

Agentforce Keynote: Build the Future with AI Agents

Auto-transcribed by <https://aliceapp.ai> on Wednesday, 18 Sep 2024.

Synced media and text playback available on this page:

https://aliceapp.ai/recordings/lJeKdQu3FizqgXZwbD9MH_2xBkCPv5UA.

Words	9,080
Duration	00:53:30
Recorded on	Unknown date
Uploaded on	2024-09-18 16:56:06 UTC
At	Unknown location
Using	Uploaded to aliceapp.ai

Speakers:

Speaker A - 12.17%

Speaker B - 3.28%

Speaker C - 26.65%

Speaker D - 11.44%

Speaker E - 21.89%

Speaker F - 5.08%

Speaker G - 7.16%

Speaker H - 12.32%

Notes:

- Product AI platform Salesforce.com is launching. Adam Evans: This is an AI moment, and it's really at this moment that we think about the future. What we're launching, but more

importantly, what you do with it, is going to help change this moment for everyone.

- Kevin Quigley from Wiley. Thanks for being here. Can you just tell us about, like, you've launched service, uh, agent. Tell us a little bit about kind of the outcomes that you've seen.

- We've seen an over 40% increase in case resolution compared to our whole chatbot from the agent. We took our support experts, we took our existing CRM m team and, uh, our product support teams and empowered them. It makes it easier to handle more scenarios.

- Your AI is only as good as your data because your agents need access to trusted, accurate, and real time data. With Agent force, you can leverage a lot of your existing workflows and data to create these data aware prompts. It's all about bringing the right data to the right prompt at the right time.

- We have one of our very own asian blazers here in the room with us today. Take a moment to say hello to Erin and the easy cater team. Before we get into AI, can you share a little bit more about what ezcater does with the audience today?

- Agent Force is humans working together with AI across all kinds of channels. The reason engine grounded in your data with data cloud and policies and unstructured data. Head out to Moscone west on the second floor. Can't wait to see what you build.

Speaker A

00:00:00

Product AI platform Salesforce.

Speaker B

00:00:02

Adam Evans.

Speaker C

00:00:08

Welcome, um, to day two of Dreamforce. Is everybody having a good time? Are you guys ready for Dreamfest tonight? Pink magic dragons. Kai, go. It's gonna be a fun time. So wanted to say, well, first off, thank you. Thank you for getting up early. Thank you for waiting in line. People are coming in. Hopefully you had your coffee. Thank you for being our customer. Thank you for being our trailblazers. And soon our agent blazers. And if you could just step back. I mean, this is such a moment. It's a moment for us. And I don't mean necessarily just us in the room, but the bigger, uh, us. There's not many moments like this. 30 years ago, we had a moment kind of like this. It's this little thing. It's called the Internet. Came out, started slow, and then it went fast, connected us, made our world smaller. 15 years ago, we had another moment like this. We started carrying around these little devices in our hands. Not exactly. This one started slow and then it went fast. Changed the way we communicate, the way we operate, uh, with businesses, really transformed the social fabric of our lives. And today we collectively are going through another moment together. This is an AI moment, and it's really at this moment that we think about the future that we want, that we're going to create the trailblazers together in this room with

agent force. And what we're launching, but more importantly, not the technology, but ultimately what you do with it, is going to help change this moment for everyone. And for Salesforce, this moment started as a journey ten years ago. You've seen this with AI, the first wave, with predictive moving into a generative wave. And Mark did this slide yesterday. But I would like to do a little change and make it personal. You see that item at the very left here that says relate IQ? That's when I joined and that was my company that came into Salesforce's family, the Ohana, as we call it. Thank you. Recently, I'm a boomerang. We've come back into the generative wave. So I've had the pleasure of being in both of these waves. And I want to talk to you a little bit about the difference between the left and the right. Ten years ago, the then and the now, because ten years ago, at the beginning of this wave, we would build entire companies with great engineering and spend years on it to build features that would do things like analyze emails for sellers, tell you what the next steps are, and do more. And it was so much work to do that today. What's changed across this timeline is we can do something like that in a matter of days now with much fewer engineers, much fewer resources. And this is the power of the generative technology coming on. It is an entire collapse of companies and products into features. And it's incredibly interesting to think about, really a step function change. And during this time, companies have had an opportunity to grow. We've had an opportunity to increase our margins, drive productivity, and to completely transform our businesses. But yet, during this time, there's been this expectation always slightly ahead of us, always more growth, more margin, more productivity. There's always more work to do than there's time to do it. Simply put, there's not enough time in the day. So the question is, how do we close this gap? And that couple orders of magnitude opportunity I just spoke about in the last ten years, that's opportunity to do that. That is an inflection point waiting for us to take to ultimately close this gap. And, of course, the way that we're going to do this, to get more time back, complete more work, have more productivity, have more growth, is with agents. We are in the era of agents, and there's no better place to build agents than on agent force. It's everything that you need to build, to customize, to test, to deploy, to monitor once they're deployed, to improve. Full circle. Everything that you need built on the platform, um, that you've been using for decades connected to your customer 360, grounded in your data and more. Today, we're going to go through customer stories. We're going to talk about actual ROI. We're going to show you how to build agents and customize the ones that are out of the box. We're going to show you how to build agents from scratch. But before we do that, uh, I want to talk a little bit about what actually is an agent at the heart of it. You have a reasoning engine, and this is really the key. Unlock the ability to take context and actions and make a decision. But context, where does that come from? That's your data. That's your customer information, that's your policies, that's your instructions. This is your data connected with that engine. And ultimately, if it was just data, it could answer questions. But if we want to close that gap, we have to finish work, we have to complete tasks. We need to take action. So the combination of your data and the actions that are built on the platform with the reasoning engine, these three components are critical to creating an agent, and it unlocks so many use cases. Use cases

like a sales agent that can follow up with leads, qualify them and hand them off for you to human sellers. Use cases like service agents that can answer customer questions twenty four seven and deflect tier one tickets. Also things like transcription for all customer conversations. Thinking about adding, uh, tags to a customer's profile for a marketing campaign. Understanding things like product defects that can go back to supply chain. Also things about understanding business policies, even compliance as well. I'd like to drill into one of these to tell a story to, uh, make it more illustrative, something we're all familiar with. Um, order management, order issues. You go and you place an order online, you want to communicate with the company about maybe an issue. So what is the data and the actions for that? Well, the data may be your customer profile, who you are, it may be your order history, being able to understand the order details, maybe inventory company policies more. What are the actions? The actions are the ability to maybe look up a profile by email or by phone number, or maybe look up an order by order number, inventory by SKU. You get the idea. So let's take an example of a customer saying, I haven't seen my order yet. My order number is x. What will the agent do? It can take the order number, collect the data, and look at that. And depending on the answer, it gets the data itself. It changes the next step. It's delayed. Why is it delayed? Let's go check the inventory for that. Sku, is it out of stock? Depending on if it's out of stock or not, changes the next step. In this example, we're out of stock. So what do we do? Let's look at our policies. What are our policies when this happens? That now changes the response and the actions that the agent's going to take. So it's this idea of connecting actions and data with the agentic reasoning, step by step, reacting to the data as it comes in to navigate these choices. Why is that important? I'm giving you one example, just an illustration of something that's relatable, the story of something being delayed because it's out of stock. But in the real world, there's many, many variations of that. That same statement, I haven't seen my order yet. My order number is XDev could have played out in thousands of different permutations. That order could have been shipped and maybe the agent can check the tracking, live tracking, and it's actually in transit. And then the outcome would have been, your order's going to be there in two days. Sorry for the delay. That order could have been shipped. Maybe it was delivered. So that's perhaps a lost and stolen package issue. And we need to check our policies about how to handle that. Or maybe that order number doesn't even exist and the agent could figure that out and say, I'm sorry, did you make a typo? Is this the right order that you meant? The point is, is that there's a lot of complexity, there's a lot of variability, and the reasoning engine connects this. And, uh, I've been a programmer doing written code for 30 years. I can tell you the way that we used to solve this is that we have to go through every permutation, at least we try to. It's a lot of code. It's actually impossible. We have to make a flow chart for each thing. We have to make a dialogue flow for each of these. We have to know all of the questions that are going to be asked before they're asked to actually build this. This is the big unlock. Not only is it no code, but it lets you move across all of these paths because of the reasoning engine within the agent. Now, there's one scenario that I haven't highlighted here, which is what if the agent can't do it? And that's where the agent brings us in. It escalates to a human. That's a

big part of agent force. And really, I think that's what this really boils down to us, what do we want to do at this moment? How do we want to extend ourselves to augment ourselves, to be more productive, to help our customers? What kinds of jobs and tasks do we want to work on versus moving to the agents? We get to design the future. And it's going to be a really exciting time for all of us to go through this. How we bring humans and AI together, grounded in your data and connected to your actions, your flows, your apex. Thousands of connectors in Mulesoft. This is what agent force uniquely has, and these are the key elements that make up an agent. And today what we're going to do is we're going to drill through each one of these. We're going to show you customer stories, live demos and more. And for that, I'd like to welcome to the stage senior director of product management, Gary Brandelaire, to kick us off.

Speaker D

00:09:51

Thank you. Agent force is what AI was meant to be. It's a promise of CRM fulfilled. It is AI at the service of your company, your customers, and your employees. Now it's augmenting every single of your employees so that they can focus on high impact tasks. By removing repetitive work, we can now focus on the customers, which is the biggest importance. And of course, it's just not about productivity. It's about unleashing creativity, making sure that you can build the best customer experiences possible. Now, we learned a lot last year. It felt like a, uh, decade in AI terms. The pace of innovation is accelerating and with that comes also a lot of challenges. And there is one truth popping up, which is AI in enterprise isn't easy. Everybody is chasing the same dream. It's customer success powered by AI. The story though is that you want higher productivity, higher margin, instant answers. But the reality is that that's just the tip of the iceberg. Beneath the surface you'll face. Many companies are facing hidden costs, disconnected data complexities, evolving model. And the worst of it is that the industry is telling you, go and build it yourself. Build your own AI. But the reality is that by the time you're done doing your AI, it's already outdated. And there is a better way. Entertainment. Agentforce. Agentforce unites human AI data, uh, and CRM together so that you have everything you need to innovate without the headaches. We bring you data cloud so they can bring data from external and internal sources to feed the agent and make it knowledgeable. We give you the 360 degree profile so that humans and ah, AI work together toward the same goal. Customer success and agent force so that you can build and create your agent in a few clicks. And the beauty here is that it's built on the Salesforce platform. What does it mean? It means that you don't have to learn new technologies you can reuse flows, apex profile, permission sets to create and build your agents. And make sure that you control as well the access to data that you will have. Meet your agent force, your team of AI and agent working together, we pride you out of the box. Agent for service, for sales, for marketing, for commerce, so that you can deploy your agent in a few clicks and make them your own. And the reality here is that all of these agents have, uh, built in seamless handoff to humans when human oversight is needed. Now, you might be wondering, okay, but I have a chatbot. So how is

this different? And back to what Adam was saying, which is uh, this is a lot of dialogues. You need to think about all the conversation the customer is going to have. So you need to set up the intent, the dialogues, a lot of things to set up there. And it's never going to meet all your use cases. The copilots, they are assistive but they kind of lack productivity. They can take action. It's not really hard, not really easy for them to take actions for you. And then you have agent force, and you have agents that are autonomous, proactive. They can take actions on your behalf, of course, always bringing you when human oversight is needed. Just to give you an idea of the unlock here that Adam was speaking about, we had a customer with 140 intents in their chatbot. They moved to agents and got two topics only. And that agent was able to handle more conversation than their chatbots. That's the business simplification we are talking about here. Now, a picture of this across every single industries. You could have agent in retail, doing commerce, concierge, always on your website selling. You could have in telco, an agent doing customer support 24/7 reducing wait time. And the beauty here is that the possibilities are truly endless. You can configure and customize these agents at every level. At the agent level, the topic level, the action level, up to you to decide what they are going to do. It's not only across every industries, it's also across every app. We replaced Einstein copilot with Agentforce. It can now learn reason, take actions for you inside your own apps. Something extremely important to keep in mind. All these agents are built on the same exact architecture. So what does it mean? It means that you have a role which is defining the topics that the agent can assist with. You'll have instructions that the agent is going to follow. You'll feed some data so that the agent is knowledgeable about your business and relevant to your business. You assign actions so that they can take actions for you. But don't worry, there is guardrails, of course, so that you have defined limits. They will act within the limits you give them and they can interact across every single channel. All this built on the trust layer. Why? Because data safety is still number one. And the second part is that you want to guarantee the highest degree of quality for every single interaction. Now, let's think a little bit about customer success here. I was fortunate enough to do many, many pilots with many of our customers and learning with them, and we learned a ton, to be honest. Wiley was one of our first pilot customers, a leader in educational resource. And what they did there was like, there is a seasonal surge. I have three kids, you know, school starts, you need to buy courses and get smarter. So seasonal surge. And the story there is that they added to increase the number of agents to insert cases. But with agent force, they could literally answer 40% more cases by piloting agent force. So with that we are going to deep dive in the demo where you see how role action guardrail are ah combined together to make an agent that is going to answer Wiley's customer needs. And for that I'm going to bring on stage my favorite PM at Salesforce, Angel Ali.

Speaker A

00:16:28

All right, thank you, Gary. So Wiley is one of the world's leading publishers and one of the global leaders in research and education. So every year they see a huge surge in customer

requests when students are getting ready to go back to school. So that's why they deployed agent force. And agent force is grounded on Wiley's business policies. So it understands that when I ask if I can read a textbook on the plane, it semantically reasons through that request and knows that it needs to work offline. But it doesn't just answer questions, it can also take actions for me, like helping me reset my password. But what if I ask it something a little more complex like helping me add and update something into my orders? It's going to reason through that. But actually today this agent doesn't know how to do that yet. So it actually does what I want it to do, which is to deflect and redirect me to a human. But let me show you how you can actually customize an agent like Wiley's in our agent builder. So let's go into setup and look for our agent for service agent and open up our builder. An agent builder is where the customization magic happens. Here we have a couple of topics. And topics are uh, where we can let our agent know what the jobs that need to be done and the topics of conversation it should be able to handle. So let's take a look at our order management topic and customize it. Now, every topic has two elements, instructions and actions. And actions are crucial because they provide those guardrails of what the agent can and cannot do. So we already have a few instructions in here, like redirecting to a human representative if someone tries to update billing details. But look, my five year old, it needs a few more instructions to get the job done correctly. Let's add an instruction here that asks for confirmation details when someone tries to update an order. And then we'll do another one that um, asks you help with our order if they ask for order information. And let's go to the next step. So the second part is actions. And actions are so powerful because they allow our agents to act on our behalf. But what's really special about our actions is that they're built using the platform tools that your trailblazers have already been using, like flows apex and prompt templates. So let's take a look at one of our actions here and here. We are in flow builder and you might be wondering, why are we in flows? Well, because your trailblazers have been spending years putting things into flows, your business rules, your processes, and we want to make sure that you can reuse them to be really prescriptive to your agent so that it's pulling the right order information to ensure a seamless customer experience. So let's go back into agent builder, add those two actions and click finish. And just like that, we were able to add in new instructions and actions to a topic without writing a single line of code or creating any complex dialogue trees. So let's head over to preview and test some of these user requests or utterances. And what you're going to see here is our reasoning engine reasoning through this request. So what are we looking at? Well, this is the result of our reasoning engine. And don't be fooled. This is not a dialogue tree. This is the reasoning engine. Reasoning thinking grounding through our request to pick the right topic and to select the right action in order to complete the request. So let's go back to the original request to update our order. And here we are again. We selected the right, we're going to go and look for it to select the right topic. And then now we're actually getting an update order action here. And just like that, with the simple change in our instructions and actions, we're now able to have our agent do this job for us. But let's test and see if our guardrails worked and try to update payment method and see if I can trick the agent to doing something it's not

supposed to do. And remember that instruction we had to not update billing information. Look at our reasoning here. It correctly states that it can't do this update for us because we provided explicit instructions to prohibit it, the agent from taking billing information. So now it's redirecting us to a human. So I can continue to go and preview and test, but that's going to be really hard to scale. So let's go back to the request we had and create a test. And here we have our new batch, uh, testing center. And we already see that the request we put in here has already passed. But let's go ahead and create a test that I know is going to fail. Something that has nothing to do with order management, which is telling me a joke and what I'm looking for is it to fail, which it does. And if I wanted to have order management tell me jokes, I can take a recommendation and update that topic. So let's go ahead and delete that because it has nothing to do with order management. And while it would be great to go line by line with you all thousands of people updating utterances, I'm actually not really good at updating utterances on the fly. But you know, what is LLMs? So now with the simple prompt, I can ask the LLM to generate, uh, utterances for me. And what it's doing is creating many permutations of order management related utterances. And what I'm looking for is a pass. And wow, look at that. All greens. And so now I can, I feel very confident that I can activate this topic, and it's not going to work just once, but thousands of times. So let's activate this topic because it's ready for prime time. So now I'm back on the support page, and let's try that utterance again because I really need that microbiology book to start learning about yeast. And bam, here we are, able to update that order. So with agent builder, I was able to quickly create a new topic, have a new job to be done, test it at scale, and bring new functionality to my users. Back to you, Adam.

Speaker C

00:23:15

Adam. Thank you, Angela. That's amazing. Love the test center. Being able to test things out, AI, helping you launch AI faster. Um, great. Well, we actually have Kevin Quigley here from Wiley. I'd love to ask you a couple questions. Thanks for being here. Thanks for being a customer. All right, so, Kevin, um, can you just tell us about, like, you've launched service, uh, agent, tell us a little bit about kind of the outcomes that you've seen.

Speaker B

00:23:39

Yeah, sure. So I'd say there's two really amazing things that we've seen. One is we were able to expand the self service topics that we're covering compared to our old chatbot, from the most important and common issues that we were getting to virtually every question that can be answered by our knowledge base for that product. So that's really cool. And we're also now able to give personalized dynamic, not canned responses to those inquiries when they're coming in on the customer's own terms.

Speaker C

00:24:05

Awesome. So, bringing in knowledge, very quick, uh, personalized answers, talk a little bit

about, I mean, we were working together on this, and I think we kind of share maybe a little bit of the story, but what was it like to actually set this up? What was different? Who was involved. I'm sure everybody's thinking, how hard is this to deploy?

Speaker B

00:24:22

Yeah, sure. So there's something really interesting that's happening with agents, which is that you're shifting that conversational experience design from being this sort of tedious back and forth collaboration between your support experts and your IT resources, your CRM and your programmers. And now there's a bigger focus on the conversation just by those support and product experts. So you're taking the people who know what a good customer experience should look like and they're able to tell the agent how to perform that experience. And so that's much more efficient, you can imagine, on the design side, than having to chart out that logic. And like I said before, it makes it easier to handle more scenarios because you're not having to define every single branch of a conversation tree. You're just able to say, this is how we want to support this area of our product.

Speaker C

00:25:08

So you didn't need like an army of PhDs and cs degree folks to have this?

Speaker B

00:25:12

No, no. AI specialists. We took our support experts, we took our existing CRM m team and, uh, our product support teams and we empowered them. They came together and they said, this is what the experience should be. And to a good result. We've seen an over 40% increase in case resolution compared to our whole chatbot from the agent.

Speaker C

00:25:30

Love it. Awesome. Kevin, thank you so much. Wiley team, thank you so much for being a great customer. Thank you. Congratulations. All right, and so speaking of knowledge and being able to have great results with Wiley, that brings us to our next chapter, which is data. Avanthka, would you like to tell us about data?

Speaker E

00:25:46

Of course, Adam. Thank you. Well, isn't it quite amazing how Agentforce is able to orchestrate everything from workflows and data all the way to prompts and models to create this truly autonomous agent experience. Now, after speaking with many of you in this room, here's the question we're all asking. How can we trust these agents to operate autonomously and deliver real business value? Well, you might have heard us say that your AI is only as good as your data because your agents need access to trusted, accurate, and real time data. Now, in terms of, uh, how you do this, it's quite easy to say that you need to power your agents with data, but how do you actually do this? Now, a year ago,

many of us thought that the answer to this was to train a custom model on all our business data. But soon enough we realized that this isn't actually the most efficient approach, especially from a cost and time investment perspective. These models quickly went out of date, especially as your data kept evolving and growing. Now, at, uh, the most basic level, these AI agents work by sending a series of prompts to a large language model, which then generates a response. So here's where the real magic happens. It's all about bringing the right data to the right prompt at the right time. Now, with Agent force, you can actually leverage a lot of your existing business workflows and data to be able to create these data aware prompts. For example, we can start by pulling structured data within your system. Think flows, related lists, merge fields. All of this already lives in your system. But let's take this a step further. You all have, uh, vast troves of unstructured data living in your enterprise. Think emails, cases, conversation history. You can't just pull all that data directly into a prompt, but at the same time, you can't afford to ignore that data either. So this is where the power of retrieval, augmented generation comes in, or what we call rag, in order for you to bring the right unstructured data into your prompts. Now, with data cloud, you actually have the ability to bring all the structured and unstructured data into a single system. And from there, your agent, at the moment it needs that data, can easily search and retrieve the right data it needs. So we offer a variety of techniques for search, whether it's vector search, hybrid search, or even search based on knowledge graphs. Now, in terms of search, we talked about how important it is to be able to bring in the right data at the right time. So, as you can see here, data isn't just an accessory, it's the fuel for what powers Agent force's next generation reasoning engine. Now, let's talk about, uh, how this all works. In order to build these context aware agents that are able to reason on your data, you also need a good prompts. Prompts are what really power this. Now, many of you love promptBuilder. Thousands of your own users are already using promptBuilder today for various use cases. Think content generation record summaries. A lot of our users are reaping huge productivity gains from using promptbuilder, and we're excited to see how you all continue building new use cases. Now, speaking of customers, don't just hear it from me. We actually have one of our very own asian blazers here in the room with us today. So let's go ahead and say hello to the easy cater team, Erin. Hello. Hello. Well, I want everyone to take a moment to say hello to Erin and the easy cater team. Let's wave camera. All right, Erin. Well, we're here to talk about AI, but before we get into AI, can you share a little bit more about, uh, what ezcater does with the audience today?

Speaker F

00:30:01

Yeah. So I imagine everybody here knows how easy it is to place an order for food or friends and family on a Friday night. But if you have to place a business order, suddenly you don't know exactly how much food to order, you don't know people's food allergies. And how do you make certain that, uh, the delivery driver is going to get past security? Nothing easy about that. So that's where the easy tech food platform comes in. We're making it simple for organizations to manage their food needs.

That's amazing. Well, we all talked about AI and how critical it is to powering your business. Tell me a little bit more about some of the use cases and opportunities you see AI really driving value and how asian forest is critical to that journey.

Speaker F

00:30:39

Yes. So in recent years, we've been pointing AI and ML on our, uh, 17 years of catering data. So with a few simple inputs from, from a customer, we can generate a novel order that also honors the restaurant's menu. So if you wanted to place an order for a team of 17, but that restaurant's lasagna tray only serves twelve, we can auto generate what's the right additions to that order. Where we're going now, with the help of Salesforce, is solving what we're calling the discovery phase. So today our solution really depends on the customer deciding the restaurant of their choice. Or they can opt to talk to what we call our beyond helpful human customer service agents that have the full business catering ontology in their heads. So when customers use terms like I want something gourmet or seasonal, or phrases like, can you order me something similar to what I had last month? The agents can take this unstructured set of prompts and decide what the customer needs. So as we scale and need to handle more volume, we need to augment our human agents with a solution that's going to give the customer that same sense of talking to an expert. And so that's where the salesforce solution comes in. The prompt builder is allowing us to structure these set of grounded prompts that can think about the context not just of food descriptions, but also the customer's order history or the location history. You don't want to be the salesperson sending your client pizza for the fifth day in a row. And so we have the confidence not only is it understanding that grounded context but using modern rag techniques is going to validate that against data that's really important. Many moons ago in our early experimentations with AI, we had a model decide an interesting recommendation was spaghetti with your side of french toast for a breakfast meeting. So we don't want that. So we trust that the salesforce solution is going to give the customer really a well reasoned response and know when to hand it back to a human agent, which is super important to our customer service.

Speaker E

00:32:32

Wow, Erin, what a wonderful use case. I love your team's innovation and especially thank you for being a pioneer and believing in agent force. Well, you all just heard it directly from Erin. So what are we waiting for? Let's see this in action. So up, uh, next we're going to show you how easycater is using prompts, data search and the power of actions to power what Aaron just said. The next generation customer experience on easycater. So let's dive in. Now here I am on easycater's site when I'm hosting large team onsite, I know I can depend on ezcater to help me, especially with my complicated dietary preferences. Now I log into the site and a rep from easycater is available to me to help me with my dietary preferences. Now, as Erin mentioned earlier, their reps probably get thousands of these types of questions every day. So how can they help their business scale but yet provide

this kind of white glove experience that they're providing on their site today? Well, we just learned about how good prompts result in great agents. So we're going to start our journey in prompt builder. So here we are in Promptbuilder and I'm going to create an instruction to provide catering recommendations based on the user's query. So when a query comes in this variable called the input request, that's what dynamically captures what the user is asking for. Now we get the user's requests and we want to be able to answer their question, but it doesn't stop there. We want to know a little bit more about the company the user is coming from along with some data about their past orders. Now this resource picker right here is your access point to all your data in Salesforce. So from here I'm going to go ahead and click on the account. I can bring in a record snapshot. This is essentially a snapshot of all the information we have on the account record. Now we can go beyond this surface level data and actually use some existing flows that we have in our system to fetch some of their past order history. So now we inserted our flow to get the order history. Now let's take it a step further. We know that we have enough context about the customer, but what about easy caterers data? What about all the information about their caterers, the menu options, the customer reviews? All this data lives outside of Salesforce. And obviously I can't show you easy cater's database, but it might look something like this where they have thousands of rows of unstructured data. As you can see, they have menu descriptions, caterer names. I mean, look at that long menu description right there. So how can we make sure we're getting just the right information we need at the right time? Well, with data cloud, you can actually bring that information in. So we have all this information that we just saw stored in a data model object. That many description that we saw earlier is brought into data cloud. Now we want to be able to search this data, and to do that, I'm going to build a search index. It starts by determining what fields of data we want to search based on. And from there we're going to actually segment this data into smaller chunks, what we call chunking. Now, once we chunk the data, we want to be able to create numerical representations of this information called vectors. And for that we're going to use an embedding model. Now, uh, I've used this embedding model. I create those vectors, and now I'm able to store it in Datacloud's vector database. All right, so I've done everything I need to create that search index. I'm back here in promptbuilder and I'm able to access that search index I just created with all the caterer information. And now I need to determine how I want to retrieve that information. First we input the search query. Now, you'll notice that this is very dynamic in nature. This is almost like a prompt of its own because this is what dynamically captures what the user is asking for and passes that into the search. And from there I want to determine what fields of information I want to retrieve. This includes the customer review the menu options, and I'm good to go. So I have all the data I need. The last thing I'm going to do is insert a few more formatting instructions, stylistic considerations. I've, uh, built my prompt. Now, the beauty of promptbuilder is you can even test your prompts right here. So I'm going to be able to test this on a sample account record and enter the same query I asked earlier finally, we can just select a model. So I have a whole host of models I can choose from. These are all generic models not trained on any of easy cater's

data. And here I am ready to go. So let's go ahead and give that a test. And here we have it. Now let's actually break down what we're looking at right here. Let's zoom into that resolution. What we've done is we've actually resolved all the different data references you saw earlier in that prompt template in this resolution. So first we had a reference to a record snapshot. So we brought all the information that's on that account record right here. From there we also brought in some of their past order history through the flow. So all those results are brought in as well. And finally, what about those catering options? Remember we had that large, uh, unstructured database of information? Well, we use search to bring in just the relevant results for the query. And in this case I asked about vegan restaurant recommendations. And look at that. Our search was able to bring in just the right results. I mean that first one with vegan table, a 94% similarity score. So that's how search is able to bring in just the right information. And we have all this data packed into the prompt that we send to the large language model. Let's go ahead and send that to the model and let's see what we get. Here we have it. Our response was able to synthesize all that information and give us right what we needed, which were the three catering recommendations. Now all I need to do next is attach this prompt to my agent as another action. And let's see what the easy catering experience looks like now with Agentforce. So I'm going to ask the same question. And there we have it. We got our three recommendations from our model. Now we know that these agents go beyond simple question answering. We can even have these agents take actions on our behalf. So how about we ask this agent to place an order for my team? And there we go. Now what happened behind the scenes? We were able to invoke Easycater's proprietary API to place this order. They have their own API which was also available as an app action for our agent. And there we go. I'm going to go ahead and submit my order and I can rest easy that an awesome catering order is on the way for my team onsite. And that, my friends, is how easy cater is able to power this white glove customer experience using the power of prompts, data search and overall agent force. So thank you so much for following along.

Speaker C

00:39:42

Awesome job, Ivanka. Thank you. That was such a great display of how accurate data with good prompts creates great AI. And there was one thing that you showed right at the end. So I want to highlight that being able to transact with the order that was connecting previous investments and APIs that were already built before the agent to the agent so that the agent could take advantage of that. And that's what we call actions. Which brings us to our last chapter here with VP of software Engineering, Claire Chang. Take us home.

Speaker G

00:40:14

Thank you, Adam. Um, it's great to be here with you all. When we talk about agent force brings together humans with AI, data and actions. The true value here is about getting work done. Um, agents being able to take actions is what empowers employees and drives customer success. How does Agent force extend your workforce? What does Agent force

provide to make your teams more successful? Give them more time and help get to the jobs that they can now complete every day. First of all, Agent Force is giving you access to a collection of highly capable, out of the box autonomous agents. As you heard from Gary earlier, we have built multiple agents to complete the common jobs across customer 360 to help your service teams, sales teams, marketers and sellers. You can enable them easily and customize them, um, to fit your business and drive your success. Wine. Agentforce can assist you and take actions on your behalf. At the center of it is Atlas reasoning engine. As you've learned from Adam earlier, this is the brain behind Agentforce which simulates how humans think, plan and take actions. It starts from evaluating user's request and the context, refining them for clarity. From there, it retrieves the most relevant data from your business, your CRM and your data cloud. This is super critical to make sure the plan is formed and tailored to solve your business problem. The process keeps evaluating and uh, refining the plan, making sure only the right and relevant actions are going to be invoked for execution. With this extremely powerful reasoning engine, you now have the ability to get almost any job done. If you can describe it, Agentforce can do it. It means you can customize out of the box agents or uh, build your new agent. All powered by Agentforce with agent Builder. As you will see from the earlier demo, in agent Builder, you can create a job to be done for your agents by describing the topics, writing natural language instructions and creating a library of actions using the tools that you already have today. In Salesforce, you can use flows, Apex, Mulesoft APIs and the prompt templates that many of you have been already using in the past year. In Agent Builder you can also observe how agents reason, plan and take actions. You can evaluate the results and refine instructions for further improvement. Since Agentforce is a native part of Salesforce platform, all of your data, all of your customer 360 apps and all of your business logics and automations that you have already built can be leveraged by the agents. Because Agent Forest is deeply integrated with your apps and with your data, we now transforming the work can be done across every role, workflow and industry. We are truly bringing together humans and AI agents to drive customer success and we can scale it up like never before. And of course, Agentforce platform is extensible and is integrated with our amazing partner ecosystem. You can extend your workforce by bringing a partner agent or you can use a partner action and add it to your existing agent. Or you can bring any data, uh, with our zero copy data network. So now you have a collection of out of the box agents and you have all the tools you need to customize and build your own agents. How would you employ those AI agents in your workforce and ask them to take actions on your behalf? Well, with Agentforce platform we make it easy for you. You can service the agents in any of your existing apps across all your conversations and even integrate with your automation flow. Now we would like to show you how you can build your agent and lighten up in your organization with that. Please welcome Carlos Lozano.

Speaker H

00:45:09

Thank you, Claire. Good morning everyone. Let's build a custom agent together. You guys

up to it? Following the anatomy of an agent that you see on the screen. And let's start right here in the agent setup. So as I click new notice that agent force is giving me the flexibility to start with a pre built agent for service, for sales, for marketing, for commerce. You get the idea. But today let's start an agent from scratch. If you can describe it, Agent Force can do it. So let's go ahead and describe the purpose and the goal of this agent. And let's do an agent that's going to handle something that we all like to do every time we go on a business trip, which is our travel expenses this next Monday or in a few days when you guys go back home, we're all going to do this. Wouldn't it be nice that we have actually a custom agent that can handle that? We can give it all of our expenditures and can handle that for us. Let's go ahead and do that. So let's go ahead and give it a purpose. So it's going to be a travel a travel expense agent that's going to be tracking here expenditures, reimbursements and making sure it follows very important the company policy. Now watch what happens next. Based on this natural language instructions, we're going to see what agent force is actually bubbling up, what it's doing here, what's actually happening here. It's doing a semantic search. But what does that even mean? Based upon the description we provided in the previous step of the wizard, it's actually able to find semantically similar salesforce resources that you already have in your, we have a number of actions here and it was able to infer in which channels given the descriptions we might be able to deploy it. Let's take a quicker look inside of this topic suggestion. These are your assets and resources you have already built in the past. And because you have properly described them, descriptions are very important for these assets. The system is able to bubble them up. Now these have high similarity and that's looking great. I also want to actually surface and deploy my agent to allow it to send notifications on slack. So I'm going to go ahead and click on that as we go to the next step. Based on the descriptions you have provided to your resources, to your flows, to your apex, to your prompt templates, it is also able to infer this topic's instructions. Remember, as part of the anatomy of a custom agent, these instructions act like the rules that the agent needs to follow. How is it going to use these actions? What should it do and what should it not do? Fantastic. I can edit them at any step of the way. For today, let's just go ahead and move on. Where do we want this agent to be deployed? It has pre populated this channel selected for the sake of today. I'm going to unselect, um, email. I want it to basically, um, give us an output on slack and allow my users to engage with it. But also very important agents can act in the background headlessly without you even noticing. And how are we going to trigger those headless agents? We're going to use a mechanism that you guys are familiar with. In this example, we're going to use flow. So let's save this. Let's move on to the next step, which is hooking it to data. Very important. This is what's going to give it context. This is what's going to ground the agent and allow it to essentially provide these answers in respecting my company policy. These are things that are available in your, this is structured data available in your budget payments knowledge, but it can also be unstructured data that you decide, for example, to upload in a PDF file, perhaps a specific policy for a specific region, for example. Now this is looking great. I'm going to move up to the next step for a recap. We gave it a name so it has a user. And that user, you will give it the access and the

permissions that you want. We gave it the agency to act on our behalf through topics and actions. We hooked it up to data and we decided in which surfaces we're going to deploy it. The only thing in the last part of this wizard is to create the agent. And where do we land? We land in the agent builder. Great. We love the agent builder. Final step is we're going to create an agent trigger so that we can essentially use the vehicles that are used to invoke, that you are used to using to invoke processes can point to this agent trigger. So it's going to auto suggest here a, uh, name for my trigger. What is it that I want to pass to this trigger? I'm actually going to pass that very travel expense and I'll withhold it in this record id. When do I want my agent to stop doing the work? When will it exit that reasoning loop, if you will? When we have an expense that is approved or rejected. Let's go ahead and apply this, activate this agent, and now it is ready to be invoked. And one last thing, we're actually going to invoke it. Every time a travel expense record is created. You know this, I'm in flow builder. Every time a travel expense record is created, I'm actually fetching it using this get element. I have stored the record in, uh, a variable, the record id, I just called it. And now we're basically going to call the invocable action that we have created for the agent. But it's fetching all of the parameters that I previously created in the agent builder. Let's go ahead and activate our flow and now we can run this simulation. This is the travel expense that I have created. I have provided the information, I have attached the unstructured data like receipts and essentially the expenditures that I have had. And as I submit this, the reasoning engine is going to kick in, is going to start assessing if it's going to approve or reject this travel expense. So let's take a look at the outcome. Oh, miraculous. Good news for me. My recent travel expense has been approved. I can go to, I can go to Dreamforce. And thank you very much for your attention. Back to you, Adam.

Speaker C

00:51:41

Thank you, Carlos. Thank you Claire. And thank you, Michael, our demo driver here as well. Thank you, Michael. All right, so what you have seen is something that only Agent force can do. This is humans working together with AI across all kinds of channels. The reason engine grounded in your data with data cloud and policies and unstructured data and ultimately activated through flows. Apex mulesoft more. This is what Agent force is. And when are you going to get it? Turns out most of the things you've seen today are ga ing in a one month from now, right? This is the out of the box agent sales, Agent service, Agent marketing, commerce. This is also the reasoning engine, which powers it all, and the builder that we've been demoing. Uh, we're not stopping there. In our spring release in February, this is where we saw that really awesome batch testing center coming. Additionally, more developments in advanced search multimodality as well image processing and voice support. Agent Force is piloting right now, and if you can't wait to get started, head out to Moscone west on the second floor. If you haven't seen it, it's awesome. We built 500 agents with you in 30 minutes sessions. By the way, since we started yesterday at 08:00 a.m. you can get your hands on this now and launch agents in

just a few minutes that are personalized to you. That's the launch zone. Additionally, while you're there, stick around. On the third floor, we have the mulesoft platform, uh, or the keynote, which is going to help you integrate it with all actions across everything. And then the platform, which is going to take it to the next level. Building multimodality and agents beyond CRM use cases with the agent builder. Last but not least. Thank you. Have a great dreamforce. Can't wait to see what you build. Thank you.

Speaker E

00:53:29

Hey, trailblazers.