SDS PODCAST EPISODE 958: WITHOUT TRUSTED CONTEXT, AGENTS ARE STUPID (FEATURING SALESFORCE'S RAHUL AURADKAR)

Jon Krohn:	00:00	If AI models are so intelligent, why do they keep doing such stupid things? For my guest today, the answer is simple. They lack the context they need, and he's got the solution to provide it. Welcome to the SuperDataScience Podcast. I'm your host, Jon Krohn. Today's guest is Rahul Auradkar, an EVP and GM at Salesforce, responsible for the products that make up Salesforce's $7 billion unified data engine. We've partnered with Salesforce today to announce the addition of the Data Management Solution Informatica to that data engine. Hear about that and other key topics around providing the right data context to make AI particularly agentic AI effective.
	00:45	Rahul, welcome to the SuperDataScience Podcast. Where are you calling in from today?
Rahul Auradkar:	00:49	Thanks, Jon, calling in from Seattle, Washington, sunny Seattle, sunny Washington.
Jon Krohn:	00:55	I sense some sarcasm. I like that to start the episode off. All right, so we're here to talk about a unified data engine that you're responsible for at Salesforce. And this unified data engine implies lots of pieces that are unified together. And so I'd like to talk about those key pieces. So one of those is called Data360. Rahul, what does that mean? What is Data360?
Rahul Auradkar:	01:25	Data360 is our offering of what was in the previously called as CDPs, customer data platforms. We truly believe we put the C in CDP because the prior CDPs put the M, if you may, marketing data platforms, it's irony that they built I think on CDPs and they ended up making it available to marketing only as a domain. So what does it mean? Putting the CNC in Data360, it's about unlocking trap data, harmonizing it, unifying it, finding the right insights, BI and AI insights, and more importantly, the

last mile of reverse ETLing it and activating it across all touchpoints, modalities, and channels in a very consistent manner so that we can provide delightful experiences to our customers.

Jon Krohn: 02:13 To your customers, and also the point that we're getting to here is that you're providing data that will delight agents as well downstream.

Rahul Auradkar: 02:21 That is correct. It also forms the foundation. When we first launched Data360, it was called a different name. We had the foresight and the vision to talk about Customer 360, and also we had the Foresight and vision to talk about it in the context of AI. With the advent of generative AI, now this becomes the grounding engine with all that is needed for agents to be more intelligent and more grounded as they deliver automation and outcomes.

Jon Krohn: 02:49 Awesome. Yeah. So we'll get into the Agentic stuff momentarily. That's the really exciting thing for me as well as probably a lot of our listeners, but let's talk a little bit more about this unified data engine. So I said that it had lots of components. One of them is Data 360. Another one is one that I bet most of my listeners will be aware of is Tableau. So Salesforce acquired Tableau in 2019. It's an intuitive system for visualizing and understanding data. Why was that something that Salesforce acquired and how have you integrated that into your unified data platform?

Rahul Auradkar: 03:23 So Tableau invented simplistic connect and explore type analytics. Throw dataset at it, you can start exploring right away. And the out- of-ox experience is phenomenal from it. And then subsequently it became an enterprise mainstay. In terms of what customers are doing with their data sets, making decisions based on analytics, it's essentially analytics-based decisions. Now, if you start looking at where the analytics industry has moved, we

have always talked about actionable analytics. And whether you go from descriptive to discovery, to predictive, to prescriptive analytics, that spectrum we've always looked at as being the holy grail in where analytics is going. What we have done now is combine that with the assets that we have with the Salesforce platform and with Data360 that allows you to have the metadata understanding, the unification, the harmonization. So you have ... Now we've also built one of the best semantic engines or semantic data models on top of that dataset that allows you to now have not just actionable analytics, but most importantly, it's the advent of agentic analytics, which is where we are with Tableau Next.

Jon Krohn: 04:42 Really cool. And so that semantic layer, I guess it allows an understanding of what's happening in the data and maybe kind of a way of interfacing with natural language. Is that what you mean by a semantic layer there?

Rahul Auradkar: 04:54 So there's two parts to it. One is when you take a look at data sets, say you're in an enterprise and you and I are talking different languages when we say revenue. Revenue is a semantic understanding of something that's meaningful for that enterprise and for that audience to which you're talking about revenue. So you want to have a semantic understanding, which is a collation of, which is a harmonized and unified collation of datasets that come together to define revenue. So you and I should be using the same language and also the same numbers that come with it. That's a semantic model. Now you can define things like ACV, revenue, churn, things like that. We can consistently define it using models. There's other aspect to it which you brought up, which is agentic analytics, which is agent first type approaches. So agent first, meaning natural language first.

05:42 You can ask, but think about this as ask data type scenarios, ask analytics type scenarios, as in you're

asking questions off of your analytics and your data. So we have two forms of experiences. You can do it through the traditional dashboarding analytics experiences, or you can ask your data some questions and then augment that with traditional dashboarding and analytics. And from where you can then seamlessly project it back into applications like Slack, seamlessly projected back into applications like Sales or Sales Cloud. So you can action on it right away.

Jon Krohn: 06:15 Right, right. That makes a lot of sense and is exciting. Turning back to this, I've kind of been going through the different parts of this unified data engine. So we have Data360, Tableau. The next one, I won't spend too much time on this one, but is MuleSoft. Tell us about that one.

Rahul Auradkar: 06:31 So MuleSoft has been an industry leader with iPaaS, the integration platform as a service, which is essentially app to app integration. You see the app sprawl in enterprises. Customers use MuleSoft to integrate apps from an app sprawl standpoint. Then we have leadership position in API management at API governance at the edge as to how you manage APIs and interfaces between different systems and different applications. And now we are extending that out to agents. We refer to that as agent fabric, the ability for us to govern and orchestrate and manage actions and provide security for agents across in a multi-agent, multi heterogeneous environment within enterprises. Now what that allows us to do, Jon, is it allows us to then have the ability to react to changes in data, the fluidity. One of the things that we focus heavily with Data360 is data fluidity. For you to have fluid data that we are unlocking, we are unlocking from silos.

07:30 We refer to that as unlocking trap data.That data sits in warehouses and applications and lakehouses for good reason. They have been running for good reason, but we need to bring that to life in a manner in which it can be

actioned upon and manner in which it can actually drive actions. So what MuleSoft allows us to do is this iPass integration and API and agentic governance. It allows us to drive actions at the end points where automation needs to get driven.

Jon Krohn: 07:55 Cool. All right. So that gives us almost the whole grand tour of the unified data engine. Now, the key piece is that you've integrated something called Informatica, Informatica's capabilities into Salesforce's platform, and that's really what we're partnering to announce today. So tell us about Informatica rule.

Rahul Auradkar: 08:13 Absolutely. So we are so delighted to have acquired Informatica. We closed it ahead of schedule. Our M&A team has done a fantastic job in getting all the regulatory approvals. We closed on November 15th, legally closed. It's part of Salesforce now. We believe it is one of the most synergistic acquisitions you can look at in the enterprise acquisition landscape. It completely augments and synergistic to what we are doing for, especially in this new world of agentic enterprises, the agentic transformation. What Informatica brings to us is they have been a leader in several different spaces across the entire data value chain, largely in data management, whether it is data quality or data catalog or data integration. And these are the areas where they've been leaders for a long time. That is synergistic to what we do with our data foundation that we are building. As an example, they bring the enterprise data catalog that layers on top of what we have is the customer data catalog through what we are servicing our customers right now.

09:15 So they give us a super set of the metadata that is available within the enterprise. And also they have the best in class data quality tools that are needed for us to cleanse data before we start using it in a harmonized and unified way. Informatic also gives us the ability to have all

the governance lineage, which comes from the catalog. And most importantly, they also have the industry's best ETL tooling through data integration. We can do all the iPass patterns through MuleSoft. We have a ton of connectivity that we have about data connectivity, but the ETL capability that they bring through data integration is bar none. All of that then adds up to our data foundation, which becomes the found AI foundations that allows us to then drive data flowidity, which in turn leads to context that we in turn are using as a foundation for driving agentic transformation.

Jon Krohn: 10:07 Yeah. I've had episodes on the show recently where guests made the case that the key to having AI applications work effectively is having the right context, that that makes or breaks an AI application, particularly an agentic AI application working effectively.

Rahul Auradkar: 10:24 Yes.

Jon Krohn: 10:25 Yeah. I like that answer. Yes.

Rahul Auradkar: 10:29 I can go deeper in there. Look, I think most enterprise … People use phrases like data is the new oil, et cetera. So we know enterprises are data rich. We all know that. What they are, they're data rich, but they're context poor. And the AI models that exist right now are incredibly intelligent, but they're corporate stupid. They don't really understand the context in which that intelligence needs to be applied. An example of that would be you could have three silos of data and you could get three different, differing ways in which you can feed that information into LLMs through a prompt or however you're doing what you're doing. And you would get distinctly different disconnected experiences for that business application. An example of that would be, I use a favorite example of mine. You are researching and buying a car. You buy a car and subsequently agents or automation is driving

advertisements and reach out to you based on the same car that you bought.

11:39 What is the chance that you would buy two cars within a week of having bought one of the same kind? What they should be telling you more about is, "Hey, here's an insurance package, here's extra, here is a warranty, here's accessories that you want to buy." Anything related to the car. The reason is because a sales system is different from the marketing system, which is different from the customer service system. They're all different. That context has not been brought together. My personalization, what did I scroll? My profile. They don't know that I've already finished a purchase because it's sitting in a transactional system. They don't know the fact that they've already advertised to me. People have had reach outs to me. They're servicing me. The service, the car is servicing me. They don't have context of having brought all of these different schemas together, having brought these insights together.

12:24 That makes agents incredibly stupid in their responses. If you ground them with this context, then you would have delightful experiences, probably fewer experiences, fewer touch points, but they would be delightful.

Jon Krohn: 12:36 Nice. And that is what all of us want to be getting to is humans having delightful experiences with their AI interactions for sure. A large portion of data science projects fail because of inadequate data, maybe something like 80% of them. So for our listeners, for data science teams that are currently struggling with AI projects failing because of inadequate data, what are the most common data pipeline bottlenecks and how do things like the unified data engine that you've put together at Salesforce that includes the Data360, Tableau, MuleSoft, Informatica, how does that kind of

unified data pipeline together eliminate these data issues and allow projects, AI projects to succeed?

Rahul Auradkar: 13:22 So first off, we just talked about the idea that you would unlock trap data, which means you would bring the power of all the data to life and bringing them together in a meaningful manner in which all the data sets are relating to each other. That's all about context. Then you have on layer on top of it, Jon, governance and trust. You don't want just context, you want trusted context.That would mean lineage, it can have consent, it can have PII handling, it could have masking for unstructured data. Then you have in the moment and real time signals, whether they're events or transactions or behaviors or clicks and things like that. Then you have historical context, which is customers and assets and cases and interaction states and things like that. And then you layer on top of it, business specific rules, business rules, and then policy enforcement based on some level of policy that you want to enforce based on business policy as well as regulatory policy.

14:27 And then you have domain specific semantics, whether it's an industry level semantic or whether it's a workflow associated with it. Now, all of these things need to come together. And what we are providing with our data foundations for drive agents is a holistic offering that allows you to do all of the above, which is context, trusted context, trusted enterprise context with governed data, all of that being fluid, real-time signals, historical context, historical memory, business rules and domain-specific knowledge that is workflows that we have already codified through our applications.

Jon Krohn: 15:05 The solution that you're describing, it appears to integrate a lot of previously separate kinds of skill sets. So folks who understand how semantic layers work, security, governance, so many different skills now unified together,

how can our listeners or how can folks who are shaping data science teams upskill their teams effectively to be able to make use of all of these capabilities together?

Rahul Auradkar: 15:32 So first off, we have moved into a world in which a lot of what we do is no code and low code. We do support procode scenarios as well in that if you have, as an example, if you're a marketeer and you're trying to drive some campaigns using segments and activation, you can do a lot of pretty much everything that I said through configuration. So the upskilling is more about focused on the domain in which they are great at, our customers are great at and leaving a lot of what we are doing to the simple to use low-code, no-code tools that we ship. So that's one thing. The second thing is that the watermark of repeatability in terms of what can be done consistently across multiple domains is something that keeps going up. For example, our metadata layer that we have right now to drive all consistent customer interactions, now we can extend it out to the enterprise with the advent of Informatica.

16:34 The data catalog that we have that allows us to know more about customers and have lineage and consent, et cetera, now we can extend it out to the enterprise. So what we are doing is we are with our platforms and with our products, we are allowing our customers to focus more on the domain in which they are really excelling at by using the tooling that we have.

Jon Krohn: 16:57 It sounds like maybe then my question is kind of inverted. It's because so much functionality comes working out of the box with a unified solution like this that the answer is you don't really need to be upskilling that much at all.

Rahul Auradkar: 17:08 Well, there is upskilling required to the domain in which they operate and even there needs to be upskilling to use the tooling that we have, even if it is no-code and

low-code. So for which we have perhaps the industry's strongest learning community, which we refer to as our Trailhead, Trailhead Academies and Trailhead, Salesforce has a vast array of lessons and certifications online, and we keep refreshing that every so often. Third, we have what we refer to as data blazers, trailblazers, agent blazers. These are all hundreds of thousands, in some cases, millions of people out there who are educating themselves and others along the way. Not to mention our partners who do the implementations as well. So all of that, when you bring all of that together, that gives a strong foundation on which our customers can upskill, not just in the technology, but more importantly, in the domain in which they're operating.

Jon Krohn: 18:08 Excellent. Great answer. I appreciate the enumerated nature of it. One last question for you, looking out broadly, what does enterprise success look like in 2026 for enterprises that want to be having success with Agentic AI, what's the limits? What are the limits that they should be reaching for in 2026?

Rahul Auradkar: 18:35 Well, in 2026, one good canonical example would be Salesforce. So we have the concept known as Customer Zero. Salesforce running on Salesforce. And Mark tells us that we have not earned the right to ship software to our customers unless we make ourselves successful with the software that we are built. So if you take a look at all the videos out there, the YouTube material that's available of agentic enterprises and agentic transformation, some of them is a lot of that can be visible through how we have done. Salesforce has done the implementations. For example, our help.salesforce.com is agent first and agent only right now. And a vast portion of our questions with our customers now are answered through agents, and we have humans that then are augmenting the agents. And if you take a look at our customer targeting and how we do our sales pipeline and et cetera, all of that runs on unified

data from Data360 with agents running what we refer to as background marketing agents that are reaching out to customers.

19:43 We're putting predictive AI that allows us to create lead scores that happens within a 30-second window of you having really gone to our website. And I can go on a few more agents that we have. We have the SDR agent now that is shipping. Our customers are using it. We have some customers talking about it as well. So I think enterprises are starting to go in that direction. One more great example would be William Sonoma Souchef. If you take a look at William Sonoma's agent that does Soushef, that's all running on Agentforce. That's an example of where enterprises are going in the agentic transformation. It's no longer starting at the margins and the edges. It's right there at the core of the enterprise from a transformation standpoint.

Jon Krohn: 20:25 Really cool examples there, Rahul. And if people want to learn more about Agentforce 360, they can refer back to a recent episode with Tyler Carlson from Salesforce who spoke in detail about Agentforce 360, which it seems to obvious to me, Rahul, that that builds really nicely on top of this unified data context that you have been responsible for leading. And so that unified data context then allows Agentforce 360 agents to be able to be as powerful as they could be.

Rahul Auradkar: 20:59 Absolutely. That is the plan where now we have, if you take a look at our latest architecture that we laid out at Dreamforce, we have Data360 and all the Unified data, which is a context at the center of everything that we do. And then the Agentforce 360 platform is the foundation on which we are delivering all the applications that our customers use to reach their customers, whether it is sales or marketing or revenue or analytics or Slack for

collaboration or Tableau for analytics. All of that now is running on this Agent4360 platform.

Jon Krohn: 21:30 Really cool. Thanks, Rahul, for taking the time. I know your time is tremendously valuable and our listeners wouldn't know this, but we've actually gone over the time that we had allotted for you. So I'll quickly just get through the final questions that I ask every guest. So first off, do you have a book recommendation for us, Rahul?

Rahul Auradkar: 21:46 Book recommendations. So I normally like biographies. And one of the biographies that really stand out for me is Bob Woodward's green sbiography. I believe it's called Maestro. That's something that really sticks out. I read it many years back. It's pretty well written. It's very down to earth. It's anecdotes driven that stands out.

Jon Krohn: 22:11 Excellent. Thanks for that biographical recommendation, Rahul. For folks who want your recommendations on anything, biographies, on unified data, on agents, where should they be doing that?

Rahul Auradkar: 22:22 So first off, Unified Data, they can go to data.com, which also leads you to salesforce.com/data360. We can also get mulesoft.com and then Informatica and Tableau. All those four are available and they're very rich. There are customer case studies and then from there, and you can do some search on YouTube as well. You can find a lot of videos there on YouTube too. And then you have my LinkedIn profile and the Datablazer site on LinkedIn as well. Perfect.

Jon Krohn: 22:51 Thanks, Rahul. Hopefully we can check in with you again in the future to get more of your insights.

Rahul Auradkar: 22:57 Absolutely. Thanks for the opportunity, Jon.

Jon Krohn:	23:02	Nice, concise, and informative episode with Rahul Auradkar, in it Rahul covered how effective AI grounding requires not just context, but trusted context, meaning governance, lineage, consent management, and real-time signals all working together. He also described the components of Salesforce's unified data engine and how the acquisition of Informatica brings enterprise data catalog, best in class ETL and data quality tools that compliment the unified data engine's ability to provide trusted context to AI models, including to agents. I hope you enjoyed the conversation to be sure not to miss any of our exciting upcoming episodes. Subscribe to this podcast if you haven't already. But most importantly, I hope you'll just keep on listening. Until next time, keep on rocketed out there, and I'm looking forward to enjoying another round of the SuperDataScience Podcast with you very soon.