




Futures

ISSUE 3 2024

Inventing with AI Agents





Welcome

By Salesforce Futures



The real voyage of discovery consists not in seeing new sights, but in looking with new eyes..

MARCEL PROUST

In this year's Dreamforce keynote, Salesforce's co-founder and CEO Marc Benioff showed why autonomous agents are the third wave of AI. All who attended Dreamforce learned about agents in our own way, filtering our observations through the lens of both our professional aspirations and more personal ones tied to our families, communities, and society at large.

Agentforce feels like a future that, to borrow from sci-fi writer William Gibson, is here but has yet to be distributed evenly. That won't be the case for long if generative AI's rapid, unprecedented adoption paves the way for a deeper integration of agents into the enterprise.


In the first two issues of Salesforce Futures Magazine, we flagged emerging, agent-related uncertainties about the pace of adoption, the evolving relationship between humans and AI, changing customer expectations, and the impact of these forthcoming shifts on business models, competitiveness, and the organization of work. It was honestly a little mind-blowing to head to Dreamforce and see so many people having conversations about similar topics so soon. So what did we learn?

First, Dreamforce sharpened our awareness of the ways in which the future is being invented, by Salesforce and others, right before our eyes. Agentic innovations change how Salesforce helps companies connect with their customers. In our second article, we imagine the capabilities of a CRM system and how it will change customer experience.

Second, the arrival of agents requires us to think further about how to amplify their benefits and mitigate potential harms. At Dreamforce, we noted concern about job loss and automation. Included here is a piece on how we think about this issue and the questions we need to ask to come together on this issue.

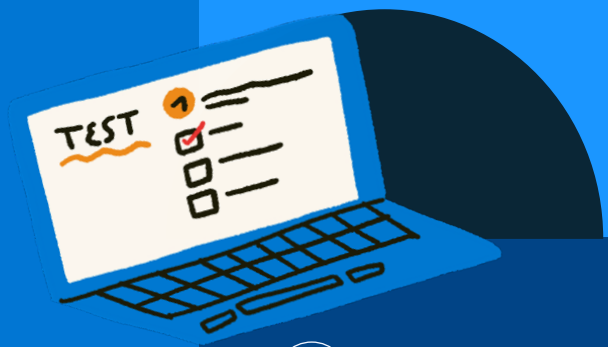
Finally, Dreamforce and the discussion of focused, task-specific agents catalyzed a conversation inside Salesforce about what happens when these agents scale. When everyone—agents included—can easily build and deploy agents, what does that world even look like? In our closing section, we look a little further out in an effort to anticipate the conversations we'll have next.

Organizational learning and leadership expert Adam Grant notes that “the future belongs to those who can connect the dots.” In this issue, our team endeavors to synthesize what we learned at Dreamforce, call out patterns, and connect dots for all of our readers during this time of great change. Please drop us a line if you have thoughts, questions, or additional ideas.



04

Inventing the Future of Agents



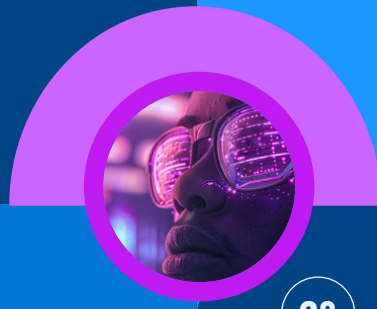
08

The Future of Customer Relationships: Agents and Aspiration

16

Overcoming the Imagination Gap:

How We Can Shape the Jobs of the Future with AI Agents



24

Looking ahead

28

Imagine if...



31 Credits

32 Acknowledgements



Inventing the Future of Agents

Computing has always moved forward in a kind of jerky, stop-start motion – what biologists call a punctuated equilibrium. New technological capabilities emerge gradually from research labs and garage tinkerers until, eventually, it feels like something new is in the offing. One or more pioneering feats of invention follow, shifting the axis of possibility. Suddenly, we find ourselves in a new paradigm and standing witness to platforms that unleash waves of innovation. Examples include the Apple Macintosh, the iPhone, and our own Salesforce Platform, which advanced an enterprise software as a service (SaaS) model that spawned an entire industry.



The best way to predict the future is to invent it.

ALAN KAY
COMPUTER SCIENCE PIONEER

Alan Kay was one such pioneer. In 1971, when Kay began his work at Xerox PARC, the computers he worked with were as big as a room, and only four of them were connected to the emerging ARPAnet, a precursor to the Internet. Kay, a gifted musician, brought a humanistic sensibility to innovation that helped him invent the Dynabook, a daring vision for personal computing, in 1972. The Dynabook was a battery-powered laptop with a touchscreen and wireless access to all the world's information and, most radically, was suitable for use even by children. Kay and his team at PARC went on to invent many core building blocks



of modern personal computing, including overlapping windows, the graphical user interface, and dynamic object-oriented programming. Later, at Apple, Kay also helped develop the vision for the iconic 1987 Apple Knowledge Navigator video, a clear antecedent of the iPad and the iPhone.



Agentforce brings the future to the present

The symbiotic dance between visioning work and the development of raw technological possibilities came to mind as Salesforce Futures reflected on the launch of Agentforce at Dreamforce 2024. A year earlier, for Dreamforce 2023, we drew explicit inspiration from the Apple Knowledge Navigator in creating our Salesforce 2030 film, which showcased a vision of humans collaborating fluidly with sophisticated, autonomous AI agents. More recently, our magazine issues featuring Personal AI Agents and Agents at Work bridged the imagination gap across the emerging capabilities of agents and how they might change business as we know it.

Agentforce, a suite of customizable agents and tools on the Salesforce Platform, offers an elegant solution to the complexity involved in trying to deploy AI tools. It targets the challenges businesses face with integrating different components of the AI stack – data, models, infrastructure, and applications into a unified and functional system.



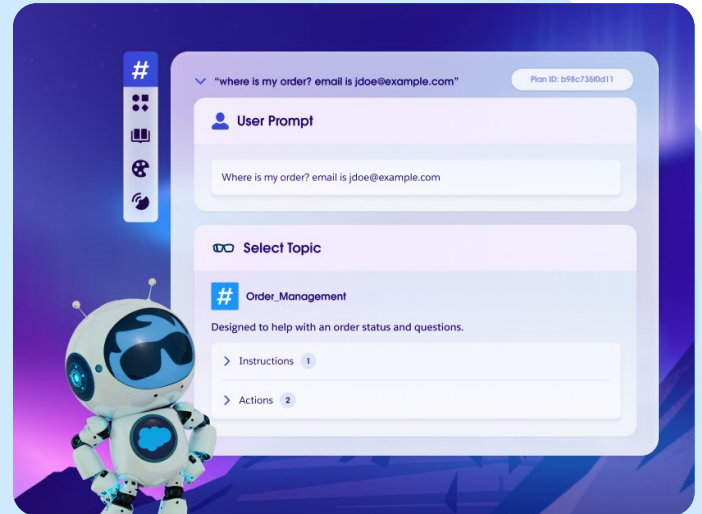
It also takes advantage of key capabilities in the Salesforce Platform, like metadata, permissions, and security. Tools like Agent Builder and Model Builder, meantime, enable organizations to easily create, customize, and deploy AI agents while Salesforce's Atlas Reasoning Engine enables agents to autonomously handle both routine and complex tasks.

“The end of the beginning” for a new AI age



Agentforce is an incredible breakthrough for Salesforce and the business world, which is why more than 10,000 people at Dreamforce 2024 raced to build their own agents. Customers walked into the “Agent Builder” experience at the show and turned verbal instructions into functioning agents in less than 15 minutes. For both our team and the customers we interacted with, it felt like the future was being invented and pulled into the present right in front of our eyes.

These aren't just another version of chatbots. They are a new breed of AI that could reshape how businesses operate and deliver superior customer experiences, leading customers like Saks, OpenTable, and Wiley to quickly embrace the technology. Kevin Quigley, Director of Continuous Improvement at Wiley, put it this way: “It's been exciting to go live with our first agent thanks to



the no-code builder, and we've seen a more than 40% increase in case resolution, outperforming our old bot.”

When we see signals like this in the present, it pulls our attention toward the plausible potential of the future. Soon, we anticipate more complex, multi-agent orchestrations solving higher-order challenges across the enterprise, like simulating new product launches or marketing campaigns.

On the consumer side, we foresee agentic capabilities finding their way into consumer devices and apps. These developments promise to shift both how work is done and how companies connect with their customers.



From imagining to shaping the future

Speculation will only take us so far. During the next year, with Agentforce generally available, powerful, agentic tools will be in the hands of hundreds of thousands of Salesforce customers around the world. We think the ability of these customers to express their business needs for agents will reveal use cases and applications we could have never imagined. Agentforce isn't just a product; it's a platform for experimentation.

Like every major technological shift before it, including the ones Alan Kay and his team at PARC had a hand in, the real magic won't be in what we've built – it will be in what people do with it. Inventing the future means acknowledging that the future is influenced by every interaction, every agent deployed, and every problem solved in new, unexpected ways. For Salesforce Futures, this signals the need for a shift from anticipating and imagining the future to diving in and shaping it.



The Future of Customer Relationships: Agents and Aspiration

Tomorrow's customers will have much higher expectations than today's. In Dreamforce stories like [this one](#) for Saks, we saw how AI agents could revolutionize customer experience, erase some of today's most frustrating pain points, and turn every service interaction into a way to wow customers and build better relationships through fast, frictionless, and personalized interactions.

Now imagine what happens when these kinds of hyper-automated, agent-first encounters occur not only in service and sales, but across all departments, functions, and customer-facing touchpoints. And what about when customers all have their own agents?

Given where technology is heading, our team believes there are a few plausible ways agents might change relationships between customers and companies.

First, agents will help companies personalize experiences for their customers with a scale and speed far greater than today.

Second, agents will help companies build trust by adapting to changing customer preferences and aspirations over time

Finally, human-to-human interactions, informed by next-gen contextual intelligence, will serve as a critical component of relationship building

Seeing what customers can already do with Agentforce—a suite of customizable agents and tools seamlessly integrated with your data and apps—inspired Salesforce Futures to imagine what an elevated customer experience might look like in the next three to five years.*

To bring this to life, we intentionally picked moving into a new community—a transition that strikes a special dread in the hearts of even those of us schooled in the adulting arts. Could agents really help make moving a much better experience in the near future? Let's jump into a fictional scenario before we unpack (no pun intended) what's happening behind the scenes.

**This is an exercise in speculative imagination and not a preview of the Salesforce product roadmap.*



First Year in a New Home

Juan and his daughters move to Dahlia Gardens.

Juan, his daughters, and their new beagle, Beni, are moving into their new home in Dahlia Gardens: A close-knit community focused on sustainability and resident experience. Juan has already given his

personal agent permission to share his family preferences around everything from lighting to study routines. Dahlia's own agents use this data to customize his new home.



As Juan and his daughters settle in, Juan starts getting gentle nudges from his agent about Dahlia's neighborhood meetups, ride-sharing opportunities, and other resource-sharing options like the community tool program and gardening co-op.

A few weeks later, Jasmine, a Dahlia Gardens customer engagement specialist, schedules a check-in with Juan to see how they're doing. Juan shares his experience in the community—information Jasmine uses to further personalize the recommendations from Dahlia's agents. For example, when Juan notes how much he likes the community resource-sharing programs, Jasmine encourages Juan to brainstorm additional ideas. Juan expresses enthusiasm, and Jasmine and Juan let their agents find a time to follow up.

A lot of the recommendations to residents like Juan are automated, and most routine questions are answered by interactions between residents' personal agents and the Dahlia's customer engagement agent.



Because Jasmine spends less time answering basic questions and scheduling meetings, she has a lot more time to connect with residents one-on-one. The direct, spontaneous feedback she gets from these conversations also helps her improve community programs in a way that's a lot more nuanced than the kind of periodic satisfaction surveys we see today.

Juan's agent reminds him from time to time about the action item to brainstorm a resource-sharing

program. One day, Juan is rushing to walk Beni before work—in any future, your kids can still overcome any objections you might have and convince you of the merits of pet ownership—when

an idea comes to him. What if there was a pet-care service where residents could volunteer to pitch in on pet care in exchange for occasional help with their own pets?

At their scheduled follow-up appointment, Jasmine and Juan start trading ideas about this potential community pet-care program.

Jasmine uses her community knowledge to help Juan build on the idea. After Jasmine and Juan refine the concept, their agents help them sharpen the description of the program, design a pilot, dynamically schedule a meeting with interested volunteers, and even start on a simple way to track who needs help and available volunteers.

The community pet-care sharing service turns out to be a hit with residents. Soon, Juan is no longer just the new guy with two daughters and an unruly beagle: He's the community innovator who helped co-create and launch a service that helps everyone with pets. To find new pet care customers, Juan and his daughters participate in more community events. The more they engage, the more Dahlia Gardens feels like home.



Unpacking this Future

What's happening behind the scenes to make this story possible?





Unpacking this Future

Agent to agent interactions

CUSTOMER RELATIONSHIPS



In this world, agent-to-agent interactions occur regularly between agents representing Juan and those representing Dahlia Gardens. These agents anticipate Juan and his daughters' needs, coordinate services, and work with Juan to arrange the family's calendar. Agents, making proxy decisions for humans, suggest community events and ways Juan might engage. Thanks to personalized recommendations and dynamic scheduling, they also improve the likelihood of suggesting ideas that align to Juan and his daughters' interests and availability.

Imagining this environment catalyzes considerations about system-level architecture requirements. Any enterprise system of the future will need to scale to accommodate the massive volume and variety of requests from agents representing human interests and the subsequent exchanges that follow.

In a localized community like Dahlia Gardens, these requests might be relatively finite. In other contexts, such as a more global commerce environment where companies sell to customers all over the world, the volume would be far greater. Coordinating agents in a complete AI system will be increasingly important in the coming years.



Personal AI agents, with better access to emotional and behavioral context, will learn and adapt to their users' preferences over time. As a result, customer loyalty with brands will be heavily influenced by the "preferences" of people's personal AI agents. Successful brands will offer trusted, hyper-personalized messaging, services, and products while adapting to the evolving needs of their customers, and they'll package information specifically for agents.

For customers like Juan, the boundaries between sales, service, marketing, and commerce interactions will disappear. Today, Juan might interact with a sales department while researching Dahlia Gardens, or engage with a service specialist when he has an issue with his home that requires attention.

Tomorrow, generalists like Jasmine, who possess strong interpersonal skills and relational intelligence, will use agents to offload mundane tasks and develop more personal relationships with the customers they serve.

Deeper, more nuanced relationships



Unpacking this Future

Prioritizing human-to-human connection

CUSTOMER RELATIONSHIPS

In the scenes with Jasmine, we see how she utilizes agents and customer data to enable the best possible experience for Juan and his daughters. Jasmine uses agents to orchestrate both agentic and human interactions with customers, giving us a preview of how human-at-the-helm might operate at scale.

In trying to imagine the future, we often find it useful to consider signals from the present that show us how things might work. Financial markets, for example, offer a preview of how the relationship between humans and machines could evolve. Many aspects of financial services are commoditized and transactional, such as trading and products. Even though most of these services and products are delivered with the help of automation and scaled to serve many customers simultaneously, trusted relationships between humans in financial services are still critical—especially for navigating tricky, nonroutine escalations or creating new, breakthrough offerings.

Fluid interfaces that adapt to the task at hand and generative “UI on the fly” are central to conversations about an agentic future. When Juan and Jasmine begin brainstorming the pet-care idea, they don’t need to open up new tabs to start building out the program—their agents automatically begin fleshing out both the concept and a plan to implement it as the conversation proceeds. Interface designs like OpenAI’s Canvas feature for writing and coding and the new artifact feature in Claude already show us how AI and humans might collaborate on a task with less swiveling between tabs and tools.

Fluid Interfaces





Unpacking this Future

Imagining positive futures with agents

CUSTOMER RELATIONSHIPS



Moves to a new community will never be completely frictionless. Things will still break and beagles will still bark. But in the Juan story, we see one way the nuanced, contextual intelligence companies gather from agents and humans alike might enable an organization like Dahlia Gardens to develop deeper relationships with the customers they serve. Dahlia uses what it learns about Juan to personalize his experience and accelerate his connection to his new community over time.

Imagining how the relationships evolve between companies like Dahlia, employees like Jasmine, and customers like Juan and his daughters is a useful way to close the imagination gap. It helps us center human experience even as we embrace exciting new technological possibilities. In a world of abundant agents, speed and convenience will be table stakes. What will really set companies apart from their competitors is the ability to adapt to changing customer needs and aspirations over time.






Overcoming the Imagination Gap: How We Can Shape the Jobs of the Future with AI Agents



The unknown is scary until you know it.

MARC BENIOFF
DREAMFORCE 2024



Thinking productively about the future is fundamentally an act of creative imagination.

However, thinking about AI and job displacement is often distorted by an imagination warp and an imagination gap.

The imagination warp is caused mainly by the powerful attractive force of dystopian narratives originating in sci-fi: Artificial general intelligence (AGI) narratives proposed by AI research labs also contribute by misleading us as to where we are now and what's around the corner. The imagination gap refers to the lack of a shared framework for understanding how technology changed work in the past that can guide our understanding of

how innovations like AI could alter our lives, both positively and negatively, in the future. We saw this warp and gap in conversations with media and analysts at Dreamforce 2024, leading to questions and concerns about where AI might take us.

We can overcome most reservations about this technology if we come together as stakeholders in a shared future, gaining a clearer view of what we can learn from the past, giving workers the skills they'll need for tomorrow's jobs, and experimenting to learn how humans and AI can work best together. We believe this effort is critical to ensure humans still possess agency – or the ability to control their own fate as their interests and values change—even in a world filled with autonomous AI agents.



The imagination warp: Existing narratives + AGI are failing us

Powerful dystopian sci-fi narratives provide easy, widely-understood references for the media. However, they also reinforce the idea that bad outcomes are not only possible but likely. In enduring Hollywood franchises like Terminator and The Matrix, machines use their superior intelligence to rob humans of their agency. This dynamic is also common in the more current AI doomer scenarios that occasionally grab headlines in the mainstream press by discussing the possibility of “civilization destruction” or “AI apocalypse.” Ironically, even the more utopian-leaning AGI predictions from language model developers like Sam Altman and Dario Amodei warn that bad outcomes are likely if we don’t put checks and balances in place to control the technology.

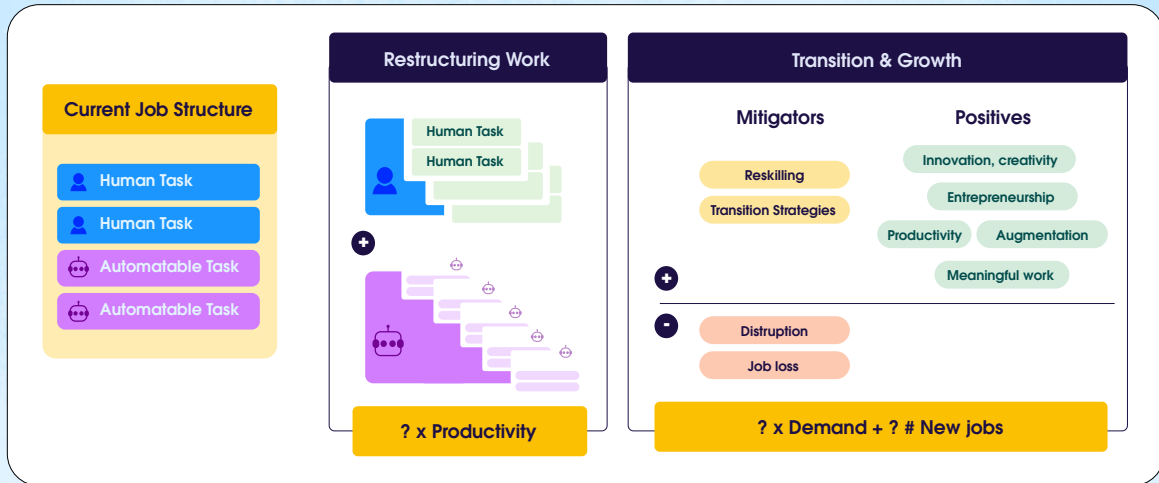
Such forecasts lead to apprehensions that any related innovations, such as AI agents, could be precursors to an AI that is so powerful and pervasive that it will make human workers obsolete. AI isn’t the first automated technology to stir fears about job loss. From the assembly lines that replaced skilled laborers in the early 20th century to ATMs that threatened bank tellers and automated switchboards that phased out operators, automation has a long history of making workers worry about their future. But in our futures work, as in life, we’ve found that negative scenarios are much easier to create than positive ones. They only require envisaging existing

systems breaking down. “All you have to do to develop negative scenarios is take reality as we know it and kick it to pieces,” writes futurist Jay Ogilvy. “Positive, optimistic scenarios, on the other hand, lack plausibility if they sound like a view of the future as seen through rose-tinted glasses.”

To be believable, positive scenarios must demonstrate how existing problems that are currently holding us back get solved.

This is challenging: If solutions were readily apparent, they likely would’ve been implemented. Part of the challenge of futures work, therefore, is to develop a deeper understanding of how complex systems might transform and pair that knowledge with a willingness to imagine novel and unexpected solutions.

If we can escape the pessimistic narratives that limit our imagination, how might we more clearly think through the effect of AI on employment? Breaking down what might happen to how we work requires a mental model shift from jobs to tasks, some of which are automatable, and many of which are not. The White House’s National AI Advisory Committee has highlighted that economists’ predictions of which tasks are most susceptible to automation from AI have changed over time, from an emphasis on manual work in the 2010s to the current focus on “intellectual and unregulated” work.



Closing the imagination gap towards a better framework for thinking clearly about AI and jobs

The key questions to ask are: which tasks will truly be automated, how that automation will reshape work, and which tasks will prove resistant to automation?

Even if AGI arrives sooner than expected and all tasks prove automatable, there are questions that remain about humans and work. First, we need to accept that there will still be hallucinations and other issues with AI to contend with. We need to experiment and learn what AI is best at and what humans are best at.

Second, even in circumstances where AI is objectively “better”, there may be scenarios where we prefer humans and human-shaped processes. Here, we think about aspects of human relationships and identity like trust, human connection, and what author Vaughn Tan calls “meaningmaking.” Framed this way, we can see how much meaningmaking happens in work—when we make tough, values-driven leadership decisions, make subjective judgments, or assess risk—and how central that work is to both personal and community identity.

If we want real conversations about better AI futures, we must understand more deeply the nuances of the human side of work.



Upsides to automating what can be automated

One consequence of the imagination gap is that we fail to recognize that the immediate consequences of automating some tasks may actually be positive for net employment.

First, in many situations, the existing supply of offerings and services is labor-constrained beyond the demand that actually exists today. We experience this when our customer service call places us in a 3-hour queue, relieved somewhat (but not really) by having a callback option. Using agents to solve for these kinds of supply gaps wouldn't displace current jobs, and service levels would improve.

Salesforce CEO Marc Benioff offers another example: Demand for access to medical advice way outstrips the 15 minutes a harried medic can give any one patient. Tight labor markets are already pervasive today, and with declining birthrates and growing resistance to immigration, MIT economist and labor expert David Autor suggests, "the industrialized world is awash in jobs, and it's going to stay that way." Another upside uncertainty relates to situations where demand is elastic to supply (i.e., as supply increases, demand will too).

Again, history helps here. Economist William Jevons first predicted this paradox with demand for coal in the 19th century, but more recently, the arrival of spreadsheets on PCs presaged an explosion in demand for the kind of quantitative analysis they enabled.

Today, AI tools are already showing promising results in increasing developer productivity. Looking forward, venture capitalists Paul Kedrosky and Eric Norlin suggest increased

developer productivity can help society escape the high costs of software, unleashing a much higher level of true demand, and solving for what they call "society's technical debt." AI can also open up coding jobs to a far broader swath of people by eliminating technical barriers, says AI pioneer Andrew Ng.



What makes us human is our ability to do things which are not-yet-understood, which requires us to be able to create meaning where there wasn't meaning before...

VAUGHN TAN
AUTHOR





Reskilling and transition strategies

Where jobs are displaced, recent economic history offers some caution. In the context of late 20th century deindustrialization, the promises of reskilling factory workers for the knowledge economy too often fell short, leading to not only the loss of millions of jobs but also persistent long-term unemployment in post-industrial areas. Beyond the exaggerated narratives of science fiction and AGI, it is these experiences that weigh heaviest on the minds of workers and policymakers alike. Thankfully, multiple reasons exist to expect more from the AI revolution.



First, the gap between the skills required for the knowledge jobs of today and the knowledge work of the future is smaller than it was in the transition from the manual, semi-skilled roles that were lost to industrial automation.

Second, AI promises us much better tools. We're increasingly capable of mapping the skills workers have today, understanding what skills they might need in the future, and plotting a personalized path from here to there. AI-powered coaching tools will also personalize learning and prove more accessible to the people who use them than past learning systems. In particular, we see enormous promise in newly-cost-effective, data-rich simulations that allow mutual testing of aptitudes and interests for emerging roles. Imagine if you could practice for a job interview as many times as you needed to before the actual event, rehearse for big meetings with much greater fidelity than today, or even run launch simulations of new ideas to refine go-to-market strategies before an expensive pilot phase.



Entrepreneurship and innovation

If past waves of innovation are a guide, AI agents may lead to a surge in entrepreneurial activity. The same AI-powered coaching and tools that help reskilling could also help people find and prepare for new entrepreneurial opportunities and provide cheaper access to expertise that could fuel an entrepreneurial surge.

In his fireside chat with Salesforce Chief Design Officer Kat Holmes, Cristóbal Valenzuela, CEO of AI video creation company Runway, pointed out how technological constraints have historically held back people who might want to tell stories but feel they lack the skills.

Already, we're seeing freelancers with unfettered access to resources, such as financial and legal expertise, that until recently would have been reserved for large companies.

We believe that a rising tide of accessibility—abundant intelligence, expertise, and support—will make it easier than ever before to start and run a business.

This also extends to more creative aspects of work. At Dreamforce, Cristóbal Valenzuela, CEO of the AI video creation company Runway, pointed out how technological constraints have historically held back people who might want to tell stories but lack the skills.

We'll also see the arrival of entirely new jobs. A study led by David Autor recently pointed out that 60% of the jobs Americans held in 2018 did not exist in 1940, concluding that new technologies create the need for new roles. AI looks to follow the same pattern. Imagine agent-builders and trainers, interpretability experts who audit decisions (and the models that produce them), and other new roles enabled by new opportunities that emerge over the next few years. The successful implementation of agents also necessitates both infrastructure and teams to monitor the security of agents at scale.





In the first issue of Salesforce Futures Magazine, focused on personal AI agents, we imagined futures where agents mediate most interactions between customers and companies. In a world where automation is abundant, the value of high-quality, human-to-human interactions should rise, especially for more novel circumstances where human attention is needed to help customers solve problems and navigate decisions. These human moments will be one of the most important ways for companies to build trust with their customers and differentiate themselves from competitors. It stands to reason future companies will invest more in roles focused on personal service, conflict resolution, and relationship building.


Actions and conversation

We believe AI can be designed and built in a way that benefits us all. Actions we announced at Dreamforce, like increasing AI literacy and providing access to AI resources, are obvious steps toward that goal.

Multi-stakeholder efforts to anticipate and protect against harms and to maximize opportunities are another tool, and Salesforce is proud to participate in the U.S. National AI Advisory Committee, National Institute of Standards and Technology, and Singapore's Advisory Council on the Ethical Use of AI and Data. The government has an important role to play in driving the successful adoption

The successful implementation of agents also necessitates both infrastructure and teams to monitor the security of agents at scale.

of AI, and designing and implementing programs that broaden access to these powerful technologies. For inspiration, we can look to concrete initiatives like ASU's \$34.6 million digital equity initiative to democratize access to technology, or the responsible AI partnership between the Block Center at Carnegie Mellon and the State of Pennsylvania. Further afield, Singapore's national AI strategy seeks to harness the power of AI for the public good by focusing on expanding the AI talent pool, accelerating adoption, upskilling the workforce, and designating public sites to nurture the AI community.



As Salesforce and other tech companies launch agents into the world, we have a responsibility to pay attention to the impact this new technology has on humans, workforces, and communities. We'll not only learn what agents are capable of and how humans and agents best collaborate; we'll also learn about agents' limitations. Widescale adoption of agents will reveal unintended consequences, second-order effects, and surprising use cases we may not have even imagined. Businesses must learn rapidly and share our insights far and wide, so that all stakeholders – policymakers and educational institutions included – can make better decisions about workforce impact.

We know AI will set off yet another wave of creative destruction.

We don't know how exactly this wave will crash or where the balance will fall between jobs created and lost, but the lessons of the past suggest we may well be surprised on the upside.

With clearer thinking and more engagement, we can move beyond both pessimistic proclamations and techno-optimistic boosterism, rebalance the conversation, and recover our agency.

This is the best way to a future of AI and humans working together.





LOOKING AHEAD



**Where are we going next?
We're going to a world where agents talk to agents.**

PETER SCHWARTZ
DREAMFORCE 2024

The agents we introduced at Dreamforce are already helping businesses make huge strides in automating tasks and unlocking the potential of digital labor.

This future is being invented before our eyes, and new systems and behaviors are sure to follow. As a futures team, we've been thinking about the broader implications of agents for Salesforce, our customers, and society at large. This work requires a future-back, outside-in orientation and a willingness to consider a range of both short-term and long-term possibilities.

When we do this, we find ourselves imagining a fast-approaching future where agents can control other agents, buy goods and services on behalf of humans, negotiate with one another, and even create new agents. Here are some of the big themes we're exploring as we look ahead to the near future.





A rise in personal agents and bring-your-own AI

New advances will place ever-more-powerful AI agents in the hands of consumers. Agentic capabilities in consumer devices and apps, like Apple Intelligence and China's super-apps that offer a one-stop shop for users, are already emerging and aim to leverage personal context to help users get things done effortlessly. As people begin to bring their personal AI agents to the workplace, companies will need to find ways to integrate AI into the enterprise more quickly. We call this trend BYOAI (bring your own AI agent), mirroring the earlier BYOD (bring your own device) movement that saw workers bringing personal smart phones to the office. And as they become mainstream, personal AI agents will forever change how companies connect to their customers. Successful companies will offer trusted, hyper-personalized messaging, services, and products at the moments when they matter most, and they'll package everything in easy-to-consume ways for the agents that increasingly serve as customer proxies.



New interfaces for a multi-agent world

Don Norman, the author of *The Design of Everyday Things*, defined an interface as “the place where human and machine meet.” Note that Norman’s definition imagines a single human interacting with a single machine. A world of agents interacting with other agents pushes us to imagine humans overseeing more complex networks of machines. This raises new questions about confidence and control. How will humans maintain a sense of ownership over outcomes? As conductors of symphonies of agents, what is the interface that best ensures we play the right music? How should companies redesign themselves to be agent-oriented in order to take advantage of the incredible affordances of this moment? If things go wrong, how will we interrogate the rules, hierarchies, and inferences in order to fix things?

Agents, with their new affordances, do necessitate new interfaces, but predicting how these interfaces evolve is notoriously challenging in absence of consumer feedback.

We anticipate a period of trial and error until workflows solidify and confidence builds. Think about Google’s ten blue links, the smartphone touchscreen, the scrollable social media feed, or the shopping cart in ecommerce. For each example, a need necessitated a new design that established





itself in our daily lives after a period of trial and error and discarded directions. UI metaphors will likely need to migrate towards time and threaded conversation based feeds as opposed to the app-dominated landscape today.

For each example, a need necessitated a new design that established itself in our daily lives after a period of trial and error and discarded directions.

Early ideas we've imagined include the presence of a chief of staff agent whose job it is to oversee all other agents, a concept that addresses the potentially overwhelming job (for humans) of monitoring machine armies.

More recently, Adam Evans, EVP & GM of Salesforce AI, began discussing the presence of the right introspection, dials, and controls to fine tune agents over time and ensure that humans are able to customize them to suit their needs—particularly when it comes to optimizing algorithms within the enterprise as strategy evolves.



LOOKING AHEAD



Gen Z and older generations drive AI adoption at work together

Younger generations historically serve as a signal of how adoption of new technologies might unfold. Millennial attitudes and uses of smartphones, for example, provided early insight into how transformative the combination of web-based apps, location services, social media, and ecommerce on mobile phones would become.

It follows to ask what's going on with Gen Z. Salesforce's AI Connected Customer research finds that Gen Z is more willing than older generations to use AI agents to shop for them (32%) and to create more personalized and useful ads and content (44%). And research from Slack Workforce Lab finds that 55% of workers ages 18 to 29 show enthusiasm for AI and automation to handle part of their work.

Gen Z's optimistic enthusiasm for AI across personal, work, and societal domains signals a future where AI tools become as integral to daily life as smartphones are now, with Gen Z leading the charge in adopting, normalizing, and shaping AI use in ways that will influence everything from workplace productivity to customer expectations.

But while younger users are surely establishing some new behaviors around things like in-app search and companions (see the following section) that doesn't mean more experienced workers should wait for younger workers to "show them how to use AI." Research cited by Ethan Mollick points out that managers may be best positioned to incorporate AI into their work because of their domain expertise and the speed with which AI is moving. As Mollick says, "The technology is advancing quite quickly and you need to use it to understand it."

Agents transform search

Agents are poised to disrupt the search market beyond recognition, further threatening the erosion of Google's dominance and presenting significant opportunities as consumer behaviors shift and new chokepoints emerge.

Today, search engines help users find information and discover content. Google dominates this space for the open Internet; Amazon's search engine helps users discover products; and enterprise search engines help employees find experts and documents inside companies. Sometimes, finding information is the user's main goal.

LOOKING ALICAD



More frequently, finding information is one of many tasks in service of accomplishing a larger goal such as when a consumer buys a product or a knowledge worker prepares a presentation.

As agents' capabilities advance from being able to precisely answer user questions to completing complex tasks, agents are poised to disrupt how brands interact with consumers and how employees collaborate to get things done. These shifts mean every brand and company seeking to connect to its customers will have to invest in new agent-compatible channels of communication to connect with their customers.



AI companions and the evolution of social norms

Strong adoption signals in the “AI Companion” space point the way toward a future where it's more likely that AI presents as human.

It should open our eyes when OpenAI includes a warning with their recent voice capabilities that humans may grow “emotionally reliant” on the feature.

We know that, even before the pandemic, there were unmet needs for companionship for the elderly, young people experiencing loneliness, and others. The pandemic exacerbated these needs, and we are seeing AI companions fill them.

There is cause to worry about overreliance, declining social connection, and other associated ills related to companion use, but this is also a moment to imagine new social norms and possibilities whereby companions help their users and even connect them to other humans with greater intelligence. In May, the founder of Bumble imagined AI concierges who could interact on behalf of their human proxies to enable better matches. We could extend this concept to networking, volunteering, and other domains and imagine plausible benefits.


Similar to recent conversations about job-related impact, we believe that, given the arrival of AI agents, imagining better futures related to social connection will be increasingly important work.

With more companies pivoting to consumer-facing AI products, it becomes increasingly necessary to ensure human-AI relationships benefit humans.



LOOKING AHEAD

Imagine if...



More competition from small players

The cost of deploying agents plummets, making them accessible to even the smallest businesses. Access to advanced AI capabilities is democratized, empowering small and medium enterprises to suddenly compete on a more level playing field with larger organizations.



Superapps go west

Apple, OpenAI, X, WeChat, Meta, Amazon and new entrants like You.com and Perplexity.ai race to take advantage of the new opportunity to serve as a superapp for all agentic customer interaction. Each service offers a conversational interface (including voice), AI search, creative tools for content generation, the ability to delegate personal tasks to agents (including transaction & payments), and much more.

Real time demand response

The ability to capture and analyze the vast amounts of data generated by agent-to-agent interactions, semantic twins, and simulation capabilities unlocks new business insights and opportunities, revolutionizing market research, trend forecasting, and product development by enabling businesses to anticipate customer needs and respond to market shifts instantaneously.





Agents develop their own language

Agents develop their own languages for agent-to-agent communication that prove more efficient and effective than human languages for certain tasks. While these languages may promise benefits for speed and precision, they could also increase opacity and raise ethical concerns about privacy, security, and human oversight.

Your customer is an AI agent

AI agents overtake websites, social media, and toll free phone numbers as the primary and default interface for connecting companies with customers. What is the new channel, and how does it change the relationship between companies and customers?



Personal agents elevate trailblazers

Personal AI agents become great equalizers by delivering access to personalized education, coaching, and even therapy, and catalyzing a new wave of super-empowered entrepreneurs, generating a similar impact to that of mobile phones a generation ago.





references and inspiration

FUTURES CURATORS

David Berthy, Mick Costigan, Marc Escobosa, Angela Gleason, Daniel Lim, Peter Schwartz

SALESFORCE ARTICLES CITED AND INSPIRATIONS

"Agentforce Is Here" Salesforce | [LINK](#)

Details the Agentforce platform's features and its role in simplifying AI integration for businesses.

"What Is a Reasoning Engine?" Shipra Gupta | [LINK](#)

An overview of Salesforce's Atlas Reasoning Engine and its role in enabling AI agents to autonomously handle complex tasks.

"What Makes a Complete AI System?" Patrick Stokes, Salesforce Blog | [LINK](#)

Discusses the essential components of a unified AI system and their integration to drive business innovation.

"Personal AI Agents and Agents at Work" Salesforce Futures Magazine | [LINK](#)

Our second issue of the magazine explores how personal AI agents and workplace agents are closing the imagination gap and reshaping the business landscape.

"Removing Barriers to Creativity with AI" Cristóbal Valenzuela | [LINK](#)

A fireside chat from Dreamforce 2024 about how AI tools can democratize storytelling and creativity.

"From 'There's an App for That' to 'There's an Agent for That': How AI Is Transforming Our Lives"

Adam Evans, Salesforce News | [LINK](#)

Explores how AI agents are reshaping daily life and transforming customer and business interactions.

EXTERNAL ARTICLES, BOOKS, AND PAPERS CITED

"Personal Computer for Children of All Ages" Alan Kay, Xerox PAR | [LINK](#)

Explores Alan Kay's vision for the Dynabook, a groundbreaking concept for personal computing suitable for both children and adults.

"The Role of Meaning-Making in Work" Vaughn Tan | [LINK](#)

Examines how humans create meaning in their work and why this is critical in the age of AI.

"Labor Markets and the Automation Paradox" David Autor | [LINK](#)

Explores how automation impacts labor markets and the economic dynamics it creates.

"Emergence, Story, and the Challenge of Positive Scenarios" Jay Ogilvy, World Futures | [LINK](#)

Explores the complexities of creating compelling and believable positive future scenarios through storytelling and emergence.



references and inspiration

"The AI Skills Gap" National AI Advisory Committee, White House | [LINK](#)

Highlights evolving predictions on which tasks are automatable and the importance of reskilling.

"AI in the Public Good" Singapore National AI Strategy | [LINK](#)

Details Singapore's initiatives to upskill the workforce and leverage AI for public benefit.

"The Design of Everyday Things: Revised and Expanded Edition" Don Norman, Basic Books, 2013 | [LINK](#)

Seminal book on design that offers updated insights into user-centered design principles.

Ethan Mollick, LinkedIn post cited from November, 2024 | [LINK](#)

Discusses executives' perspectives on AI adoption and its impact on business operations.

"Character.AI: The Companion Product Leading the Market" Andreessen Horowitz (a16z) | [LINK](#)

Highlights the rise of AI-driven companion tools and their rapid adoption among users.

"AI Concierges and Human Connection" Whitney Wolfe Herd, Founder of Bumble | [LINK](#)

Discusses AI concierges' potential to enhance personal connections and networking.



ACKNOWLEDGEMENTS

HUMAN ACKNOWLEDGEMENTS

Salesforce Futures is indebted to the following individuals for their direct and indirect contributions to this work: Max Kirby, Adam Evans, Silvio Savarese, David Autor, Left Field Labs, Vaughn Tan, Slack Workforce Lab, Emily Witt

AI ACKNOWLEDGEMENTS AND DESIGN CHOICES

The AI tools in the production of this magazine include: Adobe stock (including stock AI-generated), Adobe Firefly (custom AI-generated), Midjourney, and DALL-E for imagery; ChatGPTo was used to develop concepts, test headlines, and copy-edit.

The creative approach to this issue anchors on a hopeful future, pulling from playful visuals to juxtapose deeper reflections on both a personal and professional level as AI agents evolve their role in our day to day lives. Connecting to the Salesforce brand, we use a color palette, shapes, and graphics that underscore the tone of creativity and optimism. Our focus is on empowering the imagination of individuals and teams in a world where nuanced collaboration emerges between what is familiar, novice, and unexpected.

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 guaranteed.



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