IBM Technology Atlas

AI Roadmap

AI transforms how work gets done.

Realized
In progress

	2026	2028	2029+
The journey	*AI becomes software ("AI grows up").*	*Software becomes AI (the lines begin to blur).*	*The nature of work changes.*
What are we doing and how we are doing it	AI is driving a fundamental shift in software applications from consumption of standalone models as chatbots to multi-agent systems that can plan, act, and adapt in real time. However, in 2025, the agent landscape is the wild west. The process of building and deploying agents is hit or miss, undisciplined, and exposes businesses to new kinds of security vulnerabilities if agents are not built with the greatest of care. It is possible to create stellar agents, but it is neither easy, nor guaranteed, and it is expensive. Hugging Face CEO Clem Delangue once said "when it works, it's software." In 2026, generative AI will "grow up" and the process of building and deploying AI will become more disciplined, more orderly, safer, and (in a good way) a little more boring. Middleware and frameworks will emerge that make agent building less art and more science, resulting in more predictable outcomes. At IBM, we are pursuing an ambitious research agenda under the banner of "generative computing"—the idea that generative AI can be woven into software systems, and that deterministic traditional software and generative AI can not only coexist but can fundamentally complement each other's strengths and weaknesses. In generative computing, agents are just one example of generative programs, and new programming models will emerge to make usage of AI more predictable and secure. The capabilities of models will continue to improve, following a Moore's Law-like progression where every 9–12 months we see a 10-fold reduction in the size of a model required to achieve a certain level of capability. This will make AI increasingly ubiquitous. We will see a broader range of software quietly using AI, often in less visible but equally powerful ways, as AI becomes commonplace enough to suffuse software of all kinds. We expect all major IBM products to adopt agentic systems capable of reasoning and reaching into enterprise backend systems. For instance, in the context of data systems, agents will begin to autonomously tackle data systems design and operations tasks currently performed by engineers and analysts. For example, autonomously designing and validating novel data products based on demand, generating dataflow pipelines to discover, index, catalog, clean, and validate data for consumption by generative AI, and remediating dataflow issues as workload or resources evolve. Data systems will evolve to better address agentic data access patterns, which will be more iterative than current OLAP/OLTP queries, as the agents proactively assist users in refining their intent and the data to achieve their analytic goals. To address the cost overhead associated with data processing at enterprise scale, agents used at design time will produce efficient flows that minimize model and agent use at flow execution time. Entire ecosystems of AgentOps will emerge, focusing on monitoring, debugging, and orchestrating agents at scale. Enterprises have embraced a hybrid platform and infrastructure strategy as fundamental to their IT strategy. Hybrid AI systems will integrate into hybrid applications, platforms, and infrastructure as agents will increasingly integrate with existing ecosystems via tools. To guarantee operational control, AgentOps will be integrated in existing observability and operations strategies and solutions. Security concerns, including agent hijacking, prompt injection risks, credential theft, excessive permissions, tool manipulation, and insufficient monitoring and access controls will be mitigated via the use of ephemeral agent identities, just-in-time tokens, and delegation frameworks. Sandboxing, red/blue testing, and real-time policy enforcement will provide autonomous defense.	As AI becomes more woven into the fabric of software, the lines between software and AI will blur. Developers already use AI to create software today, but increasingly, AI will write software—including software that uses AI. AI will generate programs to solve problems better suited to traditional software, but it will also generate software to orchestrate its own operation, allowing it to overcome some of its own inherent weaknesses, accelerating progress. Application development will take new forms. Development might start with the specification of an agent, but that agent will produce code as it operates, leaving behind a trail of deterministic software as it charts a path to meeting the applications needs. As an application matures, the agent will only resurface when corner cases are goals, patching and hardening code to fill gaps. Human developers and AI will collaborate on the creation of applications, but the balance will shift over time until most human developers play an increasingly high-level supervisory role, along the lines of a product manager. Average workers will increasingly produce ephemeral automations as a routine part of their work, though it won't seem any less natural than using a spreadsheet. Legacy software will be tended by AI systems that first encircle these legacy software systems and eventually become a part of them (though never fully managing to eradicate them). Agent-to-agent communications will move beyond being novelties and will be become the de facto way that IT systems interact with one another. We have already seen AI perform at the level of a silver medalist in the International Mathematical Olympiad; address benchmark problems in math, science, and programming; help find conjectures and rule out counterexamples; check proofs; and power agents to improve pieces of algorithms by changing code. We will see this trend gain steam as these tools are further developed to be used to guide human intuition when tackling science and engineering problems. Security will become a sophisticated arms race, with new attack surfaces, new insider threats, and the equivalent of social engineering attacks, but aimed at AI systems, will become commonplace. The agent landscape will become a battlefield of increasingly active countermeasures, and some critical systems might begin to exclude the use of AI or even be partially air-gapped to protect them, forming a paradoxical retrograde trend. As multi-agent systems become prevalent in the enterprise, we will provide advances in agent-to-agent authentication and encrypted inter-agent protocols. Multi-agent observability will track issues of multi-agent collusion and coordination attacks, providing anomaly detection and consensus validation. Agents will continue to operate in existing enterprise ecosystems, hence the need for true hybrid operations and observability will remain. Agents will determine their own observability needs, based on feedback from the environment that they operate in, and auto-instrument themselves for operational visibility and action. Agents with pervasive observability and memory will lead to a much more dynamic and optimized environment for data which blurs the line between data and APIs. Agents will use this memory to learn the user's needs and will prepare the needed data in a just-in-time approach, minimizing effort expended. Further, the end user will be unaware of whether data is coming from a data repository or an API as agents autonomously determine the best way to provide users with needed information.	As AI subsumes IT, the nature of work and society will change, as it has in previous waves of technological change, from the invention of agriculture, to the mechanization of physical labor, to the rise of the internet. The mix of jobs undertaken in the economy will have begun a transformation by this point, with associated disruption of social orders and an evolution of public institutions, as in previous waves of technological change. AI will be an inextricable part of almost every occupation, just as the internet and telecommunications today are a part of practically every job, but human workers will remain as important as ever, despite futuristic predictions of singularity and the full replacement of human labor. Ubiquitous physical robots, perennially "just around the corner," will gradually become more common and cost effective beyond the niches they occupy today, though their development cycles will remain sluggish relative to the progress of AI. Significant fractions of business operations will be automated (or at least, automatable), though legacy systems will remain difficult to fully shake free of, and availability of energy and natural resources will serve as a brake on an otherwise brisk technological acceleration (barring a major energy breakthrough, which seems unlikely). AI tools will increasingly become a part of scientific discovery, accelerating progress in many fields, but that progress will disappoint those holding more exuberant techno-utopian visions. Science and invention will remain slow, methodical endeavors, albeit empowered with new tools, as in past waves of innovation. As physical AI or robotics emerge, become more prevalent, and interact with existing enterprise systems depending on their specialty, there will be platforms to observe and govern them, integrating across multiple physical AI vendors. As AI becomes increasingly capable of tackling tasks that previously resisted automation due to inherent complexities, it will open the door to fully autonomous data systems. That is, multi-agent systems will build entire, self-sufficient data stacks from scratch, ranging from data infrastructure design, to data discovery, integration, governance, and even proactively pinpointing relevant insights. Towards this goal, human interventions, both at design time and ongoing operations, will focus on clarification of goals, approval for state-changing actions, and receiving results.

Implications for:

- Automation
- Data
- Security
- Systems