Deloitte. Digital | salesforce

MuleSoft

OCTOBER 2025

The dawn of agentic AI

Orchestration, governance, and best practices on a new frontier

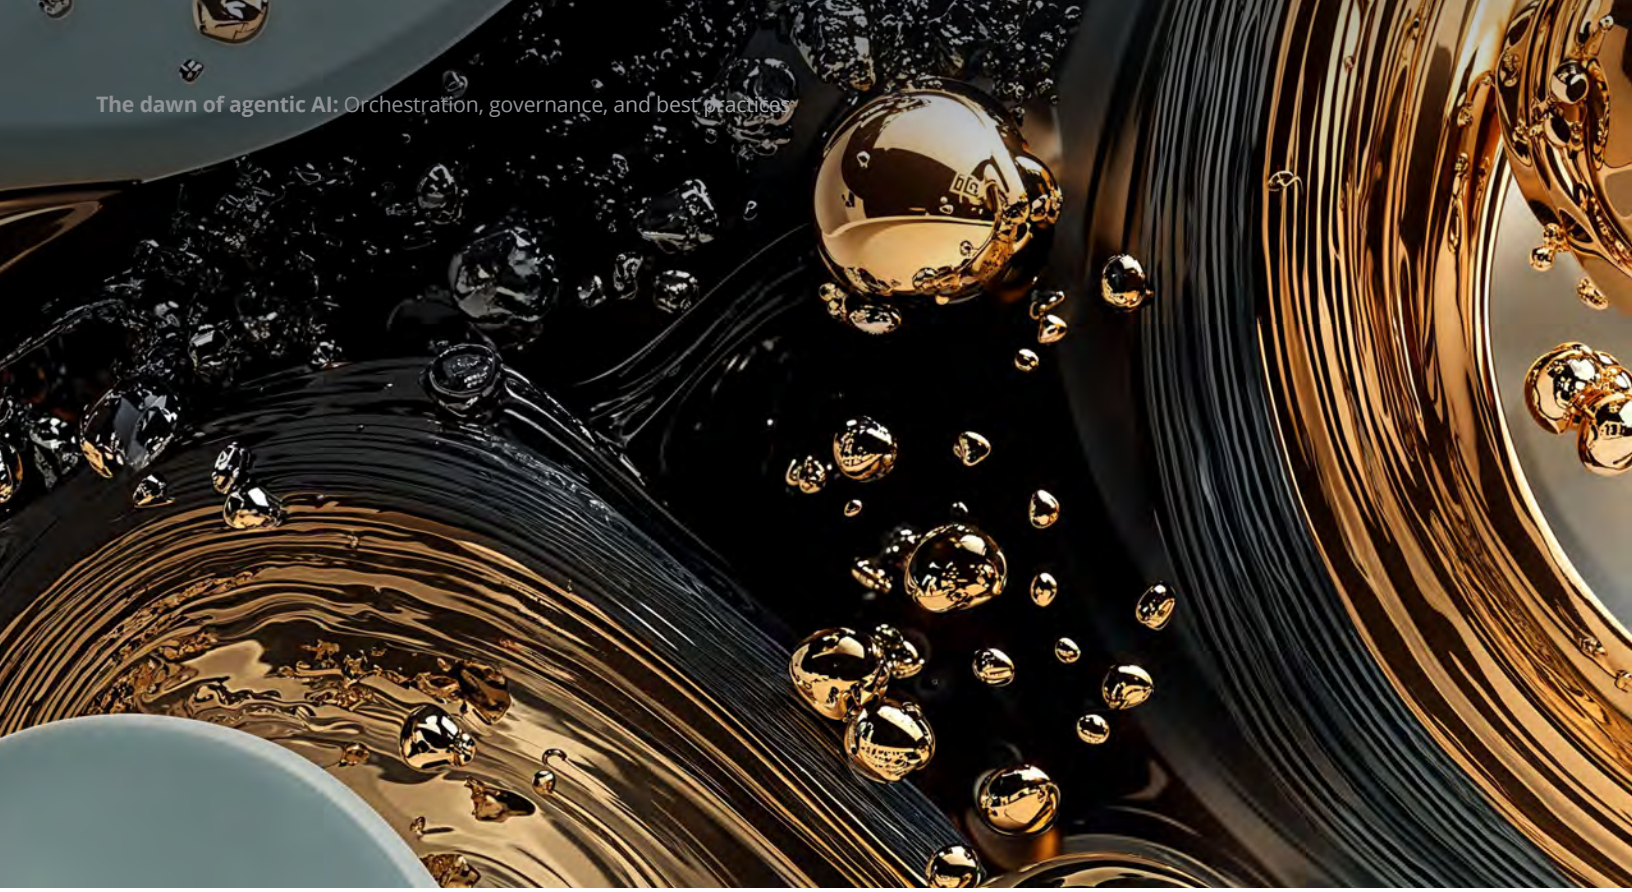

About the Deloitte AI Institute™

The Deloitte AI Institute helps organizations connect the different dimensions of a robust, highly dynamic and rapidly evolving AI ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the "Age of With".

The Deloitte AI Institute aims to promote a dialogue and development of artificial intelligence, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries, to explore key areas of artificial intelligence including risks, policies, ethics, future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, deliver impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you're in; whether you're a board member or a C-Suite leader driving strategy for your organization, or a hands on data scientist, bringing an AI strategy to life, the Deloitte AI Institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet ups and live events. Let's explore the future of AI together.

www.deloitte.com/us/AIInstitute

The sun is about to rise on the next horizon of digital labor: AI agents. As dawn arrives, we are taking our first clear look at how the combination of multiple technologies enables a degree of machine autonomy that has long been desired but just out of reach.

Agents and agentic systems set tasks, make decisions, and act autonomously to achieve defined goals. They are reasoning engines that can understand context, plan workflows, connect to external tools and data, and act with limited or no human intervention.

Many enterprises that already contend with myriad applications are now also experiencing a proliferation of agents, which introduces new opportunities as well as new challenges. There are lessons from prior technology revolutions to weigh, particularly as it relates to agentic AI orchestration and governance. For all the investment and planning around deploying agents, there needs to be commensurate consideration for how the organization will orchestrate and govern agents as scale and complexity increase.

Zooming in on an exponential curve of agent adoption, the start of the slope appears flat. In enterprises today, there are typically a limited number of agents performing a modest set of tasks, often as retrieval agents, where a human poses questions and interrogates data. At some point, however, the exponential slope will rise, and the number of agents, their function, and their scale will grow. When that happens, AI governance and orchestration become paramount and human decision-making and judgment increase in value.

> It is time to begin planning and defining how the enterprise will unify and coordinate its entire workforce and IT estate to enable the secure and controlled use of agents.

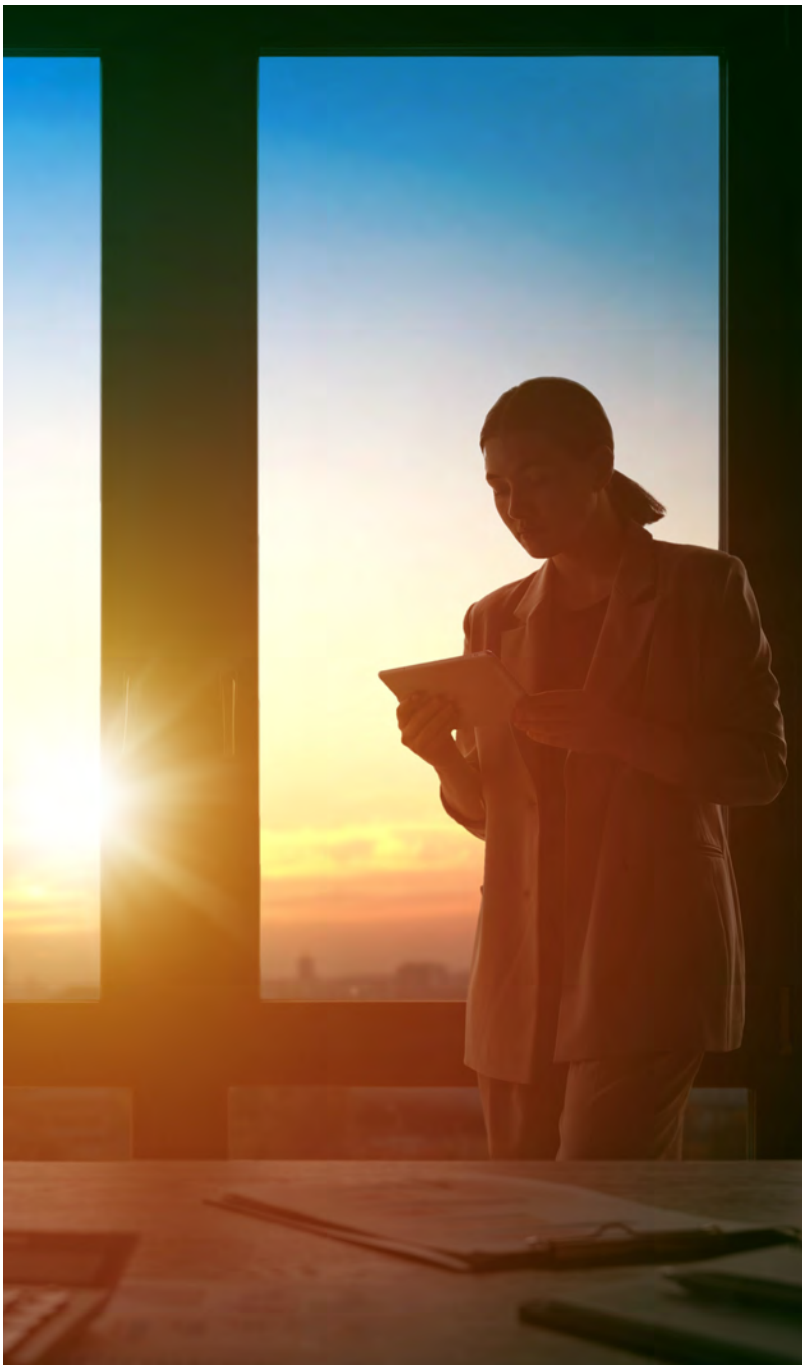

Orchestrating and governing experimentation

In some ways, the agentic dawn is similar to the advent of cloud computing. Widespread connectivity and third-party data centers opened the door to transformational data modernization initiatives.

Many enterprises leaned into the opportunity, allowing developers to experiment to learn how best to leverage the cloud. Some organizations moved forward without much concern for orchestrating and governing activity, such as tagging, controls, and decisions on who could launch or retire experiments. The result in some cases was much higher than expected costs.

Some enterprises discovered a similar challenge in building services with application programming interfaces (APIs). At many enterprises, initial experimentation ballooned into dozens, even hundreds of APIs, and IT management had limited visibility into how many existed, their function, and who created them. It was a matter of experimentation and technology development with only a passing glance at AI orchestration and governance.

The introduction of agents will again raise the specter of excited experimentation without appropriate orchestration. We can already see complexity growing. According to a Salesforce survey,[1] **organizations using AI agents have, on average, 45% more applications than those that do not (1,103 vs. 762).** In addition, enterprises deploying agents use an average of 22 AI models, while those not using agents leverage 15 models.

If agentic AI orchestration and governance are not addressed now, the cost to do so in the future will likely compound, which could slow adoption at a time when competitors are accelerating. A more effective approach is to design, from the beginning, the architecture the organization will need for its agentic ambition. This will likely include considerations around integration, orchestration, and the management and governance of a significant number of agents.

The value of human judgment

Enterprise leaders should be considering how controls and a human in the loop will be maintained as agent scale and complexity increase. In some ways, human governance is inherent in agent use. It is humans who select which agents to use and how to grant or restrict access to the data that fuels their computations. The human employee adjudicates outputs not just for accuracy but also for the more subjective matter of taste.

Imagine an agentic AI system that orchestrates the product design lifecycle in the consumer industry. A market-sensing agent analyzes real-time data from trend reports, social media, and consumer sentiment to reveal product opportunities. A concept agent uses this information to generate new product designs, while a feasibility agent evaluates designs against factors such as sourcing options, cost, and manufacturing timelines. The result of this agent activity is 20 possible product designs that meet criteria for feasibility and market demand—but which should head to production? There is no universal truth for which of the 20 options is best. It's a matter of taste. The human user decides which designs move forward to development, based on a deep understanding of their customers and the marketplace.

Human creativity and taste will become more valuable as agents proliferate.[2] Almost 30 years ago, AI became virtually unbeatable in chess, and yet, chess continues to grow in popularity. Simply doing something hard to be the best is no longer the core attraction. It is about uniquely human factors that we still value: the grit of competition, the joy of sport, and the stories and personalities of the players.

> With agents, human ingenuity and decision-making are necessary, as is adoption and a change in behavior to embrace new ways of working. Human employees need to hear the value narrative of how changing work can benefit them, and an important part of that narrative is that their judgment is becoming more valuable.

Heading toward the horizon

Among business leaders, one of the most common questions is, what are the emerging leading practices other organizations have discovered in their agentic journey? For better or worse, the leading practices playbooks are still being discovered.

Even with some uncertainty, there is another important question leaders should be asking: Were we surprised by the result from our agent deployment? **If the agent does precisely what was expected and nothing more, the organization may not be pushing the boundaries of the technology's sophistication.** They may instead just be automating a task that could probably be automated with mature AI applications.

Instead, look for surprises and moments of revelation. When Generative AI chatbots emerged several years ago, there was widespread surprise among users. The coherence, quality, and human language imitation was stunning because nothing like it had existed before. It was a eureka moment. That same degree of surprise and enticement is what organizations should be striving for with their agentic vision. And to be sure, surprise is a human reaction, underscoring again the importance of human taste and discretion.

With forethought to orchestrating agents, governing their use, and championing human judgment, organizations can confidently move forward onto a technology frontier that is still relatively untread but growing more illuminated by the day.

Authors

Beena Ammanath
Executive Director
Deloitte Global
AI Institute
bammanath@deloitte.com

Andrew Comstock
Senior Vice President
and General Manager
Mulesoft at Salesforce
acomstock@salesforce.com

Get in touch

Karim Trojette
Global MuleSoft
Alliance Leader
Deloitte Consulting GmbH
ktrojette@deloitte.de

Kurt Anderson
Managing Director, API
Transformation Leader
Deloitte Consulting LLP
kuanderson@deloitte.com

Karuna Mukherjea
Chief Marketing Officer
Mulesoft at Salesforce
kmukherjea@salesforce.com

Endnotes

1. Salesforce, "Integration key as 93% of IT leaders turn to AI agents amid soaring resource demands—new research," January 29, 2025.
2. TEDxMiami, "The value of human creativity in the AI era," video featuring Beena Ammanath, 9:41, May 15, 2024.

Deloitte. Digital