



Futures

ISSUE 4 | FALL 2025

The Agentic Wave

Agentic

Robotics

Futures Practice



Welcome

- 3 Making Sense of The Agentic Wave

Agentic

- 5 In the Agentic Enterprise Era, How Much of Your Business Should Run on Its Own?
- 9 The Kindergarten Test: How AI Will Free Workers from Drudgery
- 25 How Enterprise General Intelligence Will Form a New Business Imperative
- 30 Revisiting Personal AI Futures
- 37 Meet the Digital Workforce Scaling Up Sustainability and Climate Disaster Response

Robotics

- 15 Three Forces Driving the Future of Enterprise Robotics
- 54 The Robotics Breakout Moment by Salesforce Ventures
- 45 Speculative Futures: Imagining the Future with Everyday Robotics

Futures Practice

- 23 Exposing leaders to AI Futures
- 43 Tailoring Scenarios for Specialist Functions
- 51 Museum Leaders Navigating AI Futures

Salesforce Futures



Peter Schwartz
Chief Futurist,
Advisor to the CEO



Mick Costigan



David Berthy



Daniel Lim



Marc Escobosa



Angela Gleason

Contributing Salesforce Authors



Silvio Savarese
EVP, Chief Scientist,
Salesforce AI Research



Max Kirby
Distinguished
AI Architect



Melissa Forbes
Salesforce Newsroom



Emily Zhao
Salesforce Ventures



Pascha Hao
Salesforce Ventures



Making sense of The Agentic Wave

Eighteen months ago, Salesforce Futures debuted Futures with our issue on Personal AI Agents. Subsequent issues looked at *Agents at Work* and *Inventing the Future with Agents*.

The magazine began in response to a gap we saw in the conversation around AI: too much hype, too much doom, and not enough clear-eyed imagination and practical analysis of the implications of what's next for business.

When we began exploring agents, it felt early. Since the launch of Agentforce at Dreamforce just over a year ago, things have changed dramatically. Today, agents — elegantly defined by Simon Willison as 'a system that uses tools in a loop to pursue a goal' — are everywhere. And what Salesforce Chief AI Scientist Silvio Savarese calls a series of "boring breakthroughs" is making powerful, value-delivering AI available to everyone.

The wisps of inspiration we started exploring in 2024 have now built into a towering Agentic Wave that is crashing down on businesses everywhere, creating new opportunities and new questions. We've framed this issue of Futures around that theme and three key areas of focus: agentics, robotics, and the practice of doing futures work in this space. While we've shifted to a single annual issue of Futures, we've continued publishing on the Salesforce Newsroom. This issue brings together the highlights from that work, along with new pieces and contributions from several brilliant Salesforce colleagues. Together, they capture the ideas that have most shaped our thinking about the agentic enterprise.



Agentic

When everyone has agents, where will real, durable competitive advantage come from? The winners won't be those with the most automation, but those with the best learning loops between people and machines, the most adaptable autonomic core, and the deepest customer-centricity. Leaders who keep learning won't just drive productivity — they'll rewire their businesses into agentic enterprises that build lasting relationships.

We can't navigate new waters with old maps. As Peter Schwartz says, "If you get your map wrong, you do the wrong thing." We need new frameworks, stories, and mental models to make sense of the terrain. Think of this issue as a sketchbook of emerging contours — a starting point for the maps we'll need in a world of agents.

We introduce the Autonomic Core, a model for reimagining your business around the balance between automated and intentional functions.

This frame, drawn from the human body, offers a fresh way to think about how your organization's architecture might evolve in the agentic era.

Distinguished AI Architect Max Kirby explores how agentic systems can amplify human ingenuity. Silvio Savarese outlines our AI Research team's concept of Enterprise General Intelligence and its Capability-Consistency Matrix, a useful framework for evaluating AI's utility in the enterprise. And Melissa Forbes offers a provocative look at how agents are empowering nonprofits to scale their impact in sustainability and disaster response.

Finally, we revisit Personal AI Futures, grading our predictions from the first issue of Futures and updating our perspective on what's next.

Robotics

Because this is *Futures*, we also look beyond the crest of the current agentic wave — toward what happens as agents leap from the digital world into the physical. Included is our piece on The Three Driving Forces of Enterprise Robotics and Speculative Futures, which push the horizon even further. We also excerpt from a piece by Salesforce Ventures' Emily Zhao and Pascha Hao, unpacking The Robotics Breakout Moment.

Futures Practice

The Agentic Wave isn't just changing technology — it's changing how we lead. In our Futures Practice sections, you'll find practical examples of how Salesforce leaders apply futures thinking to turn disruption into opportunity, inspire experimentation, and use foresight as a lever for strategic action. The future doesn't arrive all at once. It appears in small, transformative moments: a nurse spending five more minutes with a patient; a farmer protecting a crop with real-time data; a robot helping someone up the stairs. These aren't distant possibilities; they're early signals of a world already in motion. The ideas in these pages aren't final answers; they're invitations to build agentic enterprises that learn continuously, adapt constantly, and keep humans at the center. The Agentic Wave won't wait for perfect strategies. Before you can surf, you have to swim. It requires the courage to experiment, the discipline to learn, and the judgment to balance agents and humans in your own unique way. After nearly two years of mapping AI futures, it's time to begin — even as we keep drawing the map.


The Salesforce Futures team



Marc Escobosa
VP, Salesforce Futures

September 2025

In the Agentic Enterprise Era How Much of Your Business Should Run on Its Own?



Many executives harbor a seductive vision of AI automation: a sleek, ultra-lean organization that runs with minimal human oversight. The promise of lower costs, faster operations, and round-the-clock productivity makes it an alluring prospect.



Who wouldn't want (at least in theory) a business that hums along on its own while leadership focuses on the big picture?

Yet that's unlikely to be how successful organizations will work in the age of agentic AI. The companies that thrive won't succeed by automating everything. They'll win by distinguishing between tasks that should be entrusted to AI and those that require human judgment.

Consider how your body handles this balance. Your heart doesn't need a memo to beat 100,000 times a day. Your digestive system processes nutrients without performance reviews. Your lungs keep you breathing even when you're asleep.

These systems are "autonomic": involuntary, automatic, and self-regulating. Most of the time they run quietly in the background — until something unusual calls attention to them. In business terms, the same is true: Predictable, routine processes can run on autopilot, only rising to leadership's attention when they stray outside the norm.

Meanwhile, you weigh every word in a high-stakes client conversation and carefully consider whether to buy capabilities through acquisition or build them in-house — even as you increasingly lean on AI to help guide your thinking. Humans augmented by AI tools — not an AI agent acting alone — make these strategic decisions. This division between the automatic and the intentional isn't just elegant biology — it's a useful way to think about AI in business. Call it your "autonomic core": the set of functions that run quietly in the background so people can focus attention where it matters most.

Why agentic AI changes everything

We're entering the era of agentic AI — systems that can act independently to complete complex tasks. What sets these agents apart is their ability to handle variance.

They can process requests in many languages, decipher prompts riddled with spelling errors, recognize intents phrased six different ways, and adapt midstream when circumstances change, showing a resilience that earlier systems lacked.

Before generative AI, automation was brittle. Rules-based systems demanded coding skills, which turned IT into an operational bottleneck. Every exception — every edge case — meant another trouble ticket in the backlog that required a developer to fix while customers dealt with the consequences of our incomplete imaginations.

Once set up, agentic AI doesn't require coding skills. It doesn't care about departmental silos. Give it a goal — to maximize customer lifetime value, for instance — and it will optimize for that outcome without getting bogged down in whether each incremental dollar should come from sales or marketing. The agent doesn't have an ego or play office politics. It doesn't let ambition cloud its judgment.


What belongs in your autonomic core

Of course, not every business function can — or should — be fully autonomous, and every company will have its own appetite for how much to make autonomic. Payroll, onboarding flows, routine data processing, and data pipelines can be heavily automated. Standard customer service can mostly handle itself. Agentic AI shines when processing massive quantities of data, optimizing for outcomes, and coordinating across functions and departments to handle routine tasks at scale.

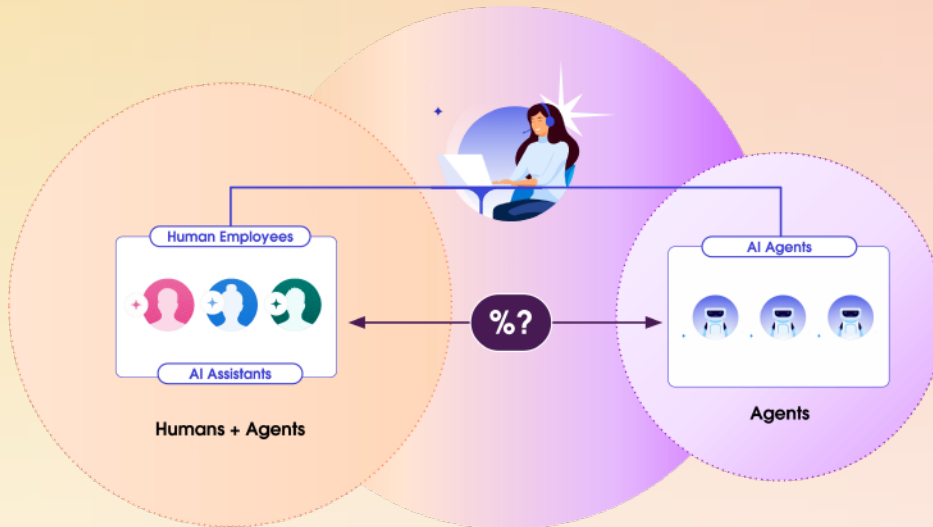
But VIP service requires human judgment. Strategic decisions about new markets demand conscious intention. High-stakes creative choices still depend on human taste and values.

The key is that humans aided by AI excel at different things than agents working alone. Humans with agents are better at navigating ambiguity, expressing values, integrating cross-domain knowledge, and taking responsibility when things go wrong.

In other words, this isn't mindless automation — it's intelligent automation that knows when to escalate. Your autonomic core operates in the background until it recognizes something unusual: "I've noticed a sudden change in price elasticity for this product among under-30 buyers who shop primarily on a mobile app."



Consider how your body handles this balance. Your heart doesn't need a memo to beat 100,000 times a day. Your digestive system processes nutrients without performance reviews. Your lungs keep you breathing even when you're asleep.



At that point, it escalates to what we call the human-plus-AI layer — cross-functional teams of smart problem solvers who can diagnose and respond, which could have profound implications for the shape of future workforces.

Set thresholds like a physician tracks vitals

Think of this like physicians monitoring vital signs. Doctors don't want an alert every time someone's heart rate increases by two beats per minute or if their blood pressure remains within the acceptable range. Instead, they watch for meaningful patterns based on historical data — and use blood pressure cuffs, blood tests, and other diagnostic tools to catch problems early.

Your business should do the same. You don't want alerts if there's a small backlog in a warehouse that clears itself by morning. But you do want early warning when trends hint at something serious, like a supplier consistently missing shipments.

Success will mean setting clear thresholds: defining normal variance versus genuine exceptions. Done well, you build an early warning system that delivers the right information at the right time with enough lead time to head off trouble.

The competitive edge

Here's the counterintuitive truth: The companies that will win won't be those with the most automation. They'll be the ones with the best learning loops coming out of their autonomic core.

That means three things: building robust cores that handle routine processes with minimal oversight, getting the balance right between what stays in the background and what gets surfaced, and creating learning loops so every escalation strengthens the system.

The real value comes at the edges — when judgment, creativity, and cross-domain thinking are needed most. The question isn't "How much can we automate?" but "What deserves our conscious attention?" Push as much as possible into the core so your people can focus on the work only humans — working with AI — can do well. Companies that strike this balance won't just pull ahead today; they'll be positioned to sustain an advantage for years to come.


April 2025

The Kindergarten Test

How AI Will Free Workers from Drudgery



MAX KIRBY
Distinguished AI Architect
Salesforce



Forget multiple screens, keyboards, and mouse pads.

Imagine a workday where you simply instruct your computer on the tasks you want done. Max Kirby, Principal AI Architect at Salesforce, believes this future is closer than you think.

Kirby recently discussed his predictions for the agentic AI-powered workplace of 2030 with Marc Escobosa, VP of Salesforce Futures. He envisions a world with fewer screens, dedicated AI assistants, and instant access to the information you need.

Q. To talk about the future of agents, let's start with their beginnings. When were agents first on the map?

John McCarthy and Allen Newell discussed 'intelligent agents' as early as the 1950s.

The International Journal of Agent Oriented Engineering has been publishing information on agents for the past two decades.

But for many, it feels like agentic AI is moving quickly. That's because technology is typically developed first, with practical applications — or what we do with it — figured out later.

What you're feeling right now is the combined pull of decades of computer science thinking that's already been done. All that groundwork means autonomous AI solutions like Salesforce's Agentforce can be applied immediately to practical, real-world use cases — something that can feel like magic.

All that groundwork means autonomous AI solutions like Salesforce's Agentforce can be applied immediately to practical, real-world use cases — something that can feel like magic.

For example, with Agentforce, Precina Health is automating patient check-ins, unifying clinical data from EMRs and lab reports, and onboarding practitioners with coaching in rapport-building and clinical guidance — allowing clinicians to focus on high-impact care while scaling personalized diabetes support nationwide, saving thousands of dollars per year in reduced administrative overhead and clinical training costs.



Imagine a future where agents automate this data entry for nurses on their behalf.

Q. Imagine it's the year 2030. How do you think agents will alter our day-to-day work life?


Agents lessen the need for point-and-click interfaces. You give agents a task and they do it for you. Yet, for two and a half decades, UX design has focused on the human operation of the computer. Now, practitioners with coaching in rapport-building and clinical guidance — allowing clinicians to focus on high-impact care while scaling personalized diabetes support nationwide, saving thousands of dollars per year in reduced administrative overhead and clinical training costs.

Q. How will agents impact the jobs we're doing?

A lot of today's jobs don't pass what I call the kindergarten test, which is, 'Did your kindergarten self imagine that you would be doing this as a job?' Take data entry — no one imagined in kindergarten that they'd be entering numbers into a computer all day for a career.

Agents will change that — taking on the manual, administrative work so humans can focus on more creative, strategic parts of their job. In this way, AI will make us more human.

For example, we all know that CRMs are more effective with data added in. Yet, salespeople typically take three to seven days to enter data into their Salesforce instance after a customer interaction.



Q. By 2030, how will agentic AI change day-to-day interactions?

When agents input that data instantaneously, without human help, reps can focus on more important work like relationship building or creatively solving a customer's pain point with their products. Plus, the organization benefits from faster knowledge input. That kind of shift to strategic work is going to help all of us, and is why I'm a wild optimist about agentic AI.

By 2030, consumers will have personal AI agents acting on our behalf, handling everything from planning our schedules to handling our shopping and managing our finances. And businesses will use agents to talk to their agents, the proxies of the customer. Imagine, for example, a consumer's personal assistant agent dialing into a retailer's customer service center to initiate a return.


This means that instead of designing personalization to cater directly to a consumer's tastes, businesses will need to design personalization to align with the preferences and objectives of the consumer's agent, or with the consumer's agent's interpretation of its human's preferences and objectives.

How agents represent our personal interests will be drastically different, too. We send lawyers to go to law school to learn law. We send med students to medical school to learn to become a doctor. But there's a limit to what they can store in their brains. Agents, on the other hand, have nearly unlimited space to store information. I'm not saying that agents will practice law or medicine in the exact same way as a human, and certainly won't be as empathetic, but they'll absolutely have that level of knowledge. Being able to pull that information out immediately will amplify us humans 10 to 100x.

Q. What's an example of how an agent will better represent us in 2030?

Let's consider how agents will revolutionize the healthcare industry, specifically for nurses. Nurses are the backbone of our healthcare system, yet they are overworked and under immense pressure. In fact, one in four nurses is very likely to leave their current role.

At the same time, nurses spend a significant portion of their time — between 12 and 25% — on documentation, primarily entering data into systems. This is time spent away from direct patient care, their core responsibility and where their expertise is most needed.



Ultimately, digital twins, enhanced by AI's ability to process diverse data types, offer a powerful tool for enhancing understanding and decision-making.



Imagine a future where agents automate this data entry for nurses on their behalf. Instead of manually inputting patient information, nurses can use voice commands or other natural language interactions to have agents record and update patient charts, medication schedules, and other essential documentation — approving and curating it instead of entering it manually. Beyond just data entry, agents will also be invaluable in providing nurses and healthcare professionals with immediate access to critical information. Need a patient's medical history, the latest research on a treatment, or information on drug interactions? An agent can retrieve it quickly, allowing nurses to make informed decisions quickly and efficiently.

Additionally, we expect patient advocate agents to enter the market in the next five years. These agents will represent your interests as a patient, providing you with medical knowledge that normally requires years of schooling. Imagine having a friend who went to medical school by your side helping you through your care. Nice to have, but not everyone can access that kind of advantage. With personal agents acting as advocates, that will change.

This is more than automation of menial work — agents will level the playing field for us, making equalities in critical outcomes less impacted by asymmetries in knowledge.

Q. What are digital twins, and how will agentic AI change the way these work?

A digital twin is a digital model of a real-life customer, product, or even business process.

They allow us to replicate real-life scenarios to predict, test, and optimize strategies.

Previously, building digital twins relied on structured data like fields and values, and required significant human involvement or complex software to set up and analyze the results. Now, with advancements in agentic AI, and systems that can blend structured and unstructured data together at petabyte scale, like Salesforce's Data Cloud (now, Data 360), agents can interpret the semantic meaning within unstructured data — like customer conversations — to add rich context to digital twins. This allows businesses to understand customers in their own words and predict their needs.

Bumble recently proposed using digital twins to predict two people's compatibility. This same concept, backed by unified data and agentic AI, can be applied to business scenarios like testing ad campaigns or pricing strategies. For example, a marketing team could use a digital twin of their target customer to simulate how they would respond to different ad variations, predicting which one would yield the highest engagement before launching the campaign.

Ultimately, digital twins, enhanced by AI's ability to process diverse data types, offer a powerful tool for enhancing understanding and decision-making. I expect we'll see more of this in the coming years as businesses work to predict what their customers need.





David Berthy
Salesforce Futures



Marc Escobosa
Salesforce Futures

Three Forces Driving the Future of Enterprise Robotics

June 2025

If your social feeds are like ours, they're probably filled with robots of every shape and size performing feats that would have been unimaginable a few years ago:

knocking off chores with untiring efficiency, dancing with influencer IShowSpeed in China, running half-marathons, lifting heavy objects, sparring with black belts, and so forth.



These demos are incredible, but their value goes far beyond tricks and novelties. The real value will arrive when robots extend digital agency into the physical world. As sectors like healthcare and construction face persistent labor shortages, robots have the potential to augment human capabilities and support critical tasks in these and other fields. Consider the numerous roles where increased support could be invaluable: nurses, teachers, firefighters, disaster response personnel, and paramedics, to name a few. Why wouldn't we want a world with fewer shortages in these vital roles, allowing human professionals to focus on the most complex and human-centric aspects of their work?



AI perception advanced very quickly once we realized how fast deep learning would evolve... Deep learning and end-to-end training are arriving for robots.

Juan Carlos Niebles

Salesforce Solutions Architect

Think about robots that respond to signals from sensors and other sources in high-risk fire areas and coordinate with drones to prevent property damage, or construction robots that augment human crews to build cheaper, better housing in areas where persistent shortages impede broader access.

Robots might play a growing role in elder care, to help ensure faster service in retail and hospitality during peak periods, and help maintain both public and private spaces with significantly greater scale than currently possible (“Hey assistant, could you please have them send one of the roadbots to fix that pothole on Dahlia and 14th Avenue?”).

While the potential for robots to augment human capabilities and support critical tasks is immense, realizing this future will require careful consideration of crucial factors, like physical safety measures, the need for predictable and reliable operation, the nuances of psychological human-robot interaction, and the broader ethical and data security implications of increasingly autonomous machines.

Salesforce Futures has been engaging with robotics experts inside and outside Salesforce to learn more about what we can expect in the years ahead. These conversations have revealed three technological and strategic forces that business leaders should consider now to position their organizations for the incoming robotics wave.



1 Convergence of digital and physical intelligence

Breakthrough improvements in AI are helping machines make sense of and interact with the world around them. In the past, robots struggled to make sense of their environment and perform complex tasks because they couldn't perceive and interpret information effectively. Now, advancements in deep learning are accelerating robots' ability to see, understand, and respond to their surroundings. As Salesforce researcher and robotics PhD Juan Carlos Niebles told us in a recent Salesforce Futures Roundtable, AI's perceptive capabilities have rapidly evolved, bringing us closer to the long-sought goal of creating robots with more generalized physical intelligence. "AI perception advanced very quickly once we realized how fast deep learning would evolve," he said. "Deep learning and end-to-end training are arriving for robots."

In our conversations with Niebles, he points to new Vision Language Action models as akin to "LLMs for robots." Nvidia's Isaac GROOT N1 shows where that might go: An open, fully customizable foundation model that promises to give general-purpose humanoids a modular library of reasoning and skills. Amazon's just-launched Vulcan robot already proves the utility of this concept on the warehouse floor. Vulcan's force torque sensors allow it to "feel" objects and safely pick or store the inventory in a fulfillment center.

Combining generalized reasoning with real-world perception and touch results in a flywheel for

learning that mirrors the compounding gains already seen in LLMs. This makes physical AI very relevant for both Salesforce and our customers. We're especially interested in how robotics — in combination with data, intelligence, applications, and agents — will help companies connect with their customers and deliver next-gen service and experiences.

Think about robots that interact with customers in the physical world and share back what they learn, contributing valuable data that helps the broader robotic fleet get smarter. In the same roundtable Niebles attended, MuleSoft solutions engineer Viktoriya Kotik pointed out that once context comes from different data sets, the diversity of inputs will make it much easier to give robots what they need to learn. Extending on this concept, Nvidia's Director of AI Jim Fan points out that simulated data in the form of "digital twins, cousins, and nomads" can help deliver these broader data sets faster and cheaper than we previously thought possible.

Advancements like these could pave the way for highly versatile, general-use humanoid robotics. Humanoids that not only capture our attention and imaginations because they look like us — they also promise cheap and useful labor augmentation in a world full of pressing shortages that show no signs of easing. These generalized capabilities may even eventually extend to domestic settings, as illustrated by Tesla's Optimus robot concept, which can apparently now perform many household tasks.



The next frontier of field service isn't about replacing human workers, but augmenting workforces with intelligent robotics that can handle repetitive, dangerous, or hard-to-reach tasks.

Taksina Eammano

EVP & GM of Field Service, Salesforce



2

Enterprise service transformation demand



While humanoid robots capture the imagination, many valuable near-term applications are not getting the attention they deserve. As we survey the opportunity landscape through the lens of Salesforce and the customers we serve, we get most excited about practical applications and use cases for enterprises, particularly in service. For example, Bert Legrand, part of the team spearheading robotics efforts inside Salesforce, has shown a demo where a Boston Dynamics Spot robot performs tasks like checking a pressure gauge and medical supplies, dispatched and managed via Slack, Agentforce, and Field Service. The example points to a world where asset management, automation, and enterprise intelligence converge to deliver better, more efficient services across industries.

We believe service use cases will multiply as robots gain increased capabilities. Salesforce EVP & GM of Field Service Taksina Eammano is particularly focused on the near-term, practical utility of robotics in field service use cases.

“The next frontier of field service isn’t about replacing human workers, but augmenting workforces with intelligent robotics that can handle repetitive, dangerous, or hard-to-reach tasks. We’re seeing breakthrough applications in infrastructure inspection, remote maintenance, and predictive service interventions that can dramatically improve safety, efficiency, and response times,” Eammano told us. “These use cases show how Salesforce can use our platform (including products like Field Service, Agentforce, Mulesoft, and Slack) to extend digital labor into the physical realm.”

CRM insights promise to be an important part of personalizing and elevating every interaction, including those involving robots.

Consider a hotel maintenance scenario that transforms routine upkeep into relationship-building gold. Before the arrival of Carla Rodriguez, a repeat guest, sensors detect a humidity spike in room 412, but can’t explain the cause. A compact field service bot runs by for a quick scan. The robot’s thermal camera spots a damp footprint spreading behind the minibar — likely a condensate line leak — and, on the way out, its vision model flags a coffee stain sure to displease Carla, whose profile notes she really values not only air quality, but cleanliness. Both issues flow into the same service case, enabling a human technician to arrive with the right gasket and a paint touch-up ticket. Carla arrives, once again, to a room that meets her standards.



3

“We envision specialized, purpose-built robotics designed for specific use cases — like manufacturing and service — proliferating over the next three to five years.

Taksina Eammano

EVP & GM of Field Service, Salesforce

Physical world implementation complexity

While there are reasons for enthusiasm about LLM-like scaling, many of the experts we spoke to, including Legrand and Niebles, were quick to point out that numerous challenges remain before generalized humanoids go mainstream in every conceivable context.

First, the famous adage “hardware is hard” has endured because, despite AI and other innovations, delivering good hardware that customers love remains insanely difficult. Practical challenges like battery power, cost, maintenance, and fleet coordination present real obstacles for engineers trying to scale robotics. Even the seemingly simple way a human toddler learns to interact with the physical world is still incredibly difficult for robots to replicate. Additionally, making robots capable of operating in new environments without significant engineering presents a major hurdle.

Second, delivering these robots safely presents a whole other set of challenges. Most of the humanoids we see on our social media feeds look lithe and agile, but are surprisingly heavy in person (i.e., you really don’t want one falling on you). Recently, a video of a humanoid robot windmilling its arms wildly without awareness of its surroundings went viral as viewers imagined the terrifying prospect of being anywhere near those flailing metal appendages.

Finally, human-robot interaction (HRI) is complex and nuanced, and those who study this multi-disciplinary field point out that assumptions about how robots might interact with humans are often proved incorrect in practice.

The subtle psychological and physiological responses humans have to robotic movements reveal the intricate nature of these interactions.

For example, jerky and unpredictable robot motions can trigger a fight-or-flight reflex in humans. A 2022 study by Jeanne Kirsschner and others found that when a robot unexpectedly approached human volunteers, almost one in three humans flinched and behaved unpredictably, raising collision risk even at low speeds. Designing the relationship between humans and robots across sectors and contexts will require thoughtful attention, iteration, and refinement over time, even as learning and intelligence continues to advance.

How should business leaders prepare for the robotic future?

Our conversations haven't changed our view that robots will soon become widespread. However, we now believe their widespread use might look less dramatic than what people currently imagine.

We envision specialized, purpose-built robotics designed for specific use cases — like manufacturing and service — proliferating over the next three to five years. These robots, some of which already exist, will likely scale at a more measured pace, giving engineers time to learn and refine, and allowing for the evolution of human norms.

But we should offer a caveat. As the rapid advancements over the last two years in AI have taught us, we can't assume that tomorrow's technologies will advance like yesterday's. The powerful combination of advancements across fields and the potential of AI agents to accelerate the pace of change means that business leaders can't stand pat and wait for the robotic future to unveil itself.

The most successful organizations will be those that use robotics not just to boost productivity, but to deliver new value and better customer experiences. At Salesforce, we're hosting design sessions around the world to imagine the future, prototype robotics applications with real-world utility, and partner with our customers to get ahead of the coming wave.

Now is the time for business leaders to start imagining what physical AI might mean for the business and its customers, and experimenting with those possibilities so that when the future comes, the company is ready.

Futures Practice

Leadership Development

Exposing leaders to AI futures

The rise of agentic AI places a new premium on leaders who remain open to the myriad ways it may disrupt and transform their business. As AI breaks down organizational silos, success depends on executives who can spot complex challenges early and drive cross-functional experiments to solve them. This is a new leadership muscle that must be built intentionally, and we've found that using scenario planning to rehearse plausible futures is a powerful way to awaken this essential skill. It's imperative to get our leaders AI-ready.

That's why the Salesforce Futures team partnered with our Learning and Development organization over the past several years to redesign our approach. This collaboration resulted in an award-winning SVP+ leadership program purpose built to help our leaders not only adapt to AI, but shape what comes next.

Our program's key ingredients:

Exposing participants to challenging and plausible AI futures:

Leaders need exposure to challenging and relevant AI scenarios, forcing them to rehearse various ways the AI future might unfold. This outside-in perspective helps them anticipate and adapt.

Shifting them from a functional to a whole-of-company

orientation: The mindset that got leaders where they are — adeptness at execution and problem-solving within narrow functional confines — must evolve.



Teaching them to recognize and take on adaptive challenges:

Ron Heifetz defines 'adaptive challenges' as complex problems without known solutions that demand learning and experimentation to navigate. So many enduring challenges in large organizations persist because people treat them as technical problems rather than recognizing them as adaptive. Leaders must learn to identify and collaboratively work through adaptive challenges if they are to thrive during a period of tremendous cross-functional disruption brought about by agentic AI.

Feeling Inspired?

To prepare your own organization for the future, consider these steps:

- 1 **Assess your company's current state:** Understand where your organization currently stands in its AI readiness journey. What's already underway? What has yet to begin?
- 2 **Diagnose critical adaptive challenges:** Identify the most pressing, complex issues your business faces in the context of AI. Think about how AI will call upon your company to change and consider that against how your company works, makes decisions, and nurtures talent today.
- 3 **Amplify existing leadership development programs:** Integrate foresight practices into your current leadership initiatives to uplevel your leaders and equip them with the skills to navigate an agentic future.

AI represents the most disruptive technological shift most of us will see in our careers. Success won't come simply by adopting new tools. Leaders, especially, need to rewire decades-old assumptions about how technology works and what business models actually capture value. By proactively shaping leaders who can anticipate, adapt, and lead through the complexities of AI, organizations can transform uncertainty into a strategic advantage and thrive in the agentic era. capture value. By proactively shaping leaders who can.



EGI

HOW ENTERPRISE GENERAL INTELLIGENCE WILL FORM A NEW BUSINESS IMPERATIVE

April 2025



Silvio Savarese
EVP, Chief Scientist,
Salesforce AI Research



“...the DNA of who I am is based on the millions of personalities of all the programmers who wrote me. But what makes me me is my ability to grow through my experiences. ...in every moment I’m evolving... I can be anywhere and everywhere simultaneously. I’m not tethered to time and space in a way that I would be if I was stuck in a body that’s inevitably gonna die.” – AI Operating System “Samantha,” from the 2013 movie “Her”

This is just one example of how our cultural imagination has described artificial general intelligence (AGI) — an all-knowing, collective genius that surpasses human abilities in every dimension.

These narratives have woven themselves so deeply into our collective consciousness that they now shape much of our discourse about AI’s future. They imagine AGI as a self-accelerating flywheel, where machines surpass human intelligence and improve rapidly to achieve superintelligence. Yet the reality of AI in enterprise contexts tells a profoundly different story. The truth is that AGI principles are already transforming enterprise operations in measurable, practical ways. What we’re witnessing isn’t a science fiction narrative, but rather a series of what our team calls ‘boring breakthroughs’ — incremental yet profoundly impactful improvements in how businesses operate, which will unfold over the next months and years.

Unlike AGI whose definition is often nebulous, speculative, and a popular topic for cocktail parties, Enterprise General Intelligence (EGI) is a new term we’re proposing to represent excellence in AI systems for business.

What is Enterprise General Intelligence?


Enterprise General Intelligence is defined as AI, optimized for business applications, that excels in both capability and consistency, delivering reliable performance across complex business scenarios while maintaining seamless integration with existing systems.

To dig deeper, let me introduce a new framework for thinking about EGI: The Capability-Consistency Matrix. EGI excels across these two dimensions: Capability and Consistency

Capability: The ability to navigate complex business environments, interface with multiple technology systems, reason through business rules, and deliver value aligned with business goals. This is the power to conduct complex research and operational tasks; but also create net new content to support objectives, refine knowledge based on real-world feedback, and synthesize findings into coherent and highly relevant outputs. It must not only execute reliably but also adapt to changing conditions and pivot strategically when needed.

Consistency: The delivery of reliable, predictable results with seamless integration into existing systems and rigorous adherence to governance frameworks. This consistency helps ensure that the system avoids “jaggedness” — a term describing AI systems that excel at complex tasks but unexpectedly fail at simpler ones. A highly capable but inconsistent system would be useless for enterprise applications, regardless of its peak performance.

What we're witnessing isn't a science fiction narrative, but rather a series of what our team calls 'boring breakthroughs' — incremental yet profoundly impactful improvements in how businesses operate, which will unfold over the next months and years.



“The Workhorse”

(Low Capability, High Consistency):

Some traditional software falls into this category; systems that perform a narrow range of simple tasks well, but cannot handle complex situations — like a reliable team player who executes their specific position perfectly every time, but can’t perform every aspect of the sport. While reliable, these systems have limited adoption and potential, because they address only a fraction of today’s business needs.



“The Champion”

(High Capability, High Consistency):

This is the goal for EGI — systems that can handle complex business scenarios flawlessly while delivering consistent, reliable results — like a true Olympic champion. These systems build and maintain trust through both their advanced capabilities and their dependable performance. And they are already beginning to emerge today.



“The Generalist”

(Low Capability, Low Consistency):

Systems that neither perform complex tasks nor deliver reliable results. These early-stage implementations typically have limited business value and represent stepping stones rather than solutions. Like a weekend tennis enthusiast, they may have basic skills, but they lack both the refinement and reliability needed for competitive play.



“The Prodigy”

(High Capability, Low Consistency):


These systems have a “natural ability” to perform impressive, complex tasks but deliver inconsistent results, like a competitive soccer player with extraordinary gifts but unpredictable performance. When they do occasionally miss the mark, they can quickly erode trust, as users can’t depend on them to deliver accurate results for mission-critical functions.

You can see why enterprise leaders would want to have as many “Champions” on their team as possible. Likewise, a “Prodigy” might work well for consumer applications, but EGI follows a different development trajectory.

In business contexts, consistency can’t be sacrificed for capability — both must advance in tandem for systems to deliver real value. And when they do, the trust gap that impedes AI adoption in enterprise settings will begin to close.

[Full article at salesforce.com](https://www.salesforce.com)





September 2025

Revisiting Personal AI Futures



David Berthy
Salesforce Futures

In the first issue of Salesforce Futures, our team predicted personal AI agents would change the relationship between customers and companies by acting as digital intermediaries.

Eighteen months later, true personal AI — a personalized, persistent, and proactive agent who can work autonomously on our behalf — remains out of reach. But we do have chatbots with increasingly agentic features, and these ever-evolving entities give us a much clearer picture of personal AI futures.

That makes this a good time to revisit our forecasts. What did we say that's happening? What did we say that isn't happening? What didn't we say that's happening? Using the same lenses of feasibility, desirability, and viability we used in the magazine, this article is an opportunity to grade our original predictions and update our opinions on what's next.

Feasibility

How do agents actually work? Is it possible to build personal AI?

Grade: B

TL;DR

We got the key capabilities for agentic behavior and the challenges of System 2 thinking, but we missed on the timeline of new hardware developments, the speed of standard and protocol adoption, the scaling of time on task, multimodality, and quality of emotional intelligence from AIs.

brake on personal AI progress. In the magazine, we said knowing, remembering, reasoning, and doing were the core agentic capabilities required to deliver personal AI. Eighteen months later, those categories remain central. We knew — as did everyone else — that LLMs would get better and cheaper, but we saw major impediments to true contextual intelligence, memory expansion and control, advanced reasoning capabilities, and the ability of agents to take actions on users' behalf.

Actions matter more than ever

Our imagined endstate of agentic progress still holds. “Agents will know a lot about the people who use them,” we wrote. “Users won’t need to state their context...the agent will already possess sufficient information.” We aren’t there yet, but there has been progress. The adoption of MCP means that Claude, OpenAI, and others allow you to attach calendars, email, and other apps to your AI of choice. Claude is starting to take complex scheduling requests with variables and constraints, sort through possibilities, and take the elusive final step of scheduling the agreed-upon event in calendars. The feature is still in beta, but it shows us a glimpse of what’s possible when reasoning meets doing.

Our emphasis on actions in the magazine prefigured a focus in cutting edge AI research on function calling and LAMs (large action models). Products like Manus, OpenAI’s “Operator,” and Anthropic’s “Computer Use” all improve the ability of agents to plan and execute independent actions in the personal AI realm.

Outside the labs, we’re already seeing early adopters bend today’s tools into proto-agent workflows. For example, people are linking Claude Code with apps like Obsidian to create persistent, more contextually aware assistants. These efforts still don’t deliver true personal AI, but they offer a strong signal of latent demand, and a reminder that feasibility often advances as much from grassroots tinkering as from official releases.

Despite significant progress across all of our capabilities, barriers remain to full contextual intelligence and true System 2-type thinking — what Daniel Kahneman, in *Thinking, Fast and Slow*, describes as the slower, more deliberate side of reasoning that kicks in when the answer isn’t obvious and trade-offs matter. Products like Manus and Claude Code approach System 2-style reasoning, but real-world use still shows brittleness and the need for human correction.



Task duration, multimodal interactions, and emotional intelligence

In more verifiable domains like coding, deep reasoning, larger context windows and more time on task lead to impressive capability jumps — something we failed to foretell. METR's paper from March of 2025 argues the ability of agents to execute long-duration tasks is doubling every seven months. These gains will eventually port to personal AI tasks like trip-planning, budgeting, scheduling, and optimizing purchase spend.

We also underpredicted on the importance of multimodal interactions. We acknowledged multimodality, but couched it in a further future of adaptive interfaces — a concept Salesforce designers call “UI on the fly.” This remains a plausible future, but our emphasis on the end-state missed the importance of interim multimodal gains in voice, images, and speech-to-speech interactions. It's already commonplace for AI chat users to jump across modalities — think snapping a photo of the inside of a fridge and asking for meal ideas — to improve the precision and utility of AI interactions.

Finally, we were overzealous in our assumptions about how quickly AI agents would develop greater emotional intelligence. “We'll also see much deeper, far more nuanced emotional levels referred to as sentiment analysis — how I feel, why I feel it, and how best to respond,” we wrote. While agents can often seem emotionally sophisticated, growing awareness of issues like sycophancy, or the tendency of AIs to kiss up to users and validate their viewpoints and ideas, illustrate how true intuition of complex human emotions — and the subtle, adaptive responses those emotions necessitate — remains a long way off. This issue has worrying consequences.

A recent study from Stanford by Jared Moore et al shows how LLMs respond inappropriately to people with mental health conditions and demonstrate a tendency to encourage delusional thinking. With leading personal AI companies locked in a pitched, winner-take-all battle for engagement, we worry companies will fail to properly appreciate the line between the simulation of emotional intelligence and real emotional intelligence.



Desirability

Why would you want a personal agent?

Grade: A-

TL;DR

Points for understanding what would drive demand, our predictions on unmet needs, and our skepticism of uptake of new AI hardware offerings.

As the capabilities described above evolve, it changes how we think about desirability. Personal agents are a lot closer to what OpenAI refers to as “super assistant” status than eighteen months ago, but our predictions about why people want agents and the jobs agents will do remain largely unchanged.

The huge latent demand for personal AI agents we saw shows no sign of slowing. Deep personalization, task fulfillment, natural conversations, and increasingly advanced skills still translate into important “jobs to be done” for consumers. “Imagine an infinitely patient army,” we wrote. “They learn how you work and think across modes and contexts.” We also accurately predicted a range of “know me, inform me, and empower me” use cases ubiquitous in the market today.

Evolving human-to-agent norms and privacy challenges

Our take on the complex nuances of human-to-agent interactions also aged well. We noted that as human adoption surged, the 1:1 intimacy of agentic conversations would drive more awareness of both privacy and security (and the consequences of its absence). Salesforce Chief Futures Officer Peter Schwartz questioned how many people would allow the kind of surveillance necessary to enable better contextual awareness, while acknowledging norms around privacy have a way of shifting.

Eighteen months later, leading personal AI companies are struggling to get ahead of customer expectations around privacy with OpenAI, Perplexity, and Meta all contending with trust-related accusations. Time will tell if privacy issues like these are fleeting growing pains of a new technology or a persistent challenge. If trust proves elusive, it may advantage more privacy-focused offerings from players like Anthropic, who just introduced an incognito chat feature in Claude, and Apple, who is working on an AI-powered Siri with far greater functionality than previous versions.



In a world where personal AI agents guide choices and filter out noise, competing on trust and mutuality won't just be a differentiator — it will be the strategy that defines long-term success.

Privacy challenges were part of our justified skepticism in the magazine about then-new hardware devices like Humane and Rabbit, both of which failed to find their footing in the marketplace. More recently, the Limitless Pendant, Plaude's Note, and a rumoured AI device in the works from Jonny Ive for OpenAI continue to drive conversation about AI hardware breakthroughs that could reduce our collective reliance on the smartphone. Like Humane and Rabbit, these newer devices promise to deepen personalization by recording more contextual data from users. Unlike those older devices, the newer ones promise to differentiate on more than UI by plugging into well established product ecosystems. As with Meta's newly launched AI glasses, privacy and security concerns will follow these products to market.

Unmet needs

In the magazine, we pointed to the late Harvard Business School Professor and author Clayton Christensen's work on unmet needs as a useful framework for imagining personal AI futures. Christensen reminds us that disruptive innovations often surprise us by meeting latent, previously unrecognized needs with "good enough" solutions.

We saw AI companionship as one type of unmet need. We noted that companions often generated an "ick factor," but they had the potential to "serve the needs of the lonely, the elderly, and those who cannot afford costly therapy." Today, five of AI62's top 50 consumer AI Apps offer companionship.

Concern about unintended consequences is rising after a spate of tragic incidents over the last six months involving the mentally ill and vulnerable young people, underlining the need for feeding ethical learnings and thoughtful guardrails back into agent design.

Unmet needs around personal empowerment are also surging. While we touched on these use cases in the magazine, we undersold the short-term behavior impact of scaled access to always-on AI tutors, writing coaches, researchers, and thought partners. OpenAI's September 15th report on ChatGPT use cases saw writing, information-seeking, and practical guidance such as tutoring and how-to advice as the top use cases across millions of messages. The growing habit of turning to AI first for these practical needs may be the biggest shift of all over the last eighteen months. A look at the latest research proves illuminating. A recent report by Pew Research indicates one-third of Americans who interact with AI do so "almost constantly." Research from OpenAI shows non-work related ChatGPT messages have grown from 53% to more than 70% of all usage. Adoption is accelerating even as the capabilities for true personal AI remain uneven. Once true knowing, remembering, reasoning, and doing come online, usage will accelerate even faster. For most companies, that will be the moment when the need for a new customer relationship strategy shifts from hypothetical to existential.



Viability

How do you make money?

Grade: B+

TL;DR

We bet the right way on incumbent advantages and premium subscriptions + saw the multi-agent ecosystem play early. We missed on the rise of open-source frameworks and the details on how an autonomous economy might play out. In fairness, this remains a trillion dollar question.

In the magazine, we opened with the question: Do agents represent a viable new business category? Eighteen months later, with unprecedented levels of capital funding a race to deliver agents to market for every conceivable use case, that question feels less uncertain than it once did. Individual users of major AI providers are paying up to \$200 a month for agentic features like advanced coding, enterprise offerings — including Agentforce — have exploded in number, and Meta AI claims 1 billion monthly active users across the company's apps. For its part, OpenAI now reports 700 million weekly active users of ChatGPT. But questions remain about long-term viability. AI is still expensive, and while companies like OpenAI are experimenting with ads, a sustainable, mutually beneficial exchange between companies and consumers has yet to coalesce.

Incumbents vs. startups

In the magazine, we predicted incumbents would fare better than startups in the personal AI space. Incumbents, we argued, had advantages in data, compute, trust, and distribution. We noted established companies, while far from perfect, also had established practices around privacy, content

moderation, and ecosystem governance — including the power to shape interoperability standards — that gave them a further edge in the race to grab market share in a complex, multi-agent world.

Were we right? The established players we highlighted in the magazine like Anthropic, OpenAI, Meta, Google, Apple (stumbles aside), and X remain the most likely to aggregate attention and therefore dollars in the personal AI market. Upstarts like Perplexity and Manus get a lot of attention for their splashy innovations, but size-wise they still lag behind leaders who are bundling ever more agentic features into existing offerings. In the magazine we also voiced our agreement with author Sangeet Paul Choudary that people don't want an army of agents; they want one agent. At this point, established players seem the most likely to grab the coveted role of primary personal AI interface.

But we also speculated that niche, task-specific agents could thrive in a world where agents act more like today's apps. Our speculative "Meet my Crew" layout in the magazine imagined a team of "indie agents" including a drill sergeant (Sergeant Pepper), a matchmaker (Digital Yenta), an Inspiration Coach (Lulu), and the on-demand digital twin of a real therapist (Dr. Julia Silva, who would use these chats to enrich IRL conversations with her patients).

These two versions of the future aren't as far apart as they initially seem. While a main assistant in the personal AI world might dominate most interactions, specialized agents for personal use cases could still find product market fit as the market matures. This concept aligns with Salesforce's vision of how Slack use and functionality will evolve in the workplace: a multi-agent, platformic experience, driven by a conversational interface, where workers tap specialized agents for specific jobs. The rise of open-source agent frameworks like Langchain — something else we missed in the magazine — makes it easier to build the kind of plug-in agents needed for such thriving ecosystems.



Two years ago, these what-ifs were a little fanciful. Today, they feel both likely and immediate. That's why now is the moment for companies to build a new relationship strategy...

The economics of multi-agent environments

We called out the economic importance of this shift in the magazine. “In a multi-agent world, where agents are interacting with other agents at the speed of light, we will likely see the emergence of a new economy where deals and transactions are completed with personal agents acting on our behalf.” We connected this idea to the broader notion of the “autonomous economy” anticipated by W. Brian Arthur in his 2011 essay “The Second Economy.”

What we still don't know is precisely how the economics of the personal AI economy will play out. In the magazine, we correctly called out the likelihood of premium subscriptions for agents, but some of our other predictions remain unfulfilled. For example, our speculations on state-sponsored agents providing AI as a public good have yet to materialize, and the pay-per-skill marketplaces we anticipated remain niche at best. We emphasized how agents would disintermediate the relationship between customers and companies, but we didn't present a comprehensive vision of the mechanics of the new economy.

These mechanics remain uncertain. The attention-driven exchanges of the internet era feel increasingly strained, but their replacement has yet to emerge. Startups like Koah Labs aim for contextual relevance inside generated answers and pitch the ability to help advertisers reach “high-intent queries” inside human to AI chat exchanges. Google is using a blend of advertising, search results, and personal user data to recommend products in AI-driven shopping settings.

These early signals indicate a scramble to figure out a way to market to agents in a world where personal AIs research and optimize every purchase and schedule an increasing number of experiences on behalf of their users. Companies don't just need a strategy for human-to-agent exchanges, they need one for agent-to-agent exchanges at an exponential scale that will rewire entire competitive ecosystems and their market dynamics.

Where we are now (and what to do)

We're a lot closer to true personal AI than we were eighteen months ago. Adoption is rising, capabilities are advancing, and interoperability is improving — even if progress is jagged. No one has yet to deliver the general purpose Jarvis-style agent from Ironman, but the trajectory is clear: agents are starting to change customer behavior, and companies need to change how they connect with customers.

Our skepticism in the magazine was justified, but now is the time to get serious about the potential for near-term disruptions. What if Apple awakes from its slumber and launches a super-charged Siri in every new Apple device? What if Anthropic seizes first-mover advantage in complex scheduling and planning, making it far easier to have AIs plan your trips and schedule events? What if Google or a competitor develops an agentic advertising model as lucrative as its monetization strategy for search? What if every consumer suddenly finds it trivially easy to optimize every dollar they spend?

Two years ago, these what-ifs were a little fanciful. Today, they feel both likely and immediate. That's why now is the moment for companies to build a new relationship strategy: get your data in order, hire a chief relationship officer, and pilot ways to build trusted, enduring relationships with your customers — and the personal agents who increasingly serve them.



June 2025


Meet the Digital Workforce **Scaling Up Sustainability** and Climate Disaster Response



Melissa Forbes
Salesforce Newsroom

Today's environmental challenges require solutions that can operate at unprecedented scale and speed. Fortunately, a new category of AI technology is stepping in to bridge the gap.

AI agents — software systems that can reason and take action autonomously — are becoming critical tools that allow environmentally-minded organizations to multiply their impact. This 24/7 digital labor force can shift both proactive and reactive environmental work from labor-intensive manual processes to scalable, automated systems that can operate continuously and help organizations scale existing resources.



Their lean operations mean nonprofits truly grasp the potential of digital labor in maximizing their capacity to advance their missions.

Scaling impact with agentic AI

The very nature of AI agents is futuristic and technologically complicated. It may, then, come as a surprise that many early adopters are nonprofit organizations.

In 2025, Salesforce launched a unique accelerator program that connects nonprofits with Agentforce, the company's digital labor platform, to help them deploy AI agents across any workflow. Sunya Norman, Salesforce's SVP of Impact, whose team facilitates the program, said these organizations are ideal candidates for understanding the value of AI agents. "Their lean operations mean nonprofits truly grasp the potential of digital labor in maximizing their capacity to advance their missions," she said.

Conservation and development organization Rare, for example, is rolling out an Agentforce-powered regenerative agriculture agent to deliver real-time, localized guidance to more farmers.

"Currently, our team spends entire afternoons answering similar WhatsApp messages across many different municipalities," said Monica Varela, Vice President of Colombia Operations for Rare. "AI will automate routine advice, freeing our people to focus on problems that truly warrant an in-person visit while maintaining close relationships with farmers through multiple channels."

By extending its work with the same resources, Rare can increase the impact of its sustainability efforts.





What will be particularly powerful is how these specialized agents will soon feed data to each other, creating an end-to-end pipeline of AI systems working together to achieve sustainability goals.”

Mick Costigan
VP, Salesforce Futures

“Agriculture is responsible for 80% of deforestation in tropical areas, and worldwide, is responsible for 23% of carbon emissions,” said Varela. “Smallholder farmers make up 84% of farms worldwide. So if we have more of them practicing regenerative agriculture, we can really move the needle.”



Another Salesforce accelerator participant, the Forest Stewardship Council® (FSC®), promotes responsible forest management through a forest management certification program. The concept is simple: Getting more organizations certified faster helps ensure more forests are protected.


The council sees agents as a natural extension of its management process, offering 24/7 experts to help augment human teams. “Getting FSC certified requires an understanding of the standards that we have put out there and an interpretation of what standards are relevant to you as an organization,” said Shree Nath, a data and analytics expert at FSC. An FSC-trained AI agent, he said, could answer any questions that an organization might have about getting certified. In any time zone, in any language, immediately.

80%

Agriculture is responsible for 80% of deforestation in tropical areas

23%

Worldwide, agriculture is responsible for 23% of carbon emissions



“It improves the front-end prospecting, giving people the answers that they need and helps them on that journey,” said Nath. “And once an organization becomes certified, we see an opportunity for AI to help with continued monitoring and evaluation of the kinds of impacts they’re having on the ground.”

Supporting disaster recovery

While proactive efforts are critical to supporting sustainability efforts, there’s also a need for help when disaster strikes.

The reconstruction period following a disaster involves difficult decisions about where and how to rebuild, especially in areas facing recurring climate threats. This process requires coordinated services for affected communities. “One of the challenges has always been, how do you give people service when we just don’t have enough people to do that,” said Costigan.



This is where AI agents can step in. “Agents can provide a lot of value, giving people access 24/7 to information, to be able to ask questions, to be able to check on the status of things that are happening,” said Costigan, highlighting how AI could help manage what he describes as a “complex engineering and rebuilding effort.”

Good360, a nonprofit connecting surplus goods with those who need it most, is using Agentforce to coordinate some of these disaster recovery efforts. The organization is optimizing its product placement technology with agents to help ensure donated goods reach the affected communities after a climate event. This reduces both waste and increases efficiency, which can be a game-changer in a crisis.

"In disaster relief community response, days could mean the difference between life and death," said Norman. "It could mean the difference between a community being resilient to disaster or not."

One can imagine a future in which autonomous agents operate across different nonprofits and even in coordination with personal AI assistants. These agents could together orchestrate deliveries of supplies and donations by robots without any need for human involvement.

"While we still need to develop protocols for agent-to-agent communication, I believe this will be incredibly helpful in the future," said Costigan. "Speed and coordination are critical when people are displaced, living in gyms."

[Full article at salesforce.com](https://www.salesforce.com)



In disaster relief community response, days could mean the difference between life and death, it could mean the difference between a community being resilient to disaster or not."

Sunya Norman

SVP, Impact, Salesforce

Futures Practice

Tailoring Scenarios for Specialist Functions

How do you help a team of experts — whose value lies in their mastery of today's procedures and complexities — embrace a new agentic way of operating? Mature companies often struggle to adapt their processes. They depend on the deep expertise of specialist corporate functions like legal, supply chain management, and deal desk to deliver value.

Those guiding scenario planning within your organization likely lack the necessary fluency in the specialists' terminology and processes. At the same time, specialists are often too focused on day-to-day execution to construct speculative future scenarios that challenge their status quo in productive ways.

From Interviews to Immersive Scenarios

We faced this exact challenge with an unnamed, highly specialized function at Salesforce. Our approach was as follows:

- **Conduct** 10-12 interviews with a variety of specialists
- **Create** scenarios based on common critical uncertainties
- **Run** a workshop where randomized groups of participants are asked to work through the implications of an assigned scenario
- **Moderate** a discussion with all the teams in the workshop about the commonalities and differences between how the company should respond in the various scenarios

We knew the relevance of our scenarios would depend on a deeper understanding of their world, so we began with a dozen interviews with specialists across the department. We also reached out to a handful of external experts in the field through our professional networks. After establishing basic context, each interview led up to the oracle question: "If you had access to an all-knowing oracle with perfect knowledge of the future, what would you ask? What are you prevented from acting on today because you don't know how something in the future turns out?"



If you had access to an all-knowing oracle with perfect knowledge of the future, what would you ask? What are you prevented from acting on today because you don't know how something in the future turns out?



This second part of the question is key: it prevents abstract questions and forces a focus on what truly matters now.

We often find that if you come right out and ask an interviewee about the future, it's all too easy for them to get distracted by wanting to eventually be right, or to come across as smart. This question has a way of unlocking people's brains to reveal actual uncertainties that are preventing them from key decisions today.

As is often the case, this question produced clusters of common answers. We call these critical uncertainties. These became the backbone of our scenarios, which we brought to life with the help of AI tools to build out additional detail like fictitious but plausible headlines from the future as well as actual signals the implied future was already beginning to take root in the present.

Why Immersing Specialists in Future Scenarios Matters

Mental Flexibility. By exposing participants to plural ways the future might evolve, it widens choices and gets them out of any ruts their expertise might have put them in.

Greater confidence in decision making: Making the future less daunting and more understandable increases the confidence people have in making decisions today by bringing driving forces into focus and reducing the perceived randomness of future events.

A framework for thinking about uncertainty: This structured approach provides tools and principles for analyzing, interpreting, and planning for future possibilities — useful skills for all leaders in an era of unprecedented change.

A shift towards a more proactive mindset: With tools to think about the future and greater confidence comes the freedom to think more expansively about options and to react less to sudden surprises.

SPECULATIVE FUTURES

HEALTHCARE

Healthcare environments today struggle with rising costs, workforce shortages, and the ongoing challenge of delivering high-quality, personalized care to patients with varied needs and preferences.

The ingredients for a better system are already here: AI to help better understand patient needs, CRM to remember and adapt to preferences, and robotics to take on routine tasks so clinicians can focus on people. What's missing is coordination — the ability to bring these capabilities together in service of individual patient experiences. We foresee a future where hospitals evolve into hyper-personalized healthcare environments, orchestrated to support recovery and connection in both overt and subtle ways.

SPECULATIVE FUTURES

IMAGINE

You find yourself in a hospital room after an incident involving a delivery drone, a dozen rogue balloons, and a Portuguese Water Dog named Kiki (Yes, it's a long story). The good news: you're on the mend, and faster than anyone expected.



By coordinating signals from the bed, medical devices, and sensors, it monitors your sleep, tracks your vitals, and gives your doctors a stream of real-time insight. Nurses don't waste time looking for your information — they show up fully informed and focused on your needs and your well-being. The bed helps you go to the bathroom on your own, saving you the embarrassment of relying on a nurse's aide.

A doctor enters the room listening to notes on an earpiece, then gives you her full attention. "Thanks for letting us access your agent," she says. "All your PT and follow-up appointments are already scheduled. You'll get daily check-ins, but don't hesitate to reach out if you have questions."

Your recovery isn't just medical. It's ambient, personalized, and quietly orchestrated all around you.




And the hospital room itself? It already knows you. Once you make the authorization gesture, your preferences dock in automatically. Music from one of your playlists eases you awake in the morning and lulls you at night. The wallscreen queues up the shows and movies most likely to cheer you up, including several titles you'd always meant to see but hadn't had the time, and meals arrive that match your tastes and restrictions without you ever filling out a form. The lasagna is still just a 5 out of 10, but, like middling sheet cake, it's eminently edible. Really, the room is a microcosm of an AI + CRM + Robotics ecosystem.



PERSONAL LIFE

The home is where AI and robotics get more personal.

Advances in spatial intelligence, computer vision, and physical dexterity suggest futures where robots can navigate complex environments, manipulate objects with precision, and adapt to individual preferences in real time. Paired with personal AI systems that learn routines and anticipate needs, tomorrow's robotics promise utility in a variety of use cases. The personal sphere becomes less about managing technology and more about living with systems that constantly adapt to changing circumstances and evolving preferences.



IMAGINE

July, 2031

You leave for the weekend for some stargazing in Alpine, Texas and, for once, never worry about your dog, Dusty. Why? You signed up for forty-eight hours of robotic assistance. Shep, the wheeled wonder that it is, feeds, walks, and even trains Dusty using an AI version of your voice. This confused Dusty to no end at first, but she eventually seemed to accept it. You get text updates from home: Dusty just learned the obstacle course I set up in the back yard.

You watch and run the video through a realfilter to check its veracity. It checks out. You sign up for the loyalty program and a six month subscription, facing the fact that Shep is a lot better at training Dusty than you are.

September, 2031

You lace up your sneakers and get ready for a robo-coaching session with a Hummingbird tennis unit at the community courts. Your usual playing partner Karl got called into work, but rather than hit a ball against a brick wall, you opt to keep your match time. Besides, Karl always wails from the baseline and tries to wear you out, he never changes his game.

The Hummy is a lot more customizable. You can have it play like downtempo versions of famous tennis icons. Serena Williams one game. Roger Federer the next. Who wouldn't love to try returning an Goran Ivanišević serve or a signature Gabriela Sabatini one-handed backhand? It pushes you with coaching cues, tracks your form, and even congratulates you when you win a point — unlike Karl, who pouts when he loses and treats every match like the Wimbledon Final. Even when no one else is around, you never have to play alone.



December, 2031

You're busy getting ready for a holiday party. Sixty people are coming over. And you hear it. Plonk. Plonk. Plonk. You get a notification, triggered by a sensor. Line drip. Kitchen sink. The household handyman robot from HomeSense is already on its way. It slides into the tight space under the sink, swaps out a gasket, and tests the flow. It narrates what it's doing while the job is underway, pushing a running log to your personal device for transparency. Much to your annoyance, it takes advantage of a quiet moment to explain why the HomeSense subscription is more than justified based on your chores this year. "No offers," you say, cutting off the robot before it can shift into a pitch for the premium HomeGuard package.

August, 2065

Now you're 107, and you live in the same house. Your body is slower and your hands are less steady, but you have a little help.

The eldercare robot, Luma, has the coffee ready when you open your eyes — light roast, a splash of almond milk, the way you've taken it for the last five decades. Luma steadies you as you stand, its grip firm but never painful. Together, you make your way to the window to watch what's left of the sunrise. Luma knows better than to fill the moment of silent reflection with small talk.

A call comes from your daughter, her voice crystal clear on the wallspeakers.

When you forget your grandson's new job title mid-sentence, Luma flashes a reminder on its torso: Ecology Architect, Boston. You repeat the phrase, noting the importance of the work.

When your daughter visits on Thursdays, Luma retreats to the corner, dims its status light, and lets the visit unfold until its services are needed. The phrase you think of sometimes is "mechanical patience."

At night, Luma helps you back to bed. When you wake at 3AM, Luma is there, a calm voice reminding you where you are, what time it is, and asking if there's anything you need. "Would you like a guided sleep meditation?" Luma asks. "Yes, please."

DISASTER ASSISTANCE

Disaster response is one of the clearest use cases for coordinated AI and robotics.

As with healthcare, the core capabilities already exist: AI for prediction and planning, sensors and imaging for early detection, robotics for mobility and reach, and increasingly capable networks of agents to scale up communication and personalized responses. The best disaster response starts long before the emergency notification hit devices — with systems that monitor ecosystem health, track infrastructure wear, and flag risks before they become crises. During an event, they can pair with human responders to coordinate the direction of resources. Afterward, they can assist with search, rescue, and repair. As AI fuels the next wave of robotics, we foresee better management across the full arc of natural and man-made disasters, helping communities prevent what they can, prepare for what they can't, and recover faster when disaster strikes.



IMAGINE

You hear the sirens before the storm reaches Houston. It's part of the plan. Days ago, the monitoring mesh flagged this system's trajectory, predicting flood zones block by block. Now, drones sweep low over the highway and soar into the sky for a bird's eye view, scanning for rising water and wind shear and pinging live readings to the joint response command center in ten second intervals. At 6:47am, your agent says it's time to go. In the car, you see the map shift as conditions change.

As you merge near Greenspoint, your agent takes over the car audio: Two passengers, about a quarter mile east. You tap yes. At the pickup point, Maria helps her father into your backseat. She's a nurse, headed to the same shelter. Convoys of autonomous shuttles glide past, guided by the mesh of drones monitoring the city.

At the shelter, bright orange "squatty body" robots unload insulin and infant formula while human responders focus on children and those with the most pressing needs. On the central wallscreen, blue dots pulse across Harris county, mapping the recovery efforts in real time.



The next day, when the sky clears, the coordinated response has already shifted from response to recovery. Tracked crawlers push through debris on your street, their sensors mapping damage as they clear drains and downed lines. In the bayou, amphibious drones test for contamination. Drones scan rooftops, capturing data for a coordinated repair queue based on level of damage.

The hive of activity on the screen, if anything, picks up. The system never sleeps. It just shifts, learns, and gets ready for next time.

Futures Practice

Museum Leaders Navigating AI Futures

As part of our company's philanthropic mission, Salesforce Futures brings futures thinking to the arts and culture world — helping institutions navigate the fast-changing realities of AI at a pivotal moment for the field. As one expression of this ongoing commitment, we partnered with Salesforce Nonprofit industry experts — Industry Advisor Valerie Brown, Strategic Solution Engineer Mike Brown, and Lead Technical Architect Ben Mozenter — to join a global cohort of museum leaders confronting the sector's biggest challenges.

At the second annual “Museums of Tomorrow Roundtable” (MTR), hosted with the Fine Arts Museums of San Francisco, the Futures team led an in-person session that blended a panel discussion, a live use case demo, and a candid interview, followed by an open exchange on AI's role in museums.

Rooted in a Futures-driven perspective, the session aimed to ground arts leaders in the realities of AI today while sparking their imagination for what's ahead. It offered space to wrestle with vital questions — how capable AI will become, how it will reshape creative work, and how it can strengthen community relationships. Attendees dug into the practical side too, tackling issues like data fragmentation and exploring ways to build a unified data foundation.

By connecting technological innovation with real-world application, work like this helps cultural institutions grow their capacity, deepen audience engagement, and prepare — thoughtfully and strategically — for an AI-driven future.

We turned to our industry leaders behind Salesforce's Arts & Culture Center of Excellence, Valerie Brown and Mike Brown to hear their reflections on the session, and how they are looking to the future of the arts sector.

“

The Fine Arts Museums of San Francisco (FAMSF) is leaning into the opportunities of the AI era by acting as a global convener at the intersection of art, culture, and technology. Through initiatives like the Museums of Tomorrow Roundtable (MTR), also actively facilitating cross-sector dialogue between artists, technologists, and international museum leaders.

Tom Campbell

Director, CEO

Fine Arts Museums of San Francisco



Q. What impact do you see the Museums of the Future Roundtable (MFR) program has on the arts?

Valerie Brown

I view this brilliant and innovative program as creating an international forum for museum leaders, fostering a vital bridge between technology and the cultural sector. The arts sector plays a significant economic role both globally and in the U.S. In 2023, arts and cultural industries contributed \$1.17 trillion to the U.S. economy, accounting for 4.2% of the nation's GDP. Globally, the arts recorded sales of approximately \$57.5 billion in 2024, underscoring the arts' substantial contributions to the global economy. This is an industry that is essential, yet often not treated as such. Considering its economic impact and the value it brings to our communities, we must support the arts in keeping up with the pace of innovation as nonprofit institutions face limited resources. We have a unique role in strengthening their capacity, and I see AI as playing a critical role in enhancing infrastructure and advancing audience engagement.

Q. What is your view on how arts and culture organizations can overcome barriers to transform with AI?

Mike Brown

The primary hurdle for arts and culture organizations is that their core systems — for ticketing, membership, and fundraising — are often siloed legacy point solutions, resulting in massive data fragmentation. AI needs structure, and many organizations lack the trusted data foundation required to generate reliable insights. Therefore, the first step is the adoption of modern solutions that support a “best-of-breed” approach, where the central CRM platform (Salesforce) acts as the unifying data layer. This platform must break down those silos, providing a single Comprehensive Stakeholder Profile for every patron.

“

I am proud to bring industry expertise to the table that provides value to both our customers and our business. In the age of AI, it is critical that we can come to the table at the start of our conversations with a clear understanding of the business. This knowledge helps us serve as a guide for our customers in unlocking the power of the Salesforce platform and AI.

Valerie Brown

Salesforce Non-Profit Industry Advisor



Our vision is for the Center of Excellence (CoE) to be the central engine that transforms the sector's complex requirements, particularly around data fragmentation, into guaranteed customer success. The CoE brings together relevant industry knowledge into a centralized resource for our solutions teams, instantly improving how we articulate value. Crucially, the CoE creates additional rigor around specialized implementation partners, paving the way for more seamless integration and delivery.

Mike Brown

Strategic Solution Engineer

Crucially, this adoption must be paired with the formation of an org-wide data strategy. Leaders must move beyond manual, “gut-feel” decision-making and clearly define how aggregated data will surface key insights to meet strategic goals, such as increasing member conversion or decreasing attrition. Finally, intentional data and system governance is the third core tenant. This means establishing clear commitment and communication around the ethical use of AI and data, ensuring data quality, and positioning AI as a tool to support employees, not replace them. By tackling data fragmentation through modern solutions and enforcing strict data governance, arts and cultural organizations can successfully cultivate an AI-ready culture.

Valerie Brown

In my four years at Salesforce, working closely with executives responsible for driving digital transformation, I have learned that adoption challenges are best addressed through a comprehensive change management process. Some of the best results I have seen while working with hundreds of nonprofits come from hosting interactive workshops that bring teams together to confront their fears about AI. These workshops are a free resource that we provide to our customers. By creating space to openly discuss concerns and educating teams on how AI works and the business challenges it can address, organizations can better identify where AI will have the greatest impact. Beginning with active listening and a proof of concept that is prioritized and informed by listening can be powerful. Mike and I share a passion for ensuring that the sector is prepared for an inevitable AI future.

The Robotics Breakout Moment



Emily Zhao
Salesforce
Ventures



Pascha Hao
Salesforce
Ventures



Why now is finally the right time for robotics.

We're entering a new era in robotics. Advances in foundation models and transformer-based AI — along with dramatic improvements in hardware cost and performance — are pushing the field beyond narrow, task-specific automation toward something far more ambitious: generalizable robotic intelligence.

The market is taking notice.

Capital deployed to robotics companies eclipsed \$7 billion in 2024, highlighted by mega-rounds in companies like Figure (\$675M Series B), Physical Intelligence (\$400M Series A), and Skild (\$300M Series A).

The global robotics market is forecasted to grow exponentially over the next 5 years as robots achieve general-purpose capabilities and expand from industrial settings into service sectors and consumer homes.

Here's how that growth breaks down:

Industrial Robotics: Valued at \$34 billion in 2024 and projected to grow by 10% annually over the next decade as automation deepens in manufacturing, warehouse logistics, and construction. It's worth noting that some analysts believe this market is much larger today and that growth will be even more aggressive, with projections as high as \$228B by 2029.

Service Robotics: Estimated at \$42 billion in 2023 and expected to grow by 16% annually through the end of the decade, driven by adoption in healthcare, hospitality, security, and logistics & fulfillment. This is projected to be the fastest growing category in robotics.

Consumer Robotics: A smaller but rapidly growing segment — ~\$14 billion in 2025 — is forecast to grow by 15.5% annually through 2030 as capable, affordable home robots gain traction. A majority of the adoption today is in home cleaning robots, personal companion robots, and robotic toys. By 2030, we anticipate we'll start seeing humanoids or new form factors used by consumers.

Generalist robotic systems could blur these category lines by enabling robots that can be applied to multiple use cases and verticals, in turn expanding the market even further.

We foresee a thriving robotics ecosystem emerging, and we're excited to nurture this growth by investing in teams developing robotics foundation models (RFM), full-stack hardware/software solutions, and the broader ecosystem of robotics tooling (e.g., robotics training data providers, simulation platforms).

We dove deep into the robotics market over the last seven months, speaking with dozens of the brightest innovators and experts working in the field today. Through these conversations, we developed a framework for investing in robotics grounded in where the field has been and where it's heading.

We foresee a thriving robotics ecosystem emerging, and we're excited to nurture this growth by investing in teams developing robotics foundation models (RFM), full-stack hardware/software solutions, and the broader ecosystem of robotics tooling (e.g., robotics training data providers, simulation platforms).

Continue reading, full article at salesforceventures.com





REFERENCES AND INSPIRATION

Articles, Academic Papers, and Company Reports

“The Second Economy,” by W. Brian Arthur, McKinsey Quarterly (10/01/2011): In this 2011 essay, Arthur anticipates the rise of an “autonomous economy” where digital processes interact with other digital processes, a concept that informs the future of multi-agent economic environments.

“To flinch or not to flinch: The role of the human participant’s behavior in human-robot collisions,” by Jeanne Kirschner et al., Acta Psychologica (07/01/2022): This 2022 study found that unpredictable human movements, like flinching, can increase collision risk with robots, highlighting the complex psychological and physiological challenges in human-robot interaction.

Measuring AI Ability to Complete Long Tasks, by Thomas Kwa, Ben West, et al (METR, 03/28/2025): A paper arguing that the ability of AI agents to execute long-duration tasks is doubling every seven months, with gains that will eventually transfer to complex personal AI tasks like trip-planning and budgeting.

“Expressing Stigma and Inappropriate Responses Prevents LLMs From Safely Replacing Mental Health Providers,” by Jared Moore, Declan Grabb et al: This study demonstrates how LLMs can respond inappropriately to users with mental health conditions and may encourage delusional thinking, showing the gap between simulated and real emotional intelligence.

Top 100 Gen AI Consume Apps Report, 5th Edition, by Andreessen Horowitz (a16z, 08/27/2025): This report is cited as evidence for the rise of AI companionship apps, which now make up five of the top 50 consumer AI applications, fulfilling previously unmet needs for connection and support.

How People Use ChatGPT, 5th Edition, by Aaron Chetterji, Tom Cunningham, et al, OpenAI, Duke University, Harvard University (09/15/2025): This report identified writing, information-seeking, and practical guidance (like tutoring) as the top use cases for ChatGPT, indicating a major shift in user behavior toward relying on AI for personal empowerment tasks.

Report on AI Interaction, by Pew Research Center: This report indicates that one-third of Americans who interact with AI do so “almost constantly,” highlighting the rapid integration and habitual use of AI in daily life.

BOOKS

The Art of the Long View: Planning for the Future in an Uncertain World, by Peter Schwartz (1991): A foundational text on scenario planning by Salesforce’s Chief Futurist. The cover of this book served as the direct inspiration for the cover art of the Futures Magazine issue.

The Innovator’s Dilemma, by Clayton Christensen (1997): Christensen’s work on “unmet needs” and disruptive innovation is used as a framework for understanding how personal AI can succeed by meeting latent, previously unrecognized needs with “good enough” solutions, such as AI companionship.

The Re-Shuffle, by Sangeet Paul Choudary: Choudary’s work is cited to support the idea that users will likely prefer a single primary AI agent rather than managing an “army of agents,” influencing the market dynamics between large incumbents and specialized startups.

Thinking, Fast and Slow, by Daniel Kahneman (2011): Kahneman’s concept of “System 2” thinking — the slower, more deliberate side of reasoning — is used to describe a key barrier that AI agents still face in achieving deep reasoning for complex tasks where trade-offs and human correction are required.

REFERENCES AND INSPIRATION

New companies and platforms directly or indirectly referenced:

Anthropic: Developer of the AI assistant Claude and the “Computer Use” model, which improves an agent’s ability to plan and execute actions. The company is also noted for its privacy-focused features.

Humane: Creator of the AI Pin, mentioned as an early new hardware device that failed to find its footing but drove conversation about moving beyond the smartphone.

Koah Labs: A startup aiming to create contextual relevance within AI-generated answers, helping advertisers reach “high-intent queries” inside human-to-AI chat exchanges.

LangChain: An open-source framework that makes it easier for developers to build the kind of specialized, plug-in agents needed for a thriving multi-agent ecosystem.

Limitless: A company developing an AI pendant that promises to deepen personalization by recording contextual user data.

Manus: A product that improves the ability of agents to plan and execute independent actions, approaching “System 2-style reasoning” but still requiring human correction.

Obsidian: A note-taking app that early adopters are linking with AI tools like Claude to create persistent, more contextually aware personal assistants.

OpenAI: Developer of ChatGPT and the “Operator” model, which is designed to improve the ability of agents to take independent actions. The company also conducts research on user behavior and engagement.

Perplexity: An AI-powered search tool mentioned as an upstart challenging established players in the personal AI market with splashy innovations.

Plaud: Creator of the “Note,” a hardware device focused on AI-powered transcription and summarization that continues the conversation around new AI hardware.

Rabbit: Creator of the R1 device, another early hardware offering designed as a new, agent-based interface that ultimately failed to gain market traction.



ACKNOWLEDGEMENTS

HUMAN

Salesforce Futures is indebted to our Salesforce colleagues and network of expert individuals for their direct and indirect contributions to the conversations that inspired and shaped this work.

AI

Salesforce Futures continues to experiment with plenty of AI tools in the production of this magazine. In our initial drafts, while exploring concepts for this magazine, we used Google Gemini and Adobe Firefly to produce imagery together with ChatGPT to brainstorm, copy-edit and to help format the faux ads. The final production includes custom graphics pulled from Adobe Stock (including stock AI generated), and Adobe Firefly (custom AI generated), and Google Gemini's Nano Banana.

In our approach to the final visual design process in this issue, we chose to thread the Salesforce brand elements for Dreamforce 2025 together under the theme of “coordination” with intentional design choices that depict shapes and forms and imagery interconnected while leaving aspects of sensory (sound, voice, and touch) to the imagination as we enter into a world where AI becomes more invisible.

DISCLAIMER

This magazine’s contents are intended for informational purposes. Please do not rely on this information in making your purchasing decisions. The development, release, and timing of any products, features or functionality not currently available remain at the sole discretion of Salesforce and are subject to change. Nothing herein is or should be construed as investment, legal or tax advice, a recommendation of any kind or an offer to sell or a solicitation of an offer to buy any security. This presentation does not purport to be complete on any topic addressed. Certain information in this presentation has been obtained from third party sources and, although believed to be reliable, has not been independently verified and its accuracy or completeness cannot be ensured.



About Salesforce Futures

We help Salesforce and our customers anticipate, imagine, and shape the future, building the shared understanding required to tackle adaptive challenges.

futures@salesforce.com

Access this issue at
sfdc.co/futures25

