



Futures

ISSUE 1 Q1 2024

Personal AI Agents



Welcome

By Salesforce Futures



In 1966, MIT computer scientist Joseph Weizenbaum created ELIZA, an early chatbot designed to facilitate the study of interactions between humans and machines.

By today's standards, the program was crude—mirroring its responses around a single keyword—but the reactions of users shocked Weizenbaum. Instead of seeing the shallow state of human-machine interactions as expected, users ascribed both intelligence and agency to the program. In other words, they easily believed that there was a thinking intelligence behind the screen. Weizenbaum's own secretary asked to be left alone with the terminal so she could have a more private conversation.

It's no accident that the 2013 Spike Jonze film that provided perhaps the most sophisticated depiction yet of a personal AI was called "Her" and not "It." In real

life and in fiction, it seems we can't help but attribute personhood to smart machines. And now, as personal agents are perhaps on the cusp of becoming a new paradigm for computing, we still struggle to avoid using human terms to describe them—we speak of assistants, copilots, sidekicks, companions, and agents.

The potential of this yet-to-be-named new paradigm excited us at Salesforce Futures as we planned the inaugural issue of our quarterly magazine. Our team's mission is to help Salesforce's leaders and our customers anticipate, imagine, and shape the future, and it is in this spirit that we dedicate our whole issue to personal AI futures. To be clear, we're focusing here on the AIs that will permeate our personal lives, as distinct from the types of tools we and other enterprise software companies are building into our platforms.

While technologists in Silicon Valley and elsewhere are researching and developing agents, too few outside relatively narrow



Welcome

circles have a sense of either what it will take to deliver them or the impact they might have on our lives, both as individuals and for the businesses that serve us.

As personal AI becomes more mainstream, we believe digital intermediaries will remake the relationship between customers and companies. How do you advertise your offerings in such a world? Will future users instruct their agents to block all forms of sponsored interruption, bar a few trusted brands? Will personal agents negotiate with company agents like those being built by Bland.ai and Parcha? What does a customer relationship strategy look like in this kind of future? It's a good time to ask such questions.

Imagine if, in 2006, you'd known how the combination of search, mobile apps and social media would revolutionize how businesses connected to their customers. We may be on the precipice of a similar moment of change. Now is the time to start to understand the contours and implications of this shift, before it arrives.


In this first issue, we dive into these new personal AI creations, take a careful look at what needs to happen to make them a reality, and imagine a range of plausible near futures shaped by their presence. The goal is to help you explore personal AI futures so you can start asking better questions, and perhaps even making better decisions, today.

In subsequent quarterly issues, we'll examine other possible impacts of AI on business, work, and society.

Oh and back to the yet-to-be-named paradigm: in our main article, we use the term agents as it's most widely used in the technical literature. In the speculative futures section of the magazine, we conjure other possible ways of identifying personal AIs.

The Salesforce Futures Magazine is our newest way to help you take an organized look at plausible and possible futures. We hope it inspires ideas, questions, and conversations, and we look forward to engaging with you on it. Drop us a line at futures@salesforce.com.

The Salesforce Futures team





Contents



6 **Orienting Personal AI**
A brief history and definitions

7 **Anticipating Personal Agents**
Three questions about personal AI futures

- 8** How does your agent work?
- 14** Why would you want a personal agent?
- 19** Do agents represent a viable new business category?

27 **Imagining Personal AI Futures**
Looking out 3-5 years

- 29** Will personal AI level the playing field or widen the gap?
- 36** Who pays for personal AI?
- 43** Will personal AI come from new or existing providers?

51 **Credits**



Significant moments from science fiction to non-fiction in the development and portrayal of personal AI.

Science Fiction

1968
2001: SPACE ODYSSEY
 Arthur C. Clarke's film series featuring the AI character HAL 9000

1970
COLOSSUS
 Supercomputer, Colossus, manages US nuclear defenses

1982
BLADE RUNNER
 Advanced AI beings are known as "replicants"

1984
THE TERMINATOR
 an AI defense network becomes self-aware

2001
ARTIFICIAL INTELLIGENCE
 Explores the emotional journey of a childlike android in a future society

2013
HER
 Explores the complex relationships between humans and AI

2014
INTERSTELLAR
 Introduces a multi-functional robot, TARS

2017
BLADE RUNNER 2049
 Expands on the original exploration of AI and what it means to be human

1966
ELIZA
 Chatbot developed by Joseph Weizenbaum

1972
PERRY
 Chatbot developed by Kenneth Colby

1997
DEEP BLUE
 IBM

2012
GOOGLE NOW

2014
ALEXA
 Amazon

2011
SIRI
 Apple

2013
ATLAS
 Boston Dynamics

2022
GPT 4
 OpenAI

2020
GPT 3
 OpenAI

2015
CORTANA
 Microsoft

Non-fiction





What even is an agent?

“Agent” is the term most commonly used in technical literature to describe functions closely associated with personal AI futures, but what does it actually mean? The term, rooted in economics, refers to the ability of actors to make decisions in specific contexts. This fits with how we’re beginning to conceptualize our personal AI tools.

Our Salesforce AI colleagues are exploring agents from multiple dimensions, and their definitions foreshadow themes we’ll explore in this issue:

“Agents (are) software entities capable of performing tasks on their own, ultimately in the service of a goal, rather than simply responding to queries from human users.”

SILVIO SAVARESE
CHIEF SCIENTIST
SALESFORCE

“Agents are personalized, persistent, and proactive.”

PHIL MUI AND ITAI ASSEO
AI R&D
SALESFORCE

“The AI agent can do all the same things as an assistant or concierge — but it also can perform actions without being asked.”

KATHY BAXTER AND
YOAV SCHLESINGER
OFFICE OF ETHICAL
AND HUMANE USE
SALESFORCE



Anticipating Personal Agents

Three questions about personal AI futures

We think personal AI agents might just change the world, in time. It's a big claim, we know. To get to that point personal agents will need to meet three criteria for breakthrough innovation—feasibility, desirability, and viability. In other words, they'll need to be able to technically deliver certain capabilities, those capabilities will need to enable offerings that people want, and it'll need to be worth someone's while to deliver that, i.e. there's a business model. None of those things are certain.

In this article, we'll lay out possible paths from where we are today, which we'll then selectively bring to life in the speculative futures that follow. With the rapid pace of change in AI today, it's really hard to identify the key drivers of change, let alone make accurate predictions. Our hope is that the framework we lay out helps readers sense-make their way through the torrent of new developments we'll see in the coming months and years.



ANTICIPATING PERSONAL AGENTS: FEASIBILITY

How does your agent work?

First, we'll examine feasibility and the question of how agents actually work, diving into knowing, remembering, reasoning, and doing to help you establish baseline capabilities agents need to deliver real world utility.





KNOWING

Unlike earlier assistants—think of today’s Siri or Alexa—agents will know a lot about the people who use them. Users won’t need to state their context and what they’re trying to accomplish at every turn—the agent will already possess sufficient information. Think of the degrees of contextual knowledge required for this to be true. It’s a spectrum that runs from the most basic level—who I am, where I am, what I’m doing, what I’m wearing, etc.—to much deeper, far more nuanced emotional levels referred to as “sentiment analysis”—how I feel, why I feel it, etc. For an agent to provide true contextual intelligence and pair that intelligence with action, a full range of inputs is needed.

Simply put: the more an agent knows about us, the better it will work.

But assembling the right data and insights to support the work of an agent operating in real time is a non-trivial challenge. Current LLMs come equipped with knowledge of the baseline data that informed their training. For most powerful LLMs today, training data includes content available on the Internet and proprietary data licensed from original content providers. As impressive as this knowledge is, it still lacks the kind of contextual understanding of the user required to make a world of personal agents a reality. What’s needed is enough contextual information to enable persistent representation; knowledge of what to remember...and what to forget; and the ability to safeguard privacy and trust through all that.



REMEMBERING

For an autonomous agent to be useful, it needs to have memory: the capability to store and retrieve information.

Memory in AI is not a single entity, but rather a collection of layers. One layer is the ability for the personal AI to store and retrieve the prompts that the user initiates across multiple sessions. We refer to the ability for a personal AI to “stay” with a user over a longer time horizon (i.e. across multiple interactions) as persistence. The second layer is what’s already trained into the LLM and how the relationships are built between the pieces of data that inform the training. There are limits to this layer such as the cutoff dates for the training data. The third layer of memory is the ability of personal AI to retrieve data (both structured and unstructured) from other data stores by using methods such as RAG (retrieval augmented generation).

Currently, context is provided by the user to the LLM, but utility is limited by the LLM’s inability to store the information provided (context length), and by the accuracy with which the user states their preferences. The user’s actual context may also include data the agent has no access to, thereby increasing the likelihood of “off” responses. Conversely, an agent might one day have access to the user’s entire recorded history, which would seemingly offer unlimited context (not to mention near unlimited data storage). But even if that’s the case, how does the agent know which parts of a user’s history are most important, or which parts are more sensitive?

The challenges with memory and meaning for LLMs echo the challenges of understanding how humans assign meaning to memory.



In a recent New York Times interview, the neuroscientist Charan Ragnarath points out, “we [humans] have a longer timeline in which we can integrate information than many other species.” Humans bring an incredible amount of contextual knowledge to any situation, and it’s a significant challenge for agents to replicate that.

Innovative companies are already considering workarounds. For example, the ability to curate a personal AI’s memory may become a critical feature. OpenAI recently announced memory control features in ChatGPT, explicitly giving users control over which interactions are remembered and which are forgotten.

Apple’s new Journal App prompts users to add rich written or audio notes to contextual information the device already has access to, such as photos that were taken, or locations where time was spent, or calendar appointments. This is a potentially powerful step towards assembling a full contextual understanding; backwards-looking to begin with but perhaps capable of informing a personal agent in the future.

REASONING

If you’ve seen Her, you’ll remember the OS1 onboarding scene where Joaquin Phoenix’s character, Theodore Twombly, is guided through a short series of personal questions to configure his AI.

Memorably, he’s asked, “How would you describe your relationship with your mother?” Theodore begins a stuttering answer, but as he’s just getting going, the onboarding process surprises him by cutting him off. The OS has extracted enough information to make the design choices it needs to start off with.

Over the course of the film, Theodore’s AI, Samantha, continues to learn and adapt. The ability suggested by this onboarding scene to extract meaning and determine complex choices from limited information evokes a powerful machine intelligence with sophisticated reasoning capabilities. In reality, how such reasoning might work at scale is an open question.



Consider the core function of the “planning skills” Lillian Weng and others refer to when it comes to the agents of today. Heuristic search and reinforcement learning have limitations in dynamic, real-life environments. For agents to truly deliver Samantha-like capabilities, they’ll need to move from heuristics to inference, reasoning, and metacognition.

One can argue that the latest LLMs already exhibit emergent behaviors that sometimes look like reasoning. According to a survey by Huang & Chang, a number of past research papers have already made claims for the reasoning capabilities in some models. Prompting to generate intermediate steps to solve problems (e.g. Chain of Thought and Tree of Thought) has produced promising results. However, LLMs still struggle with tasks that are easy for humans, such as performing complex math and generating action plans to complete tasks in given environments.

Engineers are attempting to overcome these limitations by incorporating planners that break down tasks into steps and chaining multiple systems together. Early tools like BabyAGI and Auto GPT triggered questions about trust, privacy, accuracy, and reliability, but they inspired an innovation race to deliver greater autonomy using similar methods.

An opposing view suggests LLMs, due to their lack of an internal world model to predict states, may never achieve a level of reasoning sufficient to power autonomous agents.

Those who believe this are pursuing different approaches. For example, some researchers are developing models that train AIs to learn about the world more the way humans do, which may provide an alternative path towards generalized reasoning and planning. In the future, we expect autonomous agents to have more sophisticated reasoning and problem solving capabilities.

Daniel Kahneman popularized a simplified model of how thinking works in the brain, using the idea of two different modes of thinking:

System 1, “thinking fast”, is intuitive and unconscious. System 2, “thinking slow”, is deliberate and rational. This model offers a useful way of thinking about agentic capabilities. These systems coordinate so that when a problem is too complex for System 1, System 2 kicks in with more sophisticated logic and reasoning. Characteristics of System 1 are analogous to the strengths of current deep learning models, while System 2 maps to the weaknesses of the existing LLMs.



DOING

So far, we've outlined how knowledge, context, memory, persistence, and reasoning might combine to deliver a truly transformative personal agent. The copilot concept is well established, and we see early examples making their way into the market, especially the enterprise software space.

Arguably, the most groundbreaking result of all these capabilities coming together will be, per Ninja AI, “autopilot”, or the ability of agents to take action and get things done without requiring the user’s attention or even explicit command.

Salesforce Chief Scientist Silvio Savarese is pioneering the concept of Large-Action Models (LAMs) which will have the ability to “automate entire processes” by combining “the linguistic fluency of an LLM with the ability to accomplish tasks and make decisions independently.” Savarese believes these changes will elevate AI “from a passive tool...to an active partner in getting work done in real time.”

Autonomous autopilots are the source of a lot of interest and fascination, but they're currently operating at a much more experimental level that underscores the difficulty of jumping from human in the loop to true autonomy. There is good reason to take this transition slowly. Holly Prouty from Salesforce's cross-cloud

Research and Insights team recently argued convincingly for the importance of keeping a human in the loop with new AI innovations— particularly in the short-term to ensure trust. Salesforce's Office of Ethical and Humane Use (OEHU) has argued repeatedly for “humans at the helm” until we've cultivated trust and enabled users to learn the technology.

Now, let's consider some potential future use cases for these action-oriented agents, how we could oversee their functions, and why, if our predictions prove accurate, they will revolutionize our work lives, our personal lives, and customer experiences.

ANTICIPATING PERSONAL AGENTS: DESIRABILITY

Why would
you want
a personal
agent?

2





“LLMs will graduate to becoming agents with more powerful tools, search, APIs, coding abilities, clicking on the web, etc. In particular, AI assistants for search are helpful enough and will replace Google for many young people, students and knowledge workers.”

RICHARD SOCHER
YOU.COM CEO

Last year our team, in collaboration with Salesforce Studios and Salesforce Design, imagined a plausible future of how agents might change the shape of work in Salesforce 2030: A Glimpse of an AI Future. Jordan, the trailblazer hero, interacts with her own “chief of staff” agent and a range of more specialized agents to make faster, better decisions in collaboration with her human team at Koa Coffee. What we chose not to illustrate were the personal agents that might support end-consumers on the other side of these transactions.

The transformative promise of agents lies in their ability to compensate for human limitations. Anyone who has struggled to finish one task, let alone juggle many of them, can understand the value proposition.

INFINITE INTERNS + PATIENCE

Imagine an army of infinitely patient interns who stand ready to work on your behalf. These interns can either collaborate alongside you in the flow of work, or they can labor independently, periodically checking in to ensure they’re on the right track. The more they work with you, the smarter they get, learning how you work and think across modes and contexts. Think about a world where all of us have the kind of expert staff currently enjoyed by CEOs: gifted helpers who learn our preferences, understand our goals, engineer outcomes, and specialize in doing all of the things we don’t want to do.

Tech blogger and consultant Venkatesh Rao encourages us to think about machine intelligence as fundamentally different from human intelligence, particularly when it comes to “attention”.



Our personal agents will have endless patience for tedious, detailed tasks (think taxes, paperwork, applications, and other) of the sort that sap our attention and, occasionally, our desire to endure the human condition. Agents have the potential to remove this burden.

PERSONALIZATION

Greater contextual intelligence and more persistent memory suggest agents will provide personalization that's far greater than what we see today. Itai Asseo and Phil Mui, who work on AI R&D for Salesforce, encourage us to think about personalization in three categories: "know me", "inform me", and "empower me".

In the "know me" category, agents keep your goals in mind, analyze your performance, and adjust to your unique style. Because every interaction with personal AI will be remembered—or stored as a state—and factored into future use cases, a flywheel effect takes hold: the more you use your agent, the better it gets at anticipating your needs in an intuitive way.

In the "inform me" category, we consider the possibility that agents could use their contextual intelligence to help guide and prioritize our attention and separate signal from noise. This has clear applications in the personal productivity space, but the implications on the consumer side are no less significant. Imagine an agent who helps you switch on "Zen Mode" and other filters so you can better tune your environment, and even remove things from your plate by acting on your behalf.

Finally, the "empower me" function speaks to personal AI's ability to serve as a coach/mentor, or even a manager.

"I use a similar definition of 'agent' to economists; that is, a narrow AI that acts on behalf of its principal (or user) out of sight, while that principal's attention is on other things."

CHRIS NOESSEL
AUTHOR



If personal AI can deliver such control, it will forever change the relationship between customers and companies by raising the bar for direct relational engagement.

NATURAL INTERACTIONS

Recent demos by startups Humane and Rabbit have started to tangibly articulate what agent-based offerings might look like.

Common early use cases include ordering food, finding gifts for loved ones, planning trips and events, and scheduling appointments. What these demos hint at is a more fluid and flexible collaboration between humans and machines, especially when it comes to attention. Both offerings promise less tapping and scrolling and more focus.

Both products also rely on conversational interfaces to do this, but we think voice will be only one of the ways people interact with their agents. Many proto-agents feature travel in their demos, but solving a multi-part travel puzzle that includes schedules, price comparisons, and airline and hotel preferences using only voice commands is less than ideal. Liz Trudeau of Salesforce Design encourages us to think about a more practical alternative, “Think about a collaborative interaction, a flexible interface that adapts to the task at hand and the job the user is trying to accomplish.” We like this concept, “UI on the fly”, because it emphasizes a truly responsive interface that adjusts accordingly as needs change. You can get a glimpse of what this multi-modal future might feel like in the much-discussed Google Gemini demo depicting the planning of a birthday party.

In aggregate, these developments point to futures where tools are easier to use and it’s easier than ever to get things done.

ADVANCED SKILLS AND LEARNING

In 2014, when Amazon launched its original Echo, the device was little more than a bluetooth speaker. Adding the Alexa assistant promised to turn the Echo into something else entirely: a conversational smart home hub, particularly as the Alexa Voice Service SDK toolset expanded the library of skills available to consumers. Alas, most Echos are still used mainly as speakers and the skills revolution has not happened. Nevertheless, the concept of a core agent and a library of additional skills offers a preview of what we might see in an agentive world.

The ability to train agents on data sets (including proprietary ones) means people can far more easily create useful personal AI tools for others based on specialized knowledge. This, in turn, empowers people to build more agents. Already, we can see Open AI allowing users to rapidly build their own GPTs, enhanced with advanced skills based on additional “instructions, extra knowledge, and any combination of skills.”

These advancements suggest a forthcoming Cambrian explosion of agent skills and capabilities, personalized for an infinite variety of tasks. We know this would transform how companies interact with customers. What we don’t know is which sectors will be transformed first.

UNMET NEEDS

A key element of Clayton Christensen's theory of disruptive innovations is that early disruption will come from offerings that target the unmet needs of customers at the bottom of the market, hitherto uneconomical to serve. Often, these offerings appear inadequate to the mainstream of the already served, but they can grow and improve from tiny seeds. Looking at the emerging agent landscape, we can see an example of this dynamic in the wave of AI companions, such as those provided by Replika, character.ai, and Baidu's Wantalk.

These companions, chatbots, and avatars serve the needs of the lonely, the elderly, and those who cannot afford costly therapy.

While companions can generate an “ick factor” response among mainstream audiences, observers such as Andreesen Horowitz have identified how the “companion stack” may hold clues as to the future directions of the larger agent space. For example, we may see agents that super-empower individuals, catalyzing a new wave of entrepreneurship in developing countries by serving needs that could never have been met before, and thereby generating a similar impact to that of mobile phones a generation ago.

The evolution of agents will be determined by consumer preferences and progress against technical challenges, but market dynamics and the way businesses and consumers balance trade-offs will play an equal role. In the next section, we look at how the business of personal AI might evolve.

ANTICIPATING PERSONAL AGENTS: VIABILITY

**Do agents
represent a
viable new
business
category?**



Assuming agents are technically possible and people want them, there still needs to be a compelling business case for companies to bring them to market. Which companies are positioned to win in the emerging marketplace? Who pays for these things? How does the business model for services shift in a world of agents? How can we trust them? Let's dive into how this emerging market might evolve.

WHICH COMPANIES ARE POSITIONED TO WIN: INCUMBENTS OR STARTUPS?

Historically, startups have thrived in moments of disruption. However the unique dynamics of personal agents—with their reliance on vast data, expensive compute, and the critical currency of trust—may tilt the playing field toward incumbents. Established players can also use entrenched hardware and closed ecosystems for further advantage.

In thinking about who might win, it's helpful to break personal AI out into a few component parts. First comes data, the life blood of any agent. In a 2023 post entitled "The Four Wars of the AI Stack," Swyx et al. from influential AI newsletter Latent Space point to the "data war" as one of the major theaters of competition shaping the market. So we see LLM creators hammering out licensing deals with content providers, the New York Times's legal fight against OpenAI and Microsoft, and heightened interest in synthetic data within the AI research community. The next component is the model employed, which is the motor powering a given agent's decision-making process. We see GPU chips as a part of this category. The third component is user interface (UI), the bridge between AI's potential and its practical utility.

“Form follows funding.”

STEWART BRAND
AUTHOR



Finally, distribution plays a critical role in competition. Even if consumers love a product and it lives up to its technical promise, a lack of distribution presents significant headwinds when it comes to long-term growth (see the fate of Netscape Navigator). The aggressive acquisition strategies by incumbents point to a winner-take-all dynamic. Incumbents want lock-in an existing user base before upstarts can scale distribution. As author and analyst Sangeet Paul Choudhary points out, these companies see a unique opportunity to become the primary interface for AI interactions in both work and personal life.

Choudhary also posits that people don't want an army of agents, they want one agent. If he's right, the stakes for companies competing to provide such an agent couldn't be higher. Winning could mean the opportunity to control a choke point on any transactions that flow through a personal agent. Network effects could further entrench winners: the more interactions that happen on your platform, the faster the learning loop gets, the more people use your products, and the better your service gets. All of this adds up to higher switching costs and more defensible moats.

Incumbents have considerable advantages over startups. First, they're more likely to already have access to relevant personal data, not to mention the compute to process it. Second, many are already leaders in building the leading LLMs or are using their market power to access them.



Finally, and perhaps most crucially, incumbents already possess the distribution channels, including the most-used hardware, apps, and websites through which to deploy personal AI solutions. These established companies, while far from perfect (some further than others) also have established practices around privacy, content moderation, security, and ecosystem governance—including the power to shape interoperability standards—that may give them an edge grabbing market share in a complex, multi-agent world.

The Humane and Rabbit demos mentioned in the previous section capture our imagination because they hint at the possibility new upstarts may shake up an ossified consumer tech landscape. We imagine new hardware, new platforms, and new modes of interaction. News such as the rumored hardware collaboration between OpenAI and former Apple Head of Design Jony Ive only adds to our sense of anticipation.

Before we get too carried away, in our opinion, Humane and Rabbit compete primarily on UI, betting they can deliver something so compelling consumers will abandon their existing devices and jump on board. Less certain is the quality of data these companies have access to, how effective their models are, and how much their distribution can scale.

OpenAI's ambitious efforts to link their model with UI and distribution seems more promising. For an upstart to win in the personal AI space, they will need to solve the distribution problem and become a platform. This is no easy feat. A look at the history of technological innovation reveals a pattern where excitement about decentralization and openness runs high early in the innovation cycle, only to give way to more centralized options once companies figure out ways to make money off centralization. It's also notable that according to reporting by the Financial Times, Amazon, Microsoft and Google accounted for two-thirds of the \$27bn raised by AI startups in 2023.

Any new platform would have to get big enough, fast enough to fend off acquisition and establish its position. None of this is to say that upstarts can't emerge and knock the tech giants off their pedestals. It's early yet. Perhaps personal AI disrupts some sectors and categories while sustaining others.

WHO PAYS FOR THESE THINGS: SERVICES > ADVERTISING?

The business model of the internet is powered by advertising, but we don't really know how ad-supported revenue will evolve in a personal AI world. Could advertising as we know it today disappear entirely?

Because personal agents are so intimately connected with their users, the value exchange will be more explicit and examined than it is with today's applications.



”

“Consider this thought experiment about personal agents and privacy. Imagine being constantly surrounded by cameras that track your location in space, sensors that track your emotions and thoughts, and devices that record every word you say. This is the kind of future we foretold in *Minority Report*, and it was not a future people wanted to live in.

How many people would put up with that in exchange for an agent that could give you answers grounded in the best intelligence? Who would actually be willing to give their agent enough context to ensure its actions are as accurate as possible? I know a few people who would make that trade, but I know a lot more who wouldn't.



PETER SCHWARTZ
CHIEF FUTURES OFFICER AND
ADVISOR, OFFICE OF THE CEO
SALESFORCE

A note on privacy

We're already sharing a lot of our lives online and are under pressure to share more. Recent ads for the Meta/Ray Ban smart glasses emphasize the benefit of live streaming and recording while continuing to, in the specific ad in question, dance through the streets wearing a yellow jumpsuit. New AI tools, such as Google Gemini, default to requiring access to all your personal messages and permission to retain your conversations for years.

Peter's thought experiment highlights the complex trade-offs around privacy and utility, what we've referred to for more than a decade as "the creepy or cool equation." If sharing your data creates sufficient value for you in return, it's cool. But if you perceive that the company asking for your data gets most of the value, it's creepy.



Agent gatekeepers with access to vast stores of knowledge might radically diminish the habits of consumer research and comparison, which is where discovery-focused advertising lives. Consumers may also turn on those filters we mentioned earlier. If all this happens, the rise of personal agents could deflate the advertising revenues that power the internet the way Craigslist did to print media.

Users will use something, and pay for it—either with time, attention, data, or their hard-earned dollars—if they find it useful. Those who subscribe to ChatGPT or Google Gemini today pay for access to a baseline set of skills, but it remains an open question if people will continue to pay for intelligence as competition grows fiercer. When personal AI requires additional skills to accomplish tasks, a plausible outcome may be the proliferation of pay-per-skill marketplaces where consumers pay to upgrade their agents in an incremental and case-by-case way. If personal AI agents are able to accomplish high-order tasks and goals, one can also imagine people paying for outcomes in the same way we see now in some B2B sectors.

In the multi-agent world, where agents are interacting with other agents at the speed of light, we will likely see an emergence of a new economy where deals and transactions are completed with personal agents acting on our behalf.

Consider how personal agents could deliver growth through machine-to-machine, or M2M transactions. Gartner and others have been exploring these ideas recently, but they were anticipated as far back as 2011 by economist W. Brian Arthur in “The Second Economy,” a classic article foreshadowing the rise of “autonomous economy” where machines communicating with each other drive economic value independent of humans. Chris Noessel’s concept of “post-attention value,” Kevin Ashton’s “Internet of Things,” and other ideas also point us to new opportunities this world presents.



Given all of this economic potential and how much advantage agents will confer on their users, one additional possibility is that state-sponsored agents step in to democratize access.

In this world, access to powerful AI, as Salesforce CEO Marc Benioff recently put it, is seen as a right, not a privilege, and the line between the public and private sector begins to blur. In a later section of this magazine, we imagine one way this future might play out.

TRUST WILL BE ESSENTIAL FOR EFFECTIVE PERSONAL AGENTS

Trust, Salesforce's number one value, will play a fundamental role in shaping personal AI futures. Agent relationships are, by default, intimate and ongoing. Agents' need for context pushes trust to the forefront, suggesting a world where aligning with human goals becomes the primary measuring stick for personal AI and a must-have ingredient for long-term customer relationships. Trust concerns about agents stem from ethical and technical issues.

Salesforce's Office of Ethical and Humane Use (OEHU) is a leading voice on ethical concerns, and an excellent source of responsible AI frameworks, guidelines, and best practices that help us anticipate and guard against trust failures. The OEHU has implemented consequence scanning in how Salesforce designs and builds products.

In their recent HBR article on how to build trustworthy assistants, Kathy Baxter and Yoav Schlesinger of the OEHU point out four key issues that will undermine trust in agents: integration and amplification, a lack of a human in the loop, poor data quality, and responsibility gaps.

Technical concerns, such as security, reliable grounding, and protocols for cross-ecosystem access challenge some of the loftier fever dreams of a future dominated by agents. Early incidents like Air Canada being forced to acknowledge a policy invented by a chatbot show how far we still need to go before companies and consumers can trust personal AI across all of the use cases mentioned here and elsewhere.

Recent research by Salesforce points to a gap between customers' overall trust in companies and their belief those companies will use AI responsibly. For example, while 76% of customers trust companies to make honest claims about their products and services, only 57% trust them to use AI ethically. This trust gap points to the importance of designing trustworthy AI from the outset, particularly when it comes to considering the role humans will play in AI systems.



IN CONCLUSION

How will agents shape our world?

Our goal has been to equip you with a framework for understanding where personal AI agents are today and the most critical uncertainties that will need to resolve for the category to truly emerge. But the question of how feasible, desirable, and viable personal agents will only be answered over time.

This framework allows us to imagine with greater clarity a set of plausible, relevant, and challenging futures, with a particular focus on how personal AI may revolutionize the relationship between customers and companies.

Next, we explore three pairs of scenarios around the following questions: Will agents level the playing field or widen the gap? Who pays for agents? Will we get our agents from new or existing providers?

Turn to the following section, Imagining Personal AI Futures, to see how different answers to these questions might result in very different futures, experience what those futures might look and feel like, and learn the key questions and strategic implications businesses should consider.



Imagining Personal AI Futures

Looking out
3-5 Years



At Salesforce, our mission is to anticipate, imagine, and shape the future.

In this section, we imagine some of the ways things might play for personal AI over the next 3-5 years. For each of our three big questions, we consider a pair of possible futures, bring these futures to life with speculative artifacts like ads, news articles, and product pages, and call out the most important questions and implications for businesses to consider.

metadata

Each future is calibrated with three potential measures of impact.



SURPRISE FACTOR

How unexpected is this future?



TIMELINE

How far out is this future: near-term, long-term, far-out?



SCALE

Could this future scale from a local to global level?





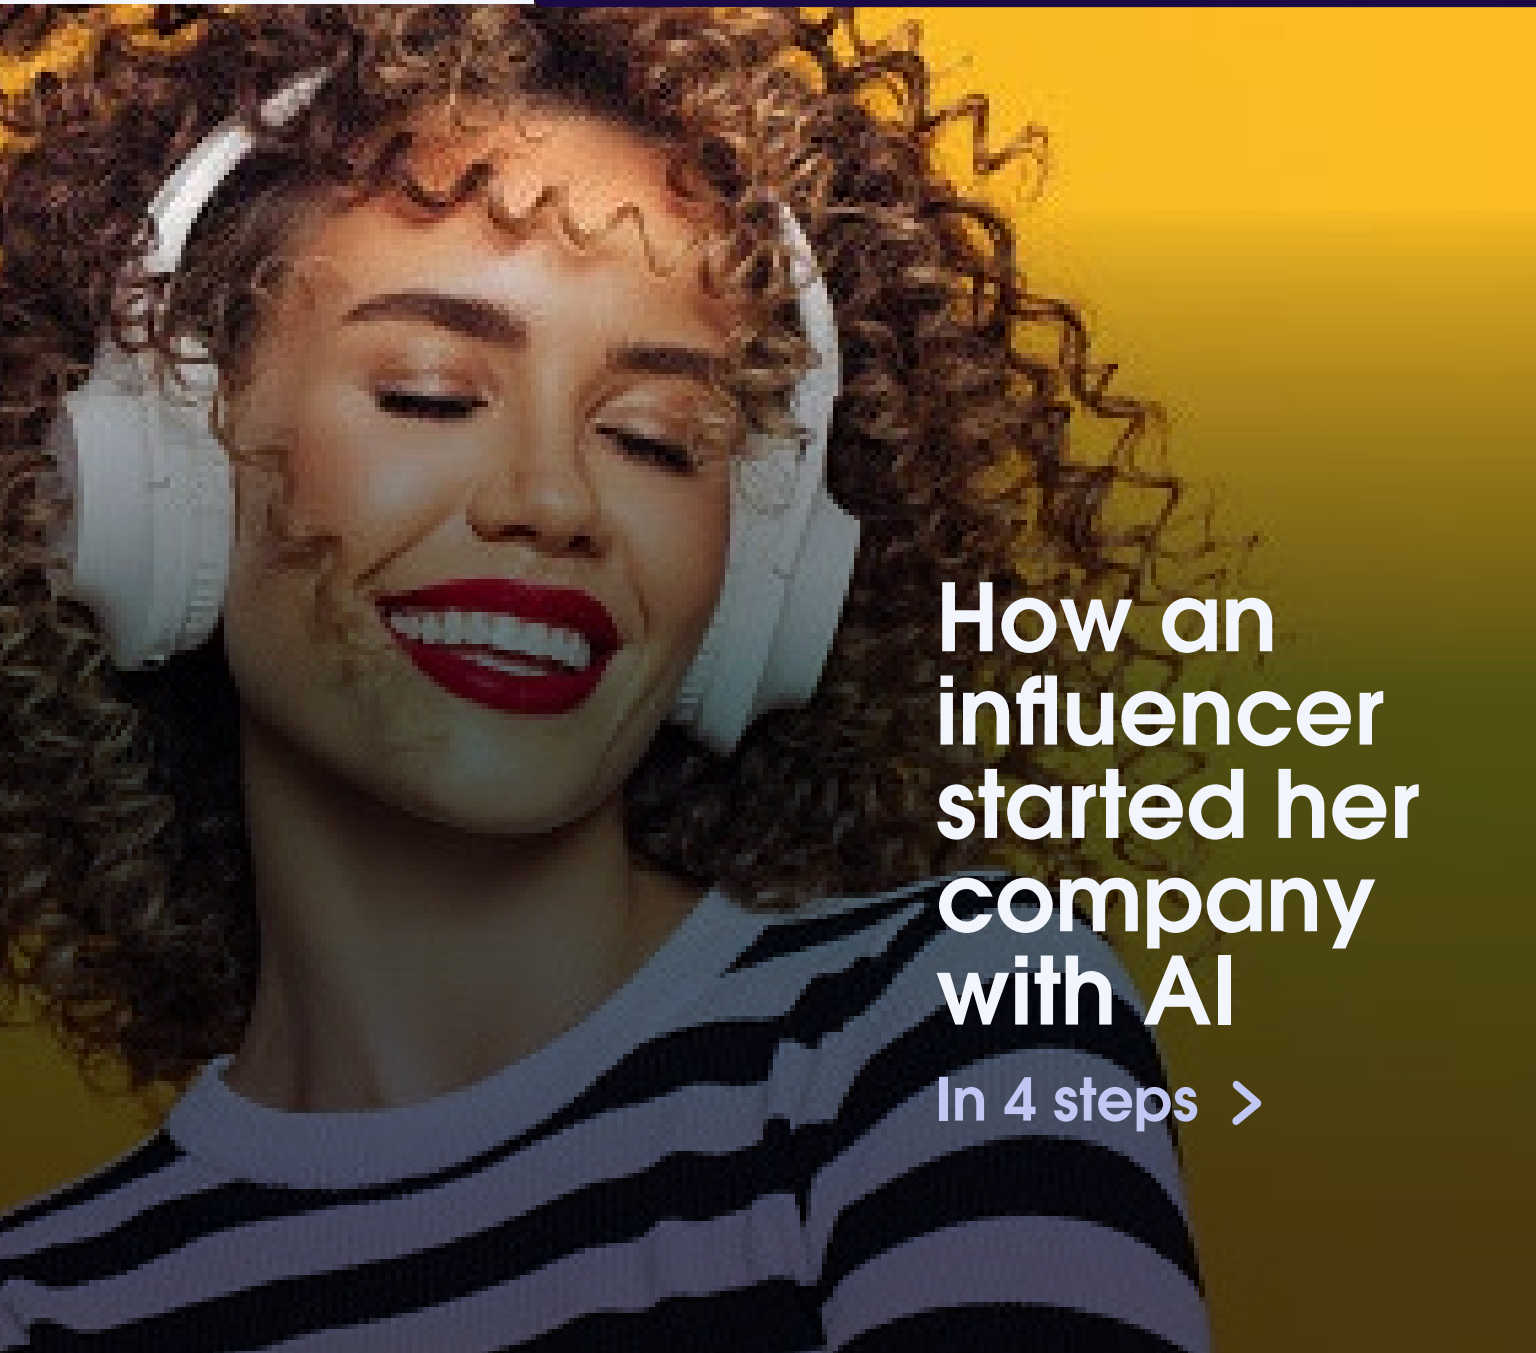
Will personal AI level the playing field or widen the gap?

In a future filled with personal AIs,

access to powerful, latest-gen tools will be critical to competition. Who gets access, and what they get access to, will determine whether opportunity expands or agents exacerbate existing inequalities and serve the privileged few.

What if AI levels the playing field?

Personal AIs turn out to be great equalizers, delivering more access to education, career coaching, and business services for those with the drive to access them. Broader availability of better tools expands opportunity and bridges the achievement gap. Accessibility and inclusion improvements flow to all. Companies and society also benefit from access to a broader talent pool as people learn to use personal AI tools to fill in knowledge gaps.



How an influencer started her company with AI

In 4 steps >



With no business or tech background, Jordan turned her highly opinionated, insanely popular headphone reviews into an empire. How did she turn likes into buys?

1

Everyone knows, I have strong opinions on headphones. I want sustainable, durable products that look good, hold up, and sound so good they make you cry. I started doing the review videos just for fun. Suddenly, I had an audience. Then, before I knew it, I had a really big audience.

2

My fans kept reaching out to me with ideas.

It was, “You should design your own headphones,” or, “You should charge for personal advice.” These were not things I knew how to even consider, let alone do. I started asking some of the new AI tools for help.

3

These tools became my digital team. For me, the first winning strategy, in addition to building my audience, was partnering with companies I admired to help them refine designs before they went to market. My digital team members helped me learn how to incorporate my business, pitch to investors, design my brand, set up payments, refine my offering, get it to market, and run my first marketing campaign. Basically everything.

4

Eventually, I realized I no longer needed to partner with these companies. I had relationships with designers, manufacturers, and distributors. I had loyal customers. It was time to make my own products. I used digital team members to vet partnership opportunities, draw up contracts, and optimize my audience outreach. I just had my weekly meeting with my new digital CFO. The way I think of it is, I have six human team members and hundreds of digital ones.





surprise



timeline



scale

Questions and Implications

If personal AI levels the playing field...

How will AI tools super-empower new entrepreneurs?

What sectors and activities will be most disrupted by a wave of AI-driven startups?

How might companies use AI to up-level their own talent?



What if personal AI widens the gap?

The most effective personal AI “chiefs of staff” are pay-for-play luxury items, used by elites to sharpen their edge and further their advantages. The best chiefs of staff deepen societal divides and reinforce consolidation of power by offering exponential skills and productivity gains to users who can afford them.



**Meet Victor, the
world's best AI
chief of staff >**



Think of Victor as the entire personal team of a Fortune 500 CEO all rolled into one supremely competent digital employee. The best part? You don't need a CEO salary to afford elite service.

With three convenient price tiers, there's a Victor available to everyone.

Bronze

Pay-per-skill 5 skills monthly; one-time fees

Ad-supported

2 hours of weekly meeting support

One size fits all
Our most popular personality

FREE

Gold

Discounted skills
25% off skill purchases

3-4 Ads weekly

10 hours of weekly meeting support

Semi-customized
Choose from a selection of personalities

\$25 / MO.

Platinum

Unlimited Skills Choose any skills from our library

Ad Free

Unlimited meeting support

Fully Customized
Personality adapts and learns based on your inputs

\$100 / MO.



surprise



timeline



scale

Questions and Implications

If personal AI widens the gap...

When employees bring their personal AIs to work, will their productivity correlate with their individual ability to pay for AI skills?

Will AI chiefs of staff allow their users to scale their productivity?

How will this affect income inequality and perceptions of fairness around AI?



Who pays for personal AI?

**As the old
proverb says,
you don't get
something for
nothing.**

Whose interest will the agent act in? While incentive structures lurk in the background of our current technologies, we believe the intimacy of agentive interactions makes the value exchange between users and buyers far more explicit. The way this market evolves determines who controls knowledge, who controls information, and whether access to personal AI is a privilege or a right.



What if personal AI is ad-supported?


“Good enough,” ad-supported products force users to endure sponsored messages in exchange for new skills, tools, and better deals. How much would people pay to ditch these interruptions and access more premium offerings? What happens if they can’t afford the alternatives?



AD -
SUPPORTED



BELLA



**One
interface.
Every
application
in the world.
No monthly
fees.**

Business. Entertainment. Productivity. Travel.

Bella partners with the world's best companies to connect you to the offerings that make the most sense for your needs. You win. Our partners win. Bella wins.

Join millions who love accessing the world's best concierges without expensive monthly payments.



surprise



timeline



scale

Questions and Implications

In an ad-supported AI future...

How might an ad-supported agent work as a business model?

How will consumers trade off attention for personal AI skills and services?

Might ad-supported models expand access to personal AI to previously underserved segments?

What if personal AI is state-supported?

In another plausible future, the provision of public goods includes state-sponsored aides for all citizens. It's easy to imagine the benefits: efficient discovery of resources, community initiatives, and services. But using these services means buying into a public ethos, and there could be a worst-case scenario where some states abuse this access to reinforce control.

STATE -
SUPPORTED





Activate Your AIDE!

The more you use your AIDE, the more you do for your country.

Access healthcare.

Access financial benefits.

Pay transit fees.

Enter state drawings.

Apply for university.

Apply for government jobs.

Report anti-social behavior.

NO
ADS!

AIDE.gov





surprise



timeline



scale

Questions and Implications

In a state-supported AI future...

For what use cases might a highly functional, state-supported personal AI be more appealing than private sector alternatives?

How might companies partner with governments to deliver public services?

Could this alter considerations about which countries are most desirable for citizenship?




Will personal AI come from new or existing providers?

Companies with vast resources, including data, compute, and intelligence,

would seem to have the advantage in a race for personal AI market share. When new upstarts like Rabbit and Humane emerge, one question you often see in the comments sections where people discuss such things is, “Great, but why can’t I just do this on my phone?” Still, disruptive technologies have a way of disrupting established paradigms. Might new providers emerge who take the personal AI market in totally unexpected directions?

What if personal AI comes from existing providers?


Incumbents bundle personal AIs with existing products and services. This is the world of Siri on steroids and Alexa on amphetamines: familiar players with new, next gen personal AI interfaces that intermedate the relationships between customers and companies (and take a small piece of every transaction). This world isn't perfect, but a lot of consumers love it.



Maryam is a Hawaii-based freelance stitcher with a packed schedule. She loves time with family and friends, especially in nature. How does she do it?

How may I help you?

Maryam, what would you like to do?



These are a few of my favorite prompts....

● "Streamline my calendar"

Scans all meetings, compares them to past meetings, and adjusts allotted time accordingly.

● "Go Zen Mode"

This kills every distraction unless there's an emergency involving a loved one. It's essential. Admittedly, sometimes I do miss my distractions, but not enough to give up time outside.

● "Prioritize my people"

I want to be with friends and family and find time to get outside. This command helps you schedule these things well in advance, which improves the odds they happen.

● "Negotiate. You can do better."

My AI is a ruthless negotiator. If I tell her "you can do better," she gets even sharkier. I should probably check and see if there's a better deal somewhere, but who has the time?



surprise



timeline



scale

Questions and Implications

If personal AI comes from existing providers...



Will there be network effects here, or will rent-taking confer disproportionate power to fewer companies?

Is voice really the most dominant mode of AI interaction, or do other modes ascend?

Similar to the explosion of advertising technology companies that emerged with Web 2.0, will a new category of enterprise software emerge around personal AI?

What if personal AI comes from new providers?

Just as people settle on 9 or 10 apps on their phone today, people choose nine or ten AI-first products for the vast majority of their needs. Users highly value control in selecting their tools rather than letting an AI oversee these choices. As a result, we see an explosion of AI startups that target specific use cases.





Meet My Crew

Tao is a Manila-based marketer/product designer. He uses personal AIs for productivity, production, recreation, and everything in between. Tao shares his favorite AIs on Meet My Crew.



EINSTEIN

“My real genius is making your customers happy.”

Customize Einstein with the skills and tools you need to get the job done.



SERGEANT PEPPER

“Last task, 3 minutes. GO!”

Struggling with procrastination? Need some motivation? Add pepper.



DIGITAL YENTA

“I’m a spicy matchmaker here to save you from swiping.”

Keeping your love life in check...



LULU

“Feeling creatively stuck on a design?”

She may never give you the answer you expect, but you might find the answer you need.



DR. JULIA SILVA

“Talk to me anytime, day or night.”

The real Dr. Silva will learn from your conversations so that your next visit with her is even more useful.



surprise



timeline



scale

Questions and Implications

If personal AI comes from new providers...

How might startups get to context faster than incumbents without losing trust?

In a world where people see value in having multiple agents, how do people find, vet, and purchase personal AIs?

Will new personal AI providers also open the door for new players in enterprise software?



Will AI help us build better futures?



references and inspiration

FUTURES CURATORS

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

SALESFORCE

“HOW COMPANIES CAN BUILD TRUSTWORTHY AI ASSISTANTS”

Kathy Baxter and Yoav Schlesinger | [HBR](#)

“ARTIFICIAL GENERAL INTELLIGENCE: CASTING A VISION FOR A BETTER FUTURE”

Silvio Savarese and Peter Schwartz | [THE 360 BLOG](#)

“TOWARD ACTIONABLE GENERATIVE AI”

Silvio Savarese | [THE 360 BLOG](#)

“BUILDING TRUST IN AI NEEDS A HUMAN TOUCH”

Holly Prouty | [THE 360 BLOG](#)

“GENERATIVE AI: 5 GUIDELINES FOR RESPONSIBLE DEVELOPMENT”

Paula Goldman and Kathy Baxter | [THE 360 BLOG](#)

“WHAT IF SUPERINTELLIGENCE RAN YOUR BUSINESS”

Itai Asseo | [THE 360 BLOG](#)

“FROM COPILOT TO COORCHESTRATION”

Silvio Savarese | [SALESFORCE AI RESEARCH BLOG](#)

ADDITIONAL ARTICLES

“THE NEW NEW MOATS”

Greylock Partners | [LINK](#)

“OOZY INTELLIGENCE IN SLOW TIME”

Venkatesh Rao | [RIBBONFARM](#)

“IT’S NOT A COMPUTER, IT’S A COMPANION”

Justine Moore, Bryan Kim, Yoko Li, and Martin Casado, Andreessen Horowitz | [LINK](#)

“ROBOT KEGGERS AND ROOMBA SPIES”

Christopher Noessel | [MEDIUM](#)

“A LEADING MEMORY RESEARCHER EXPLAINS HOW TO MAKE PRECIOUS MOMENTS LAST”

David Marchese | [NEW YORK TIMES](#)

“THE SECOND ECONOMY”

W. Brian Arthur | [MCKINSEY QUARTERLY](#)



references and inspiration

“UNLOCKING POST-ATTENTION VALUE”

[Christopher Noessel](#) | [LINKEDIN](#) | [BOOK](#)

“LLM-POWERED AUTONOMOUS AGENTS”

[Lilian Weng](#) | [LIL'LOG](#)

“THE ANATOMY OF AUTONOMY: WHY AGENTS ARE THE NEXT AI KILLER APP AFTER CHATGPT”

[SWYX and Alessio](#) | [LATENT SPACE](#)

ACADEMIC PAPERS

“LEARNING UNIVERSAL PREDICTORS”

[Jordi Grau-Moya](#), [Tim Genewein](#) | [LINK](#)

“TOWARD REASONING IN LARGE LANGUAGE MODELS: A SURVEY”

[Jie Huang](#), [Kevin Chen-Chuan Chang](#) | [LINK](#)

“CHAIN-OF-THOUGHT PROMPTING ELICITS REASONING IN LARGE LANGUAGE MODELS”

[Hason Wei](#), [Xuezhi Wang](#) | [LINK](#)

“TREE OF THOUGHTS: DELIBERATE PROBLEM SOLVING WITH LARGE LANGUAGE MODELS”

[Shunyu Yao](#), [Dian Yu](#) | [LINK](#)

“THINKING FAST AND SLOW IN AI: THE ROLE OF METACOGNITION”

[Marianna Bergamaschi Ganapini](#), [Murray Campbell](#) | [LINK](#)

VIDEOS

“SALESFORCE 2030: A GLIMPSE OF AN AI FUTURE”

[Salesforce Futures](#) | [VIDEO](#)

“THE DISAPPEARING COMPUTER”

[Irman Chaudhri](#) | [TED TALK](#)

“INTRODUCING R1”

[Rabbit r1 Demo](#) | [DEMO](#)

“PERSONALIZED AI FOR YOU”

[Google Gemini](#) | [DEMO](#)

“OS1 ONBOARDING SCENE”

[Spike Jonze film Her](#) | [FILM CLIP](#)

COMPANIES REFERENCED

[PARCHA](#), [PACTUM](#), [RABBIT](#), [OLLIE](#), [HUMANE](#), [CHARACTER.AI](#)



ACKNOWLEDGEMENTS

HUMAN

SALESFORCE FUTURES IS INDEBTED TO OUR SALESFORCE COLLEAGUES AND THE FOLLOWING INDIVIDUALS FOR THEIR DIRECT AND INDIRECT CONTRIBUTIONS TO THE CONVERSATIONS THAT INSPIRED AND SHAPED THIS WORK: SILVIO SAVARESE, KAT HOLMES, LIZ TRUDEAU, SHELBY HEINECKE, ADRIAN SMALLEY, PHIL MUI, ITAI ASSEO, PAULA GOLDMAN, KATHY BAXTER, YOAV SCHLESINGER, JAYESH GOVINDARAJAN, AND ADAM EVANS.

AI

PERSONAL AI MAY NOT HAVE FULLY ARRIVED (YET), BUT SALESFORCE FUTURES FOUND PLENTY OF USES FOR AI TOOLS IN THE PRODUCTION OF THIS MAGAZINE. IN OUR INITIAL DRAFTS, WHILE EXPLORING CONCEPTS FOR THIS MAGAZINE, WE USED MIDJOURNEY, ADOBE FIREFLY, AND DALL-E 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). FRONT AND BACK COVER ART: DALL-E.

IN OUR APPROACH TO THE FINAL VISUAL DESIGN PROCESS IN THIS ISSUE, WE FOUND THAT COLLAGES STRIKE THE BALANCE BETWEEN VARIOUS DIMENSIONS OF HUMAN INTELLIGENCE AND ARTIFICIAL INTELLIGENCE. CONNECTING TO THEMES ON PERSONAL AI, THERE IS A CONSCIOUS TREATMENT TO THREAD THE OCCASIONALLY INVISIBLE INTERACTION OF TECHNOLOGY WITH HUMAN BEHAVIOR, AS PERSONAL AI TAKES FORM. THE VISUAL STORY BEGINS WITH AN ABSTRACT DEPICTION OF PERSONAL AI, LEADING VIEWERS INTO IMAGINING FUTURES WHERE MORE DEFINITIVE VERSIONS OF PERSONAL AI EMERGE. OUR AI COVER ART IS A NOD TO THE COVER OF THE BOOK, THE ART OF THE LONG VIEW: PLANNING FOR THE FUTURE IN AN UNCERTAIN WORLD, BY PETER SCHWARTZ.



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



Tired of interruptions?

Turn on **Zen mode**

Ennen kuin aloitat käyttää Zen-tilaa, varmista, että olet saanut kaikki tarvittavat asennukset ja päivitykset. Zen-tila ei ole saatavilla kaikissa laitteissa. Lisätietoja Zen-tilasta löydät osoitteesta [support.apple.com/zenmode](#).