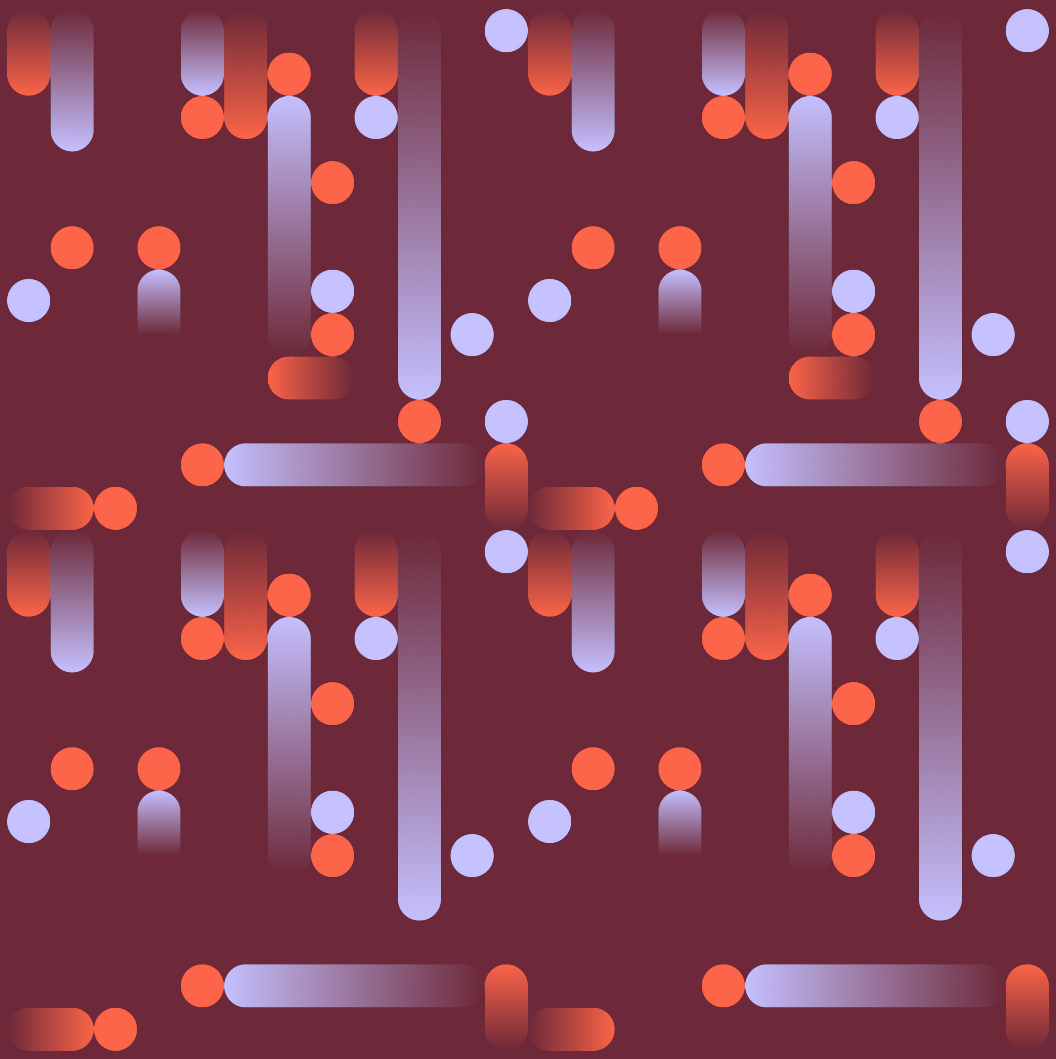

## What exactly is an AI agent? In essence, AI agents are digital systems designed to perform tasks, make decisions, and solve problems by interacting with their environment — sometimes fully autonomously and sometimes with a human-in-the-loop.

These agents, often serving as copilots, take on an assistive role, automating routine tasks, processing data, and providing recommendations. By adopting a 'jobs-to-be-done' perspective — focusing on the specific tasks they are designed to accomplish — they are seen as complementing rather than replacing human roles.

In our previous report on Generative UI, we explored how AI-powered interfaces simplify complexity and make cutting-edge AI tools accessible to the broader workforce. These assistive agents dynamically generate user interfaces, offering real-time responses tailored to specific tasks, further reducing cognitive strain and enhancing decision-making.

Small Language Models (SLMs) are compact and efficient AI systems designed to perform specific tasks with high speed and low resource requirements. Unlike their larger counterparts, SLMs excel in targeted applications, particularly in environments where computational power and memory are limited. This makes them ideal for edge deployments, where AI operates on devices like smartphones, IoT systems, or embedded hardware rather than relying on centralised servers or cloud computing.

Despite their size, many SLMs now outperform OpenAI's original GPT-4 LLM. In multi-agent systems, SLMs complement larger models by handling routine, resource-light tasks, ensuring that the overall ecosystem remains scalable, adaptable, and cost-effective.

Central to the success of these agents is selecting the right model for each job. Rather than relying on large, resource-intensive models for every task, multi-agent systems benefit from a range of model types that suit different functions. Small Language Models (SLMs), for example, are especially well-suited for targeted tasks, delivering advantages in speed, energy efficiency, and lower memory requirements. This makes them ideal for tasks requiring rapid, specific responses where the power of a full-sized Language Model (LLM) is unnecessary.

For instance, SLMs can power prompt-and-response agents that handle straightforward interactions, offering real-time processing that fits seamlessly within an agent ecosystem while LLMs take on sophisticated tasks, such as analysing large datasets or generating in-depth insights. By focusing on specific language tasks, these smaller models achieve faster response times and lower costs, providing practical and efficient solutions for businesses.

Far from replacing people, AI agents act as master orchestrators within the enterprise. Operating on behalf of the organisation, they streamline repetitive tasks, integrate across systems, and free employees to focus on higher-value work, fostering a colla§borative, intelligent ecosystem.

But this is just the beginning. The agentic enterprise will be built around a constellation of AI agents — ranging from simple prompt-and-response tools to semi-autonomous and fully autonomous systems — all working together to drive productivity.

Far from replacing people, AI agents are master orchestrators within the enterprise.

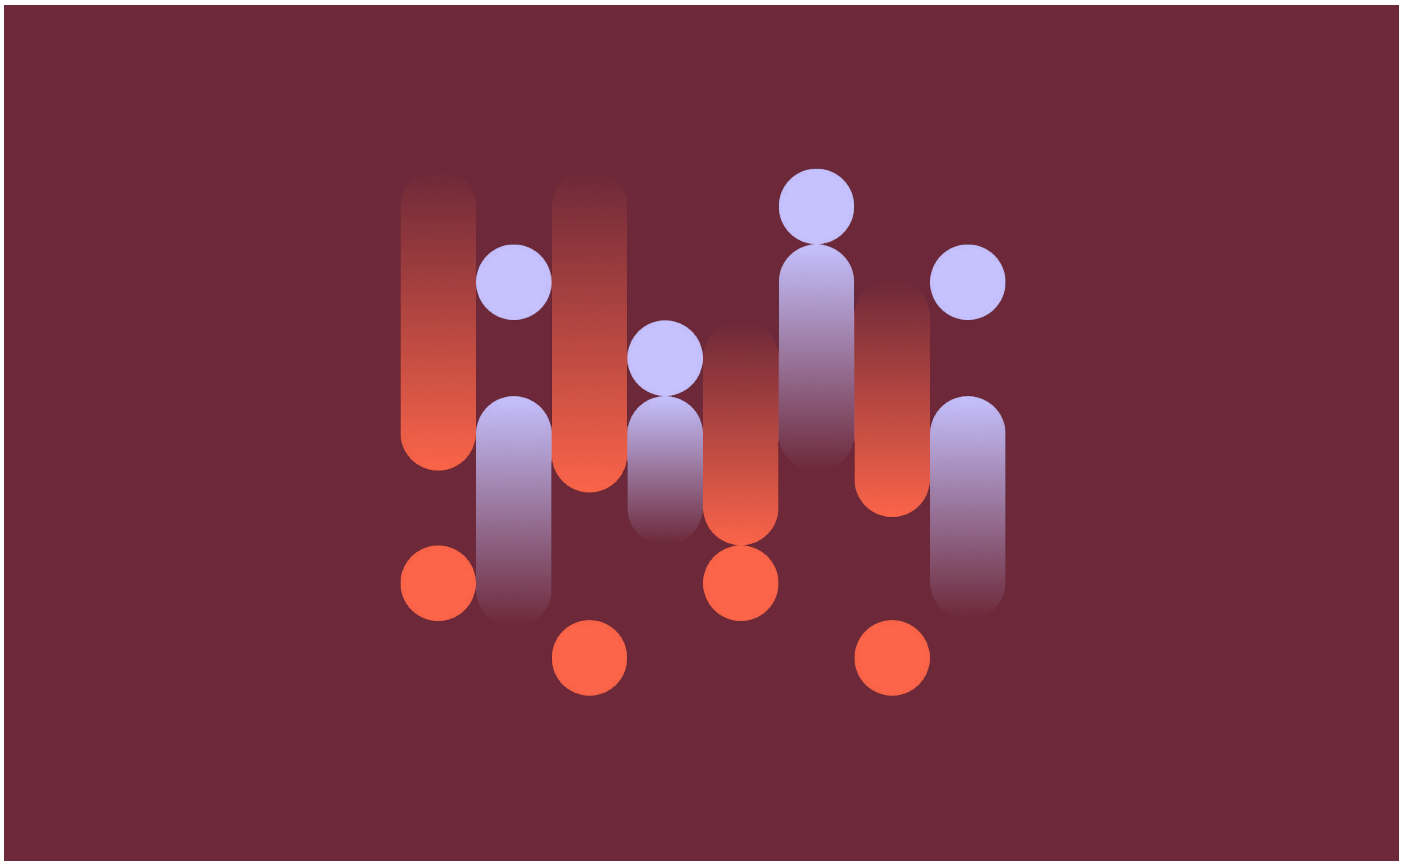

# Semi-Autonomous Agents and Human-in-the-Loop

Central to the success of multi-agent systems is the ability for agents to interact with their environment in dynamic and intelligent ways. Claude 3.5's 'Computer Use' feature from Anthropic is a prime example, demonstrating how an AI agent can control a user's desktop, move the cursor, and type in real time — executing complex, multi-step tasks based on human instructions.

It is also worth mentioning here how the "Computer Use" feature highlights the emerging capabilities of multimodal agents. Unlike traditional LLMs confined to text-based interactions, these agents can interpret and act upon various forms of data — such as text, images, and screenshots. According to initial reports, Claude 3.5 selects from pre-defined tools (e.g., clicking, typing) based on the user's prompt, generating structured commands that allow it to seamlessly interact with computer interfaces.

'Computer Use' works by relying on an iterative agent loop, where each action is executed, analysed, and fed back to the model to determine the next step. This allows Claude to perform tasks like navigating websites, interacting with desktop applications, and gathering information from documents or spreadsheets — essentially bridging the gap between AI and human-computer interaction (HCI).

Despite its advanced capabilities, Claude's "Computer Use" feature does not yet reach full autonomy because it cannot independently define objectives or take action without human prompting. While Claude can perform complex tasks and execute commands efficiently, it requires explicit instructions from a human operator to initiate actions. This means Claude operates within the scope defined by user input, responding to requests rather than proactively engaging in tasks (unlike a fully autonomous system, wherein the agent would be capable of self-directed decision-making, more on this below).

Amazon Q's Developer Agents are another strong example of semi-autonomous agents, providing productivity boosts by performing complex, multi-step tasks in software development autonomously. These agents, used by companies like BT Group and National Australia Bank, generate and troubleshoot code, compile updates, and even flag security issues independently before presenting their results for human review.

More recently, Nubank leveraged Devin — perhaps the most well-known commercial agent — to accelerate one of its largest engineering migrations. Faced with the challenge of refactoring millions of lines of code, Nubank initially planned a multi-year effort requiring over 1,000 engineers. With Devin, engineers could delegate much of the tedious refactoring work, achieving an 8-12x efficiency improvement and 20x cost savings.

Like Amazon Q, Devin operates semi-autonomously — executing complex, discretionary tasks but keeping a human-in-the-loop for oversight and validation. It learns from past migrations, optimises its approach, and even generates its own helper scripts to improve efficiency. However, final approvals remain with engineers, ensuring quality control and mitigating risks.

This approach enables semi-autonomous agents to manage lower-stakes, repetitive tasks independently, while deferring high-stakes decisions to human oversight. By maintaining this balance, organisations ensure that quality control remains intact while maximising efficiency. The result is a model where AI accelerates execution, but humans retain control over critical decision-making — unlocking productivity gains without compromising reliability.

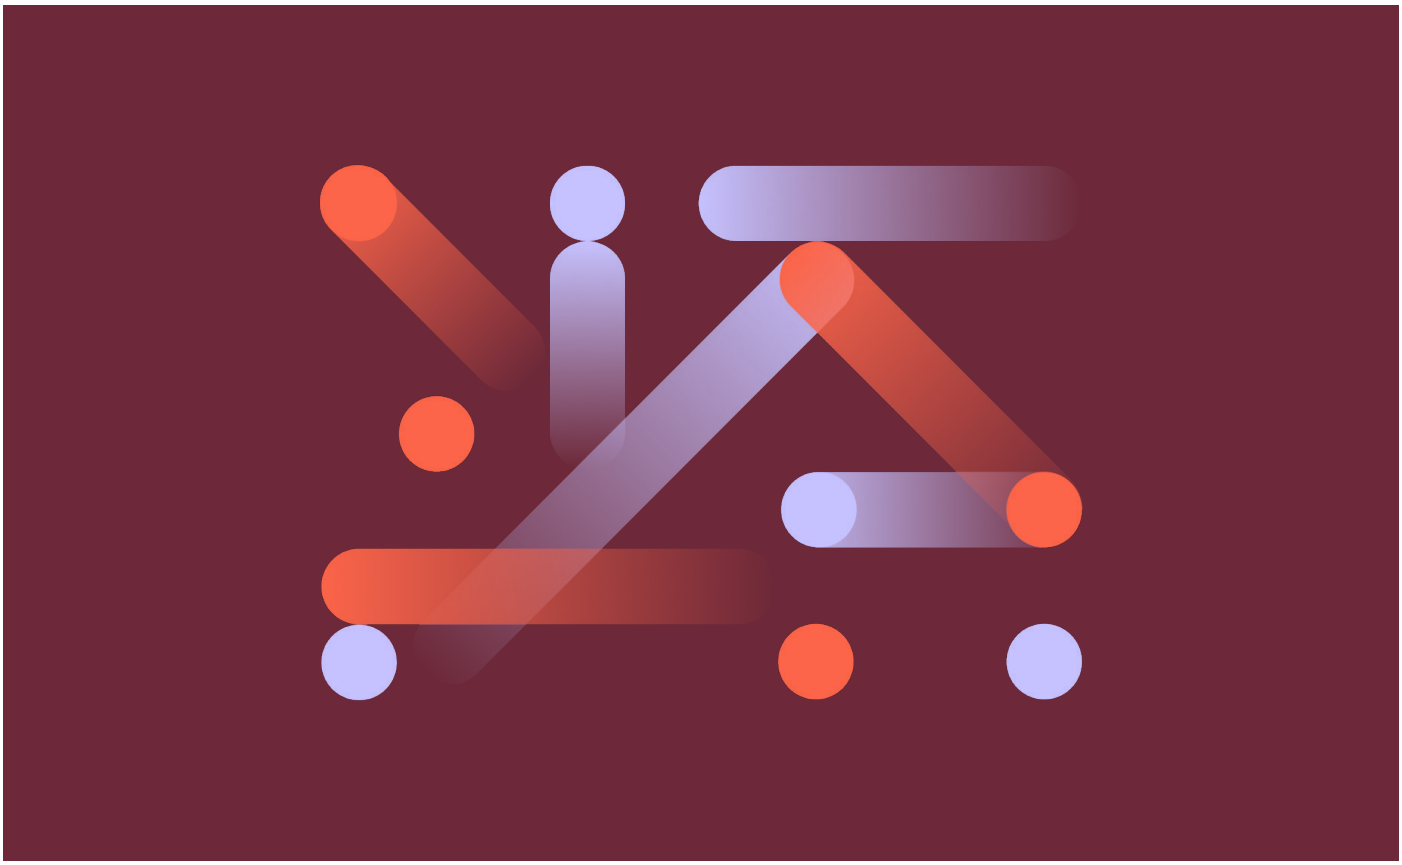

# Fully Autonomous Agents and Memory

This brings us to fully autonomous agents. Unlike semi-autonomous agents, which operate within structured boundaries and rely on prompts or human oversight, fully autonomous agents function independently once a task is initiated. They represent a significant leap in AI, capable of perceiving, planning, acting, and learning within complex environments with minimal external input.

While semi-autonomous agents seek input when encountering unknowns, fully autonomous agents are designed to adapt dynamically to new and unpredictable scenarios, operating independently without human intervention.

In short, while semi-autonomous agents seek input when encountering unknowns, fully autonomous agents are designed to adapt dynamically to new and unpredictable scenarios, operating independently without human intervention.

Autonomous agents achieve this by continuously learning from their environment, refining their approaches, and adapting their strategies based on past experiences. This ability to self-improve over time allows them to tackle increasingly complex tasks with greater efficiency. This adaptability makes them ideal for dynamic environments where conditions and requirements are constantly changing, enabling them to respond in real time and make informed decisions based on accumulated knowledge.

While memory plays a role across all AI agents, fully autonomous agents depend on it to retain and apply past experiences, enabling them to learn, adapt, and refine their decision-making over time. This memory-driven learning allows them to handle increasingly complex and evolving tasks, moving beyond simple rule-following to genuine strategic adaptation.

To push beyond individual learning, some AI agent systems leverage external knowledge sources to solve more sophisticated queries. This is where Retrieval-Augmented Generation (RAG) comes in — allowing agents to access real-time, external information for precise, contextually relevant responses.

Integrating RAG into AI workflows enables these agents to engage in more analytical, planning and reasoning as they engage with diverse data sources. Known as Agentic RAG, this approach provides agents with the capacity to assess, prioritise, and strategically leverage information — much like a skilled researcher.

Memory is essential for autonomous agents, allowing them to learn, adapt, and make informed decisions. By storing and applying knowledge, these agents evolve beyond simple task execution into adaptive systems.

Different memory types support this adaptability: working memory helps agents process current tasks and respond in real time; episodic memory allows agents to recall specific past interactions, refining their responses over time; and semantic memory provides a foundation of general knowledge, ensuring contextually relevant and consistent actions.

Together, these memory types empower autonomous agents to tackle complex tasks, collaborate effectively, and operate with minimal human guidance.

When memory and RAG are combined, agents not only retain their own experiences but also tap into vast external data sources, adapting fluidly within a networked system of intelligence.

But intelligence doesn't exist in isolation. Even the most advanced fully autonomous agents don't operate alone — they interact with other agents, systems, and data sources to execute tasks more efficiently.

This takes us to multi-agent systems (MAS) — where multiple decision-making entities work together in a shared environment to achieve their goals.

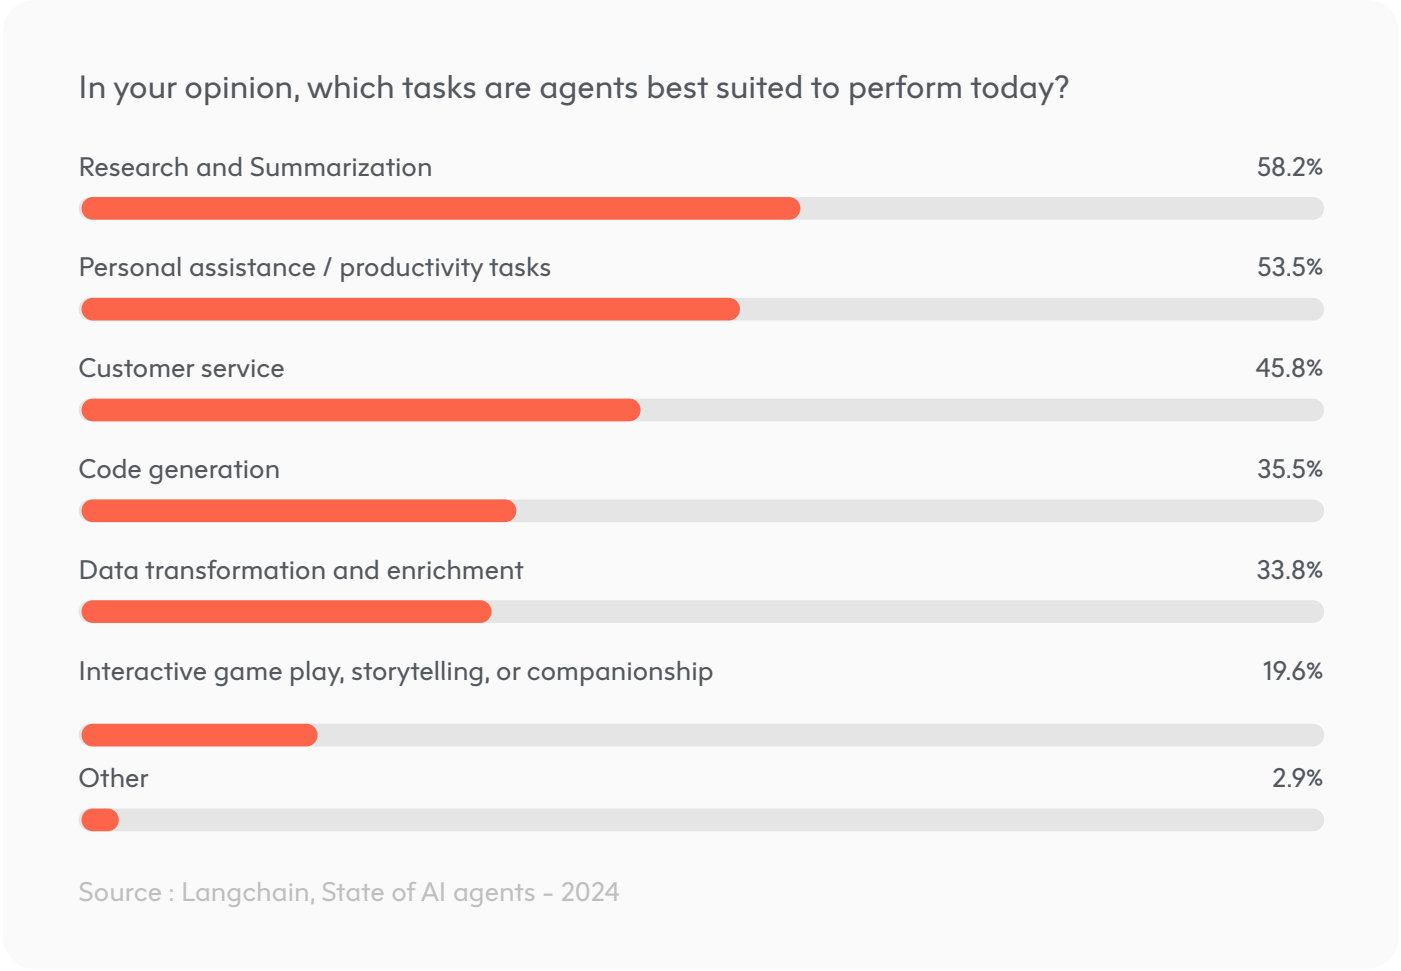

# What are Multi-Agent Systems?

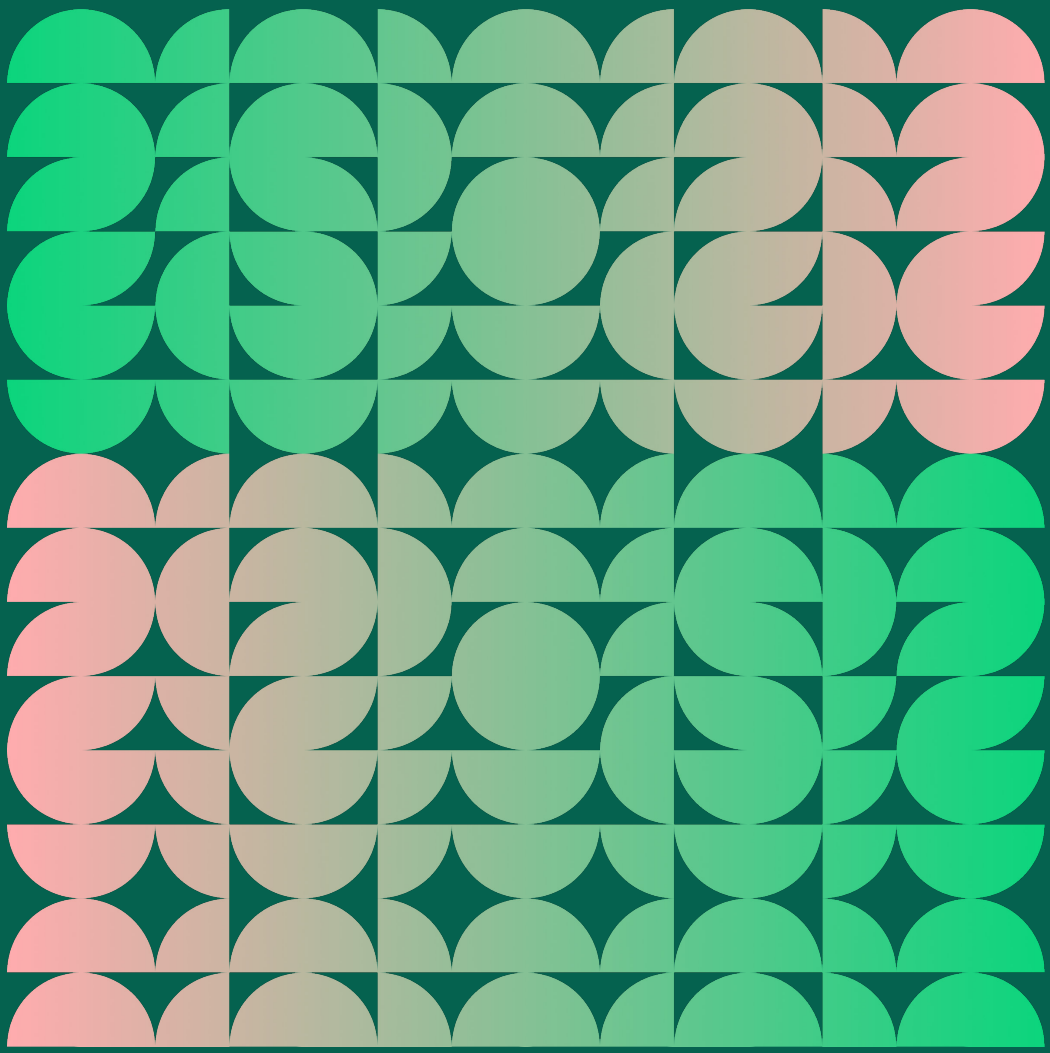

Multi-agent systems have been a foundational topic in AI research for decades, but in simple terms, a multi-agent system is made up of multiple decision-making entities — whether AI agents, robots, or even humans — working together in a shared e nvironment to achieve their goals.

While MAS enables AI-driven collaboration, it does not inherently imply autonomy or self-learning. Many MAS operate within predefined coordination structures rather than making independent decisions. However, as AI advances and new frameworks emerge, MAS is evolving toward more agentic systems — where agents not only collaborate but also reason, adapt, and act independently.

This evolution is driven by intelligent multi-agent frameworks such as LangGraph and ReAct, which enable agents to manage specialised tasks autonomously while coordinating seamlessly across functions to solve complex challenges.

The core challenges in multi-agent systems revolve around enabling agents to plan, learn, and collaborate effectively. This involves designing systems where agents can communicate, share information, and even adapt their behaviours to align with overarching objectives.

"As Generative AI adoption accelerates, multi-agent systems are shifting from theory to practice."

As we have seen, when talking about AI, at the heart of these systems are the AI agents themselves — the individual entities responsible for executing tasks, making decisions, and interacting autonomously (but not always) within their environment.

As Generative AI adoption continues to accelerate — Wharton recently found that there are now 73% weekly AI users in 2024 (up from 37% in the year prior) — multi-agent systems are shifting from theory to practice, making their widespread application inevitable. Driven by productivity gains, cost savings, and streamlined operations, these systems are poised to transform business operations, enhancing efficiency, fostering innovation, and enabling smarter decision-making across all functions.

Does your company currently have agents in production?

Yes 51%
No 49%

Are you currently developing an agent with plans to put it into production?

Yes 78%
No 22%

Source : Langchain, State of AI agents - 2024

# The Agentification of the Enterprise

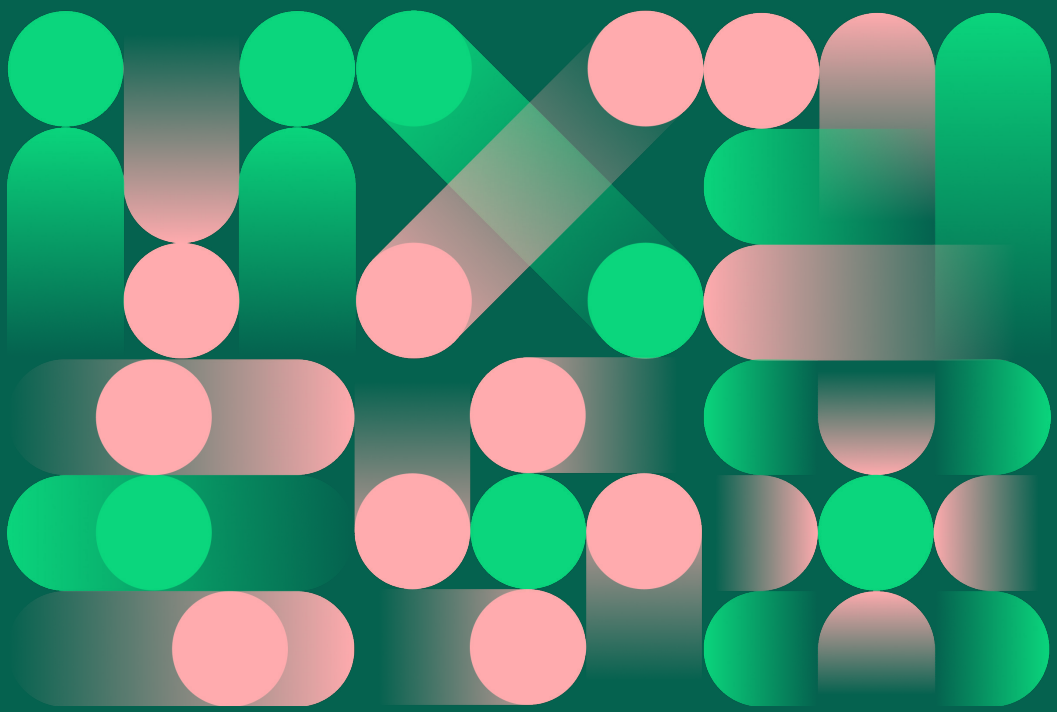

## Multi-agent systems (MAS) are central to agentification — the process of deploying artificial intelligence (AI) in the form of semi-autonomous or fully autonomous agents.

MAS provide the structural foundation that enables AI agents to interact, collaborate, and coordinate tasks at scale. As these systems advance, they incorporate greater autonomy, reasoning, and decision-making capabilities, allowing AI to move beyond static automation into adaptive, agentic workflows.

As we have seen, this transformation is driven by new multi-agent system architectures, enabling enterprises to build AI-driven ecosystems where agents operate autonomously (or semi-autonomously) while retaining the flexibility to collaborate dynamically.

Within such an agentic system, each entity possesses the ability to plan, reason, and use tools independently. While maintaining autonomy, agents engage in structured coordination to tackle complex, multi-layered, and multi-stepped challenges.

Effective inter-agent communication and distributed problem-solving are essential to making these systems responsive and adaptable. Through multi-agent reinforcement learning (MARL), agents share real-time information — such as sensor data, actions, and episodic memories of past interactions — streamlining knowledge transfer and minimising redundant learning. This coordinated exchange enhances adaptability, allowing agents to make informed, collaborative decisions that drive efficiencies across functions.

Research on MARL highlights the compounding benefits of this approach. Studies show that when autonomous agents interact and share information, they naturally develop cooperative behaviours, even in competitive environments. This continuous, collaborative learning effect allows each agent to build on the experiences and insights of others, creating a compounding knowledge effect. Over time, MAS become increasingly adept at tackling complex, real-world challenges, constantly adapting through interactions with their environment and each other.

The modular nature of these systems further enhances resilience and scalability. Organisations can update, expand, or introduce new agents without disrupting the broader system, ensuring that each component evolves independently within a stable, adaptable infrastructure.

But what does this look like in practice? What use cases and opportunities do multi-agent systems bring to the enterprise today? And what types of agents form this modular ecosystem?

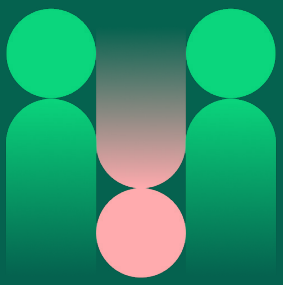

# Customer Service Agents

In the agentic enterprise, customer service agents act as strategic, conversational partners that enrich every phase of the customer experience. These agents, like Alaska Airlines' travel assistant or Best Buy's troubleshooting assistant, autonomously manage complex queries and adapt across channels, creating seamless, multi-platform engagement. By handling intricate interactions with minimal human oversight, customer agents elevate brand loyalty, increase customer satisfaction, and streamline operational efficiencies.

As we discussed in How to Design for Conversational AI, the implementation of safeguards like an "LLM-as-a-Judge" is crucial for creating reliable, brand-aligned and accurate AI experiences. This safeguard layer, a real-time LLM-based system, evaluates AI-generated responses for accuracy, relevance, and coherence, ensuring they meet stringent quality standards.

Ultimately, the agentic enterprise relies on such customer agents not only to reduce response times and operational costs but to fundamentally reshape customer relationships. Here, AI agents drive continuous engagement, learning from each interaction to enhance personalisation, anticipate needs, and ensure a unified experience that reflects brand values across every touchpoint.

# Assistant Agents

Assistant agents in an agentic enterprise operate as essential collaborators, streamlining workflows and bolstering productivity across departments. With AI-driven agents managing routine tasks — such as scheduling meetings and setting reminders — employees can focus on strategic projects, fostering a work environment geared towards innovation and efficiency.

One example is how Bell Canada's AI contact centre solutions support customer service teams by handling routine inquiries, and generating real-time conversation summaries, allowing live agents to dedicate full attention to customer needs.

These employee agents reshape team dynamics, not by replacing human input but by enhancing it.

# Engineer Agents

In March 2024, Cognition introduced Devin, dubbed 'the world's first fully autonomous AI software engineer', setting a new benchmark for code agents in development. Devin could reportedly streamline complex workflows, autonomously analyse code, troubleshoot errors, and optimise performance, giving developers the freedom to focus on creative and strategic tasks.

But Devin was not alone, across industries, companies are already harnessing similar agentic workflows and autonomous engineering agents to enhance productivity and code quality.

Wayfair has improved efficiency by using autonomous engineers to accelerate development setups and refine unit testing, while Best Buy's code summarisation tool reduces call times and enhances developer focus on high-value work.

By implementing agentic workflows and streamlining processes such as setup, debugging, and code analysis, code agents reduce the manual burden on developers, enabling them to shift their focus to more creative and high-impact projects. Perhaps more importantly, they allow many companies to bring coding capabilities in-house, meaning that they are better equipped in creating their own enterprise software and solutions from the ground up.

# Analyst Agents

RAG-powered analyst agents in an agentic enterprise serve as intelligent analysts, synthesising vast amounts of information to provide actionable insights. These agents operate across both internal databases and external data sources via methods such as RAG, delivering real-time answers and highlighting patterns that traditional methods might miss.

Companies like Bayer Crop Science harness data agents to transform agricultural data into precision-driven recommendations, boosting efficiency and sustainability in crop management.

Similarly, SURA Investments leverages AI to gain deeper insights into customer preferences, enhancing service personalisation and customer satisfaction. By bringing together comprehensive data processing capabilities, data agents make it possible for enterprises to act on insights swiftly, ensuring that decisions are not only data-driven but contextually relevant to evolving business landscapes.

# Security Agents

Autonomous security agents in an agentic enterprise act as omnipresent guardians, continuously monitoring and responding to potential threats. These agents go further than detecting anomalies; they analyse patterns, learn from past incidents, and adjust responses based on evolving risks.

Apex Fintech deploys security agents that cut down threat analysis and response times from hours to seconds, allowing the company to manage risks proactively instead of reactively. By automating these critical functions, security agents enhance organisational resilience, safeguarding data integrity, ensuring compliance, and allowing human teams to focus on strategy.

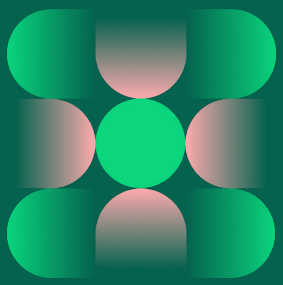

# Creative Agents

In an agentic enterprise, creative agents are instrumental in transforming content creation, personalising marketing, and enhancing customer engagement. Agents like those deployed by Puma generate localised product images and advertisements tailored to diverse markets, while Radisson Hotel Group uses AI to personalise ad campaigns, achieving faster production and higher engagement.

At the other end of the capabilities of creative agents, you have ideation, moodboarding, and concept development, where agents contribute to the early stages of creative work across various mediums (from image generation to 3D assets to video). By assisting in brainstorming and visual planning, creative agents provide inspiration and frameworks that can jumpstart the creative process.

Strategically, creative agents ensure brand consistency and adaptability across global markets by taking on repetitive creative tasks, allowing teams to focus on higher-value innovation.

# The Agentic Enterprise

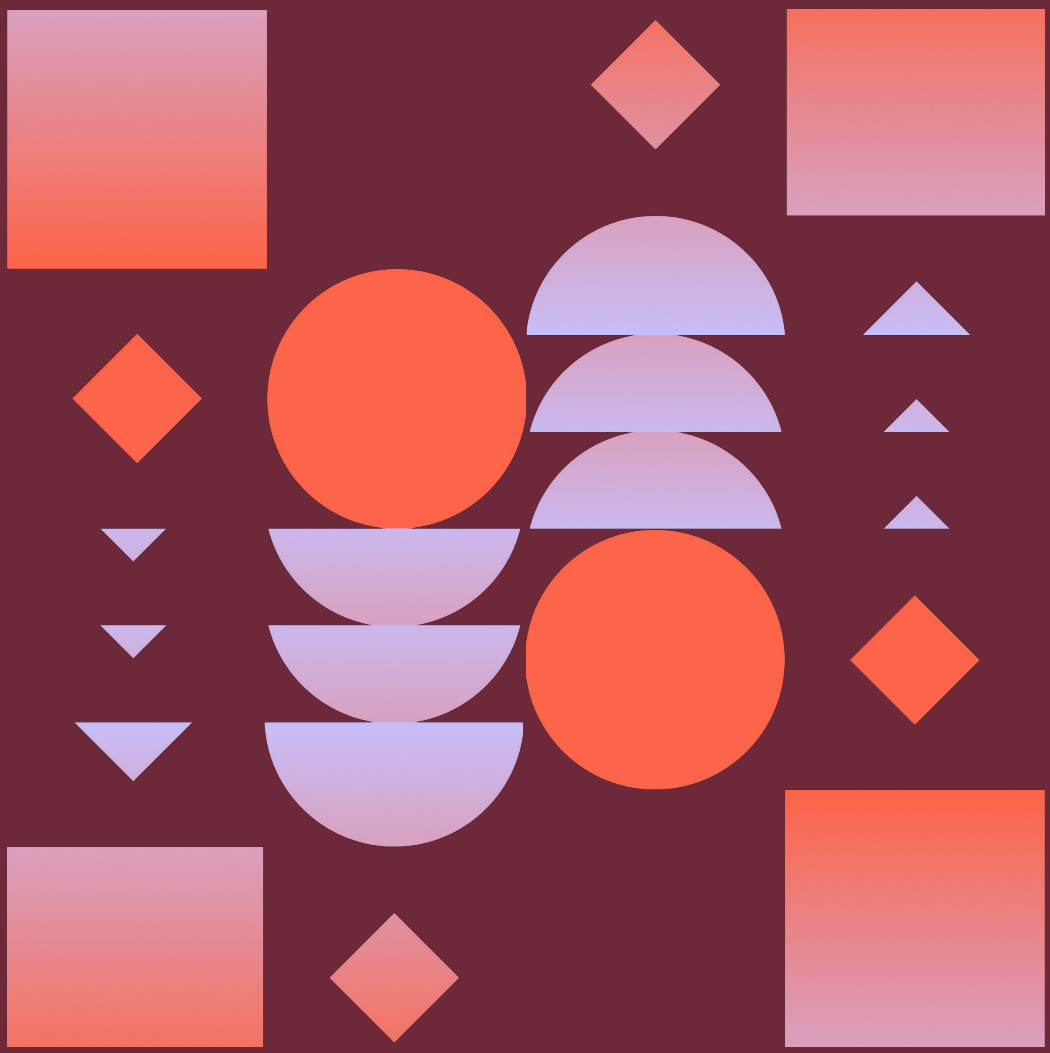

# At Elsewhen, we envision a future where every company is an AI-driven, modular agentic enterprise, with autonomous and semi-autonomous agents embedded across every function.

These modular agents — whether customer-facing, employee-supporting, code-optimising, data-analysing, or security-enforcing — work in a multi-agent system, each contributing specialised expertise while drawing from shared intelligence and memory to achieve seamless coordination.

The system remains adaptable, with human-in-the-loop oversight when needed to guide complex decision-making and ensure alignment with strategic goals. This shift from isolated operations to an interconnected, agent-powered infrastructure will empower businesses to personalise customer interactions, enhance productivity, and make agile, data-informed decisions across the enterprise.

This is a vision that represents AI not just as a tool but as an essential, strategic partner, fostering agentic workflows and enterprise resilience in every area of the modern organisation.

Get in touch today to learn how to transform your business into an agentic enterprise, where AI agents drive value, streamline operations, and shape the future of intelligent enterprise architecture.

Elsewhen

+44 (0)20 3758 3224
hello@elsewhen.com
elsewhen.com