What's new in Prompt Builder

Auto-transcribed by https://aliceapp.ai on Wednesday, 18 Sep 2024.
Synced media and text playback available on this page:
https://aliceapp.ai/recordings/BexvSbFoQRkNliWUOZaJC1gzK4jEG7df.

Words	3,206
Duration	00:18:15
Recorded on	Unknown date
Uploaded on	2024-09-18 17:55:13 UTC
At	Unknown location
Using	Uploaded to aliceapp.ai

Speakers:

Speaker A - 55.27%

Speaker B - 44.73%

Notes:

- Welcome to what's new with Promptbuilder. Salesforce is a publicly traded company, so please make your purchasing decisions based on what is publicly available today. We hope you're having an amazing first day with many more to come.

- PromptBuilder is focused on data, trust and availability. We are constantly striving to improve prompt accuracy. Today I'm going to show you two demos, one with our new data, uh, provider record snapshot, and the other Einstein search.

- A record snapshot is a combination of all of the relevant fields on an object on the running users page. It will combine all the fields that are most relevant. Only the fields the user has access to will be included and sent to the large language model.

- The next thing I want to talk to you about is Einstein search. Einstein search is retrieval augmented generation. What this means is the retrievers that you set up in data cloud can now be used in prop builder. For this demo, let's pretend that we've created a wine recommender app.

- Field based masking uses metadata in salesforce fields that you specify an object manager to mask. Pattern based uses a combination of patterns, regular expressions, and machine learning models to automatically detect certain fields. This is really, really exciting because you all now have more opportunities to protect your data through the trust layer in prom builder.

- All right, good there. Please feel free to take a picture of this as well if you want to go deeper, dives into how promptbuilder works. So we look forward to seeing you at some more of these events, uh, in the next couple of days.

Speaker A   00:00:00

Welcome to what's new with Promptbuilder. Before we get started, just a gentle reminder that Salesforce is a publicly traded company, so please make your purchasing decisions based on what is publicly available today. First, a huge thank you. Thank you for coming to Dreamforce. Thank you for coming to our session. We hope you're having an amazing first day with many more to come. Now, I'm Liz. I'm an associate product manager at Salesforce.

Speaker B   00:00:30

And I'm Sid. I'm also an associate product manager here at Salesforce.

Speaker A   00:00:34

And today we're here to talk to you about PromptBuilder. With PromptBuilder, you can embed generative experiences in the flow of work through prompt templates. With Promptbuilder, we really kept ease of use in mind. So with just a few clicks you can customize Salesforce turnkey solutions or you can build your own new prompt templates, all of this through no code. We also understand that augmenting your prompts with data is really important, so we've made it really easy in Promptbuilder. You all have spent many years gathering your business data and you want to use that in your prompts to enhance them. So we allow you to connect data from CRM, from data cloud, from external sources like Mulesoft APIs and Apex, with no model training required either. And all of this can be embedded in the flow of your work. So you build prompts once and they're extensible, they're reusable, so you can use them in our new agents enlightening web components in flows and more. And we're going to talk to you about all the amazing new features that we've released. But before I do, I'm going to pass it to sid to talk to you a little bit more about how promptbuilder works.

Speaker B   00:01:52

Awesome, Liz, thanks for that. Now let's dive into a little bit about how promptbuilder works in detail. So you always want to start with the business task. At the end of the day,

as you can see on the screen, this is generating a summary or sending an email, drafting email, whatever it might be. This underpins all the prompts. And this is what is really important to start with. Next, you have two options. You can either customize out of the box prompt templates, or you can choose to make your own from scratch with the available prompt template types that we have. After that, it is really critical to ground with the right data. Data provides relevance, data provides the context, and lastly, data provides accuracy so your end user doesn't have to go back in and do extra overhead work to make sure that the response was accurate. Now with all that being said, and once that's all packaged together, it's time for the prompt to communicate with the LLMDH. All of this is done with our uh, Einstein trust layer. So whether it's Salesforce managed models like OpenAI's chat GPT, or if it's anthropics cloud, or even bring your own models into Salesforce, all of this goes through the Einstein Trust layer. Now this is to reduce toxicity biases, any harmful content that might seep through and influence the prompt response negatively. Now once that's all said and done, it's time to go and deploy this into the flow of work. As you can see on the screen, there are a few different ways to do this, whether it's in custom Salesforce 360 such as content Connect API, or if it's invocable actions such as flow Apex, Lightning, web components, and last but not least, our beloved agent force. Whereas automated agents or assisted agents, all of these have prompts underlying everything that they do. So now let's double click into how you actually get started using prompt templates. So one thing that we wanted to design for is to make this experience as easy to use. So what we have is out of the box prompt templates. What these are are really easy to use, efficient, well thought out, generated with grounded with data prompt templates that are designed for specific use cases. So when you're sitting and looking at a blank screen, getting ready to start writing your prompt, you can start out with the prompt template. As you can see on the screen you have sales emails for sales meeting follow up emails for sales. You have summaries for cases for service, and many more industry prompt templates out of the box templates that will be coming in the next few months, such as financial, healthcare, insurance and many more. Really what this is is a launchpad for you to get started using PromptBuilder so you don't have to put in that extra overhead to come up with that initial first draft. I'll send it over to Liz to now talk about what's new with PromptBuilder.

Speaker A   00:04:24

So what have we been working on? We're really excited to announce all of these new features that have either recently been released or being released in the next few months. These are organized by investment themes. So in our product we've been really focused on the following three areas, data, trust and availability. We're going to show you demos in all three of these areas, but first, let's start with data. With data. We are constantly striving to improve prompt accuracy by providing you all more options to insert your data in prompts, whether that's through semantic search, more data cloud connections through data model objects, or data graphs through flows, and more. So today I'm going

to show you two demos, one with our new data, uh, provider record snapshot, and the other Einstein search, which is retrieval augmented generation. Since I'm talking about it, let me show you. So here we have a prompt template. This is a lead outreach template that generates an email. With an email, you have a sender and you have a recipient. And I want to ground my template, um, with relevant information about both. So I said I'm sending the sender name, sender title, the recipient name, the recipient title. And the way that I insert these fields is by going into the resource picker and clicking through all of these fields. And sometimes it's helpful to choose a specific field that you want, but other times you want more of a snapshot of the entire object. And that's where record snapshot comes in. Instead of having to manually insert each field through one click, you can create a record snapshot, and it's at the top of the resource picker. And what a record snapshot is, is it is a combination of all of the relevant fields on this object on the running users page. So when you click this, it will combine all of the fields that are most relevant. So let's see it in action. If I click a recipient, Sara, and I click preview, what the record snapshot is, is it's really a JSON object of all of these fields. Let's let it load. Um, and so what you can see here is the record snapshot for the sender includes things like full name, it includes title, company name, et cetera. And that data is used in the response, as, uh, you can see here. One thing I want to point out is this is based on the running user's permissions, so only the fields that the user has access to will be included and sent to the large language model. So we honor your permissions. The next thing I want to talk to you about is Einstein search. So like I said, einstein search is retrieval augmented generation. What this means is the retrievers that you set up in data cloud can now be used in prop builder. So for this demo, let's pretend that we've created a wine recommender app. Uh, we all like wine here. What this is going to do is it's going to take a user's review and recommend wines based on my company's products. And that data is stored in data cloud. So let's say that I like white wines that are light, crisp and refreshing with flavors of pear, apple and sometimes a hint of spice. When I click find my wine, what's happening in the background is a screen flow has started and what this is doing is it's taking that user input and it's passing it into the template called wine search and it's invoking that and then the response is what's being generated on the output. So if I go back, what you can see here is the recommendations, uh, are shown for three wines. And if you read the descriptions, you can see, uh, how it was generated with that semantic search. You can see white wine in some of these descriptions. Light and clean apple flavors, a hint of bitter almond to have the spice. And that's really exciting. But I want to show you the template that's powering this. So this is the template that was invoked here. You can see that I've written to the large language model that you're going to get a wine tasting and I'm asking it to create compelling wine recommendation based on those similar wines. For each of the wines that the retriever returns, I'm telling the large language model that you're going to get the following six fields. The wine review is sent to the prompt template through this input that we're calling description. And this is the new Einstein search retriever. When I click on this, the side panel will pop up where you as a prompt author can configure it. So you can see that this wine retriever is associated with the data model object wines. The search text

that's actually being queried for the retriever is that input which is the wine review that I gave. I can specify the six output fields that I want the retriever to return as well as the number of results. And so when I click test inputs and I enter the same description here as I did before and I click preview, it's going to invoke that retriever and get the result back. And the way that you add Einstein search is there's this new field in the resource picker. And if I look into the resolution and see what the retriever returns, it's this search result object. And you can see for every wine, it includes the fields I mentioned above, like title, winery, variety and more. And on the right you can see the response that's generated here. So this is really powerful because you can now use all of the Jado model objects and retrievers that you've set up in data cloud. In prop builder, the second big investment area is trust. As, uh, you know, trust is our number one value. And so we are always striving to enhance trust. In promptBuilder, these features, such as data masking, toxicity, and citations are either all released or coming out very soon. And today we're going to talk to you about two of them, field based masking and performance metrics. Now, I'm sure you all have heard of the Einstein trust layer. Pattern based masking in promptbuilder has already been released. And pattern based uses a combination of patterns, regular expressions, and machine learning models to automatically detect certain fields. So if I scroll down, you can see here, it will automatically detect name, email address, phone number, and anything that you specify here. But we heard from customers that pattern based might not be enough. Let's take a car company, for example. They have car objects and they have vin fields on those car objects. They want to mask that specific field specific to their company and their organization. That's where field based masking comes in. So field based masking uses metadata in salesforce fields that you specify an object manager to mask. And it's more predictable than pattern based masking because, you know, and you set that setting in object manager. And so the way that it works is you can classify those fields in object manager and then specify the compliance category, the data sensitivity level that you want masked, to be masked. So an object manager, you can click on the object and then go into the specific fields. And all of this is shown in promptBuilder. So this template, as you can see here, has a bunch of sensitive information that we've included on purpose. There's a credit card number in there, there's a phone number, there's even an SSN number. And you can also see certain, uh, fields, such as account sensitive, account field that I made custom, and this related list of account contacts. So what I've done in object manager is I've said that that field is sensitive, and certain fields on my contact object are also sensitive. So when I click Acme and I preview this, what you're going to see in both the resolution and the response is those fields are masked. You can see person underscore two, for example, as the placeholder. So instead of us sending Thomas a. Johnson to the large language model, we are sending a placeholder. And if you want to dive deeper, you can even click on this and see exactly, exactly what was masked and why it was masked. So you see the placeholder, the true value, and either field based or pattern based. This is really, really exciting because you all now have more opportunities to protect your data through the trust layer in prom builder. I'm going to pass it back to Sid to talk about performance metrics.

Thanks for that, Liz. So tying onto one of our core values, which is trust, we releasing a new feature called Prompt Performance Metrics, which will be gaining in October of this year. So what we've heard from a lot of our customers is oftentimes they'll go into prompt builder admins, they'll set up a prompt and then have their end users use it in their flow of work. But oftentimes they don't have feedback on hey, do their end users like it? Do they dislike it and not have that loop for that iterative feedback process? So we have in prompt performance metrics, as you can see on the screen, the end user can give a thumbs up or a thumbs down every time they use a prompt for each unique user. And what happens here is all of this data is now aggregated in data cloud and resurfaced within PromptBuilder. As you can see on the screen. If I look in the list view from all my prompt templates, I can see a prompt, and I can see how many people liked it and how many people disliked it. This gives me just a snapshot into how can I iterate on this? Should I choose to deploy certain prompts more and certain prompts less? So, this ties back into a core theme of trust to really allow you to understand how your end users are using this product. So now we'll touch on the last investment area here, which is availability. Our goal is to make promptbuilder as accessible to people around the world as we can, since we have customers coming from everywhere. So we want to increase locales, languages, models, and even ISv packaging within this investment area. But today I'm going to talk about expanded localization as a feature. Right now, within promptbuilder, we have ten different languages we support. What does that mean? So whatever language the end user's org is configured in, if it's one of those ten languages, it doesn't matter what language the prompt is written in, because the output of that prompt will show up in their language. And so, as of month ago, four of those ten languages we just launched Mexican, Spanish, Brazilian, Portuguese, Swedish, and Dutch. Today, I'm going to show the end to end flow for how it looks like with Mexican, Spanish. So let's get right into it. So, as you can see, here we're going to be a seller whose is configured in Mexican Spanish. Just give it a second to load here. So, as you can see, everything here is configured in Mexican Spanish. That's the language the seller is using. Let's go into promptbuilder and configure our template. We can open this up, uh, and we can click one of the templates we want to use, which is a meeting email request. What you see that's really interesting here, actually, is that the prompt itself was written in English. And that's interesting because our end user's language is Mexican Spanish. So how does this actually work? Let's give it a shot. If we take any contact, let's say Sarah as a lead and a product, let's preview this. And what you're going to see on, um, the screen just shortly is that the resolve prompt is still in English, because it could be that the author of the prompt, the admin, wrote it in English. It's an out of the box template that was configured in English. But what you're going to see on the screen shortly is that the response of that prompt is in Mexican Spanish, as you see here. So the resolution in English, the language of the prompt, and the response in Mexican Spanish for an end user to use in the end to end workflow customized for their use case. And what's really interesting here as well is that in the next few months we will give prompt admins the ability to force a prompt, template or a prompt to be in a specific language, but also have the option for their end user to automatically see it in whatever

their is in. So if you want your prompts to always output in Swedish, you can have it that way if you want. So this is all part of our mission to make promptbuilder as accessible to people around the world and make the most impact that way. So, that being said, we've wrapped up our investment areas, introductions, and some of our demos. We just want to give everyone a big thank you for coming out and spending time with us here. And we also have a quick reminder that if you scan this QR code 1st 4000 attendees get a free five dollar Starbucks gift card. So please feel free. And also we have some more enablement as well.

Speaker A   00:17:47
Um, just going to let this.

Speaker B   00:17:48
Yeah, just give. All right, good there. Please feel free to take a picture of this as well if you want to go deeper, dives into how promptbuilder works, or even take it a little bit more high level best practices for getting started, or just introductory sessions. We have all of that here. So we look forward to seeing you at some more of these events, uh, in the next couple of days. So thanks once again for coming out, and we'll feel free to take questions after. Thank you so much, everyone.

Speaker A   00:18:13
Thank you.