Interview

Customer first thinking.

The Future of Customer Data Management

An Interview with Chris O'Hara, Global Product Marketing Lead, SAP Data and Analytics

Chris O'Hara:

Chris O'Hara leads the customer data solutions practice at SAP and is the co-author of Data Driven: Harnessing Data and AI to Reinvent Customer Engagement.

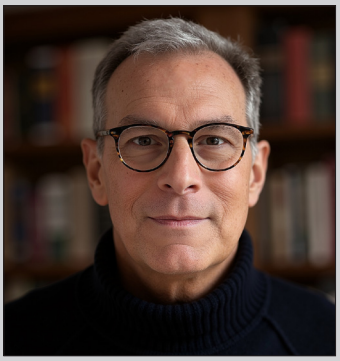

Back when mainframe computers ruled the world customer data at the individual level was kept under lock and key by the high priests of IT.

Marketers were usually given only limited access – sometimes no access at all. Computing time was just too precious to waste on non-operational uses of the data. Busy IT staff gave low priority to ad hoc queries or list pull requests coming from marketing. It was the Age of Big Iron - siloed data stores, large and very pricey Relational Database Management Systems (think Oracle) – monolithic legacy applications and systems designed to keep the lights on (think SAP).

IT resistance to marketer's pleas for data began to lessen by the early 1980s as the cost of computing began to drop. Database marketing suddenly became fashionable. And then CRM systems stormed the market in the 1990s (think Siebel), generating new sources of customer sales and service data. By the early 2000s SalesForce had proven that CRM could be served up as a cloud-based software service instead of a costly on-premise system. CRM became more affordable for a broader range of businesses. Now all kinds of personally identifiable customer data was available to marketers.

But it was the explosion of Big Data in the mid-2000s, spawned by the Internet, that began to throttle internal data management systems. Traditional data warehouses simply couldn't cope with the incoming deluge of web-based data in its many diverse forms. So a new form of database came along to deal with the problem: data lakes (later evolving into hybrid "data lakehouses") designed to serve as vast catch-all basin for raw data that could be directly queried by end users.

Yet even with all of these technology innovations, the goal of a "Single View of Customer" – "One Version of the Truth" – the so-called "360 degree" view – remained elusive, more an aspiration than a reality. Marketing had its own view of the data – Customer Service and Sales had their view – Finance and Operations had their view. Separate islands of customer information. Piecing together a common view – a unified customer profile – meant that data engineers had to work across different application siloes, extracting, cleaning, transforming, standardizing and loading the data into a single master database. Often those projects hit a wall due to the enormity of the task.

Just over a decade ago another technology came along to make the collection, integration and activation of customer data much easier. Called Customer Data Platforms (think Twilio Segment), they were initially pitched to businesses primarily as an enabler of omnichannel engagement where identity resolution and management is crucial. Soon CDPs won a preferred place in the data management firmament

of many companies, highly valued for its integration capabilities. At last, a single view of the customer!

Today there is growing recognition that a unified view of the customer, structured around first party data, is the key to competitiveness in an AI-driven Experience Economy. Even the C-Suite is starting to appreciate that a reliable, unified and dynamic profile of the customer, enriched with consent and preference data, is a strategic asset deserving of investment. A properly architected customer data foundation, many companies have finally realized, can in fact drive business growth, in part because it leads to a better customer experience.

So once again customer data management technology is evolving to support the need for real-time engagement and personalization as well as AI-orchestrated customer journeys. Rather than function primarily as a point solution, CDPs are being retooled to operate at the enterprise level – so-called "composable CDPs" - that sit astride the modern data warehouse ingesting data directly (think Snowflake). No more copying data from one database to another. The advantage to marketers: a customer data backbone that the entire company can rely upon. One that no longer has a sign on the door saying, "Marketers Keep Out".

At one time Chris O'Hara was one of those data-driven marketers knocking on the door of IT. Today he leads the charge at SAP to transform the marketing data infrastructure. He believes that in order to deliver a more personalized customer experience at scale, organizations will need to link supply management and demand creation. They can only do that, he argues, by creating a "data fabric architecture" that provides a seamless view of customer data across all sources and formats. Chris has spent most of his career as a pioneer on the front lines of customer data management, and is a recognized expert on data-driven marketing.

SS **Stephen Shaw (SS):** You're seen as a pioneer in data management and ad tech. Tell me how you got started.

CO **Chris O'Hara (CO):** In the very early days of the Internet – this was during the dotcom boom - I was running a company that had built a programmatic marketplace for advertising which enabled publishers to post their remnant inventory and let agencies buy it. So we built all this technology around it where you could append your audiences with all kinds of third party data. And then some really smart guys led by Brian O'Kelley - who knew the Internet a lot better than us - developed AppNexus1. We looked at that and thought, oh my god, they built the thing we were trying to build but just didn't know how.

During that time, I started writing for a publication called Ad Exchanger, which was starting to become popular as programmatic media grew up. I had a data-driven marketing column that got a big readership, and it caught the eye of a guy named Tom Chavez who had founded a company called Krux2 which was a data management platform that had become really popular with big publishers like the Wall Street Journal and the New York Times. He thought there might be a marketing use case for his data management platform and asked, "Do you want to run this part of the business?". And I was like, sure!

Kellogg's was our first client, and they had always suspected that they were getting ripped off with their ads. So we put all these tags on their site, tracked all their advertising. We could see Kellogg's was serving 400 to 500 ads every month to the same people. We had calculated that any exposure after about 12 ads was wasted – after that people just didn't notice the ad anymore. So we went to Kellogg's and said, you know what, if we just suppress every single person who's seen 12 ads already, you will probably save $20 million a year. And they were like, holy crap, that sounds good. And it worked. So we asked their head of marketing, Jon Suarez-Davis3 JSD, to go around the country with our sales team and we convinced every big CPG company to buy Krux. And then it was off to the races. Salesforce later bought our company for $800 million.

At Salesforce, I ran data and analytics for them. We acquired five or six companies at the time, really cool stuff. And then, I got a phone call from SAP. They had built a CDP and wanted me to come run marketing for that.

SS SAP is not exactly a household name in marketing. What convinced you to go there?

CO Just to back up, at SalesForce we did email marketing like nobody's business, having bought Exact Target in 2013, which was one of the bigger acquisitions in MarTech at the time. So they knew everything about the email consumer. By acquiring our technology (Krux), they were able to have a 360 degree view. Were they a Minnesota mom? Were they a travel intender? Were they suburban dad? Whatever.

So I said, why the hell do we have all this customer data trapped in the SalesForce Marketing Cloud? Why wouldn't it also be in CRM? And so that turned into a long conversation about data management in general. And eventually, just as I was leaving, we were thinking about a horizontal data layer sitting underneath everything - CRM, Tableau, Mulesoft, whatever. The intelligence living there rather than in the application layer.

So when I came to SAP, I looked around and I said, geez, you know, this is really similar. But in our case, ERP is the center of the universe, not CRM. We know the account, we know supply chain data, we know pricing data. People use us for finance, for hiring, for HR intelligence, for...

SS … running their business, really.

CO … yeah, for running their whole enterprise. So why do we have this really cool data management technology just in CX to power email and commerce? That's really dumb. So our CMO at the time, Julie White, put me under Dan Yu in data and analytics. So now we are building a data layer to support this idea of having one database to rule them all - what we now call and know of as a "data fabric" - a beautiful data fabric that stitches everything together, that semantically unifies data. We can build a knowledge graph that connects all these different data attributes together and knows how they relate to each other. And then ultimately, when AI really comes of age and there's this thing called agentic interactions, it will give these agents someone to talk to so they can be really effective.

Now we're not there yet, but we've migrated early versions of what we're doing, which we call Datasphere4, into a fully blown data cloud offering which we've recently announced. So we're really trying to tie everything together and give our customers a really super strong data intelligence layer that relates all these things together. And then what we want to do for the first time is let our customers use that data and intelligence layer in everything they do.

SS Before we get too much deeper into the data discussion, I do want to close out the conversation about AdTech because there's two things that have happened recently, and I'd like your perspective on the implications. One is obviously Google's retreat from phasing out cookies. Two is Mark Zuckerberg's comments recently announcing what is effectively a jihad on the ad industry, declaring that Meta is going to deliver an AI-based turnkey ad solution that basically eliminates all the intermediaries, which sounds like a massive land grab. Between those two events, what's your perspective on the future of the digital ad business? And does that put even more emphasis on the importance of first party data?

CO First of all, I'll say this, I never really was a big believer in the digital ad ecosystem. AdTech was largely about banner ads. And they're really not good ads - they're terrible, the creative's terrible, people don't really like the format of the ads. As Kellogg's discovered, people very quickly get banner blindness. So I think the future of that part of AdTech is dead.

Today I can digitally target a household or individual through my TV, whether that's through an ad serving service or, say, Samsung itself who, let's be honest, is watching and listening to the household all the time with their little $700 TV. So now I can relate my media attribution to an online purchase that I can track. So that's where I think it's going. But I think we have to go beyond the scammy world of old school AdTech and get into a world where creative comes first. How do we programatize that, get it to scale and then give someone the ability to measure it? That's where we have to go.

SS Have companies finally reached the point where they see customer data as a strategic asset? Does every company now need a customer data strategy?

CO Yeah, that's a really profound and interesting question. The notion that we can ignore customer data is absurd because the whole world can't operate without the customer. The customer is the atom, right? It's the atomic level of information in our enterprise. So I would argue it's maybe the most important asset because without your customers, you cannot survive as a business. Today we capture the data in all kinds of different applications, but we don't connect it. It never comes together holistically so that we can steer our business.

I'll give you a great example. We're about to announce that for the first time ever, we're going to put SAP supply chain data together with demand data for a major company. They might have 50 different campaigns

running for different shirts or jeans or whatever. Customers are reacting and clicking and buying stuff online in different sizes, colours. So they're getting all this demand data and yet, this company has no idea what's in the warehouse, what's being produced, what's in the supply chain, can they even get it to where it needs to go? Yet they're creating all this demand.

On the flip side, at SAP we know all the raw materials, costs and prices and what's available in the warehouse, when and where it can get shipped - but we don't have the demand data. Should we buy more cotton? Should we produce more of the red jeans or the blue jeans? How do we plan for this? Those two worlds – supply and demand - have never come together at scale. By taking those two very distinct and siloed data assets and plugging them together, all of a sudden you're giving the enterprise an insight it can use to steer the business.

SS I think you call it ERP to CRM.

CO Correct.

SS What stands in the way of an organization achieving that holistic view?

CO Yeah, well, a couple of things. One is governance and security. So, you wouldn't be surprised to know at SAP we have a lot of big finance companies who prefer to keep their data on premise and not in the cloud. There are many companies that have been very slow to adopt the newer technology, right? Their fear is justified because they're reading all the time about data being hacked and customer data being exposed and financial liability for that happening, and those are real fears.

The second thing I would point to is we have a tendency to sell technical solutions in this business and not focus on outcomes and ROI. What made us successful at Krux and at Salesforce and at SAP is that we can relate the solution to an outcome and tell a customer story, and make it real. I just think software sellers, especially in MarTech and AdTech, don't do a terrific job of that.

I also think what happens is that a software company sells a very intriguing, powerful solution, but then points to Accenture and says, okay, here's your integration partner. It will only take $10 million and five years to deploy this, so don't worry. So we have to be very cognizant that people want to own, maintain and deploy these things more easily than they have in the past. We can't over complicate the solution. And we can't overpromise and under deliver, which I think is another very common thing.

SS Who should be leading the charge in customer data strategy?

CO Yeah, that's a really good question. At Salesforce, I would sell Marketing Cloud to the CMO. But the CRO has their own 360-view that is specific to CRM. There are multiple versions of a 360 and they're all very important. And now I see that even more intensely at SAP because our primary persona is the CIO. They own the technology, they're responsible for it, they have the budget, even if they don't initiate the buy.

At Salesforce we had something called "Journey Builder" where you could design and change a journey, and it was brilliant. But all the intelligence lived in the endpoint application, not in the data platform. That was silly. Like, why are we building all the smarts into this application? That's stupid. Now we're starting to say, okay, there's this data infrastructure that connects all these things together. That's where I should build the journey. That's where I want to start. We're going to bring the intelligence to the application rather than vice versa. But here's part two. Now we're in a world where AI agents are talking to other AI agents. But to whom are they talking? They're talking to these data clouds. And the first question they ask is, "Do you have this information I need"? Then the other agent says, "let me check if I can work with you". And then if there's a handshake, a lot of stuff gets done very, very quickly. So we have to build for that future of mind. And the differentiator is the underlying data layer obviously. That's where the battle is. That's the next 10 years in software. And that's going to manifest in many, many super interesting ways. That's the future we're looking at.

SS You've said that building a "modern data management infrastructure begins and ends with mastering customer identity." You also wonder why identity and authentication management still doesn't have a seat at the table, as you put it, given its importance. What approach should companies take?

Interview

CO There are a lot of different pieces of individual identity that make up a profile. You have 50 different ways of identifying Chris O'Hara. There's several dozen cookies. There's a mobile ID. Apple has a very special Safari ID. There's your Chrome ID. Your laptop has an ID. Your tablet. All of these things have to be put together into one persistent profile of Chris O'Hara. If you're marketing to me on my mobile device and my laptop and you have no idea that I'm the same person, I'm getting overexposed to your messaging because you think I'm two different people.

So that's number one: the idea that you have to have this persistent ID is super important. Without some of that infrastructure you cannot be an effective marketer, and you cannot do effective analytics and you will never ever get to attribution because you think you have 25 Chris O'Hara's and 25 other platforms say you have now 50.

So that problem has to be solved. But there's also matching the digital identity together with the PII identity where you know my email and my phone number and my physical address while also giving me some ability to manage that with permissions. So that's governance and without that you lose the trust of the customer. Now every company thinks of this as a last order assignment, but it really is a first order of business because of safety, security, governance, intelligence, right? And two, consumer trust, without which you can't do anything.

SS Why then wouldn't identity management be more of a priority?

CO Well, I think it's a huge technical challenge to do really well. And secondly, I think it's very expensive. And thirdly, I think we got used to a very lackadaisical regulatory environment that doesn't really care that much. And we've seen all kinds of big companies break these rules and just get a little slap on the wrist for doing so. And we look at that and say okay, maybe I don't have to care that much, right? And that's really horrible. And one of the consequences of that is we've built these big walled gardens as places where we feel we can actually get to a real consumer more easily than the open Internet. And we really don't trust the open Internet that much. And we shouldn't.

SS Until now CDPs have been the go-to solution for single view of the customer. Now the integration function is moving to the warehouse level, so you have "composable CDPs". Help me understand the differences and where it's going to end up.

CO I think everyone thought for a little while that when you look at these big enterprise marchitectures, there'd be a box with the CDP in the middle and that was your data store, your process, your centralized point. And that became very untrue. As you said, I think people went right back to the data warehouse, right back to the old school. I was very skeptical of that too until I moved to SAP and I saw the power of some of these data warehousing applications like Snowflake or Databricks. So CDP lost half its appeal because I can put all my data in Databricks. CDP becomes interesting for orchestration and activation, and for very specific things. But CDPs are too narrowly focused on marketing orchestration. It's a very expensive proposition for making email better.

SS I want to understand the "data fabric" part you referred to earlier because I confess I don't fully understand it. Can you explain what it means?

CO The primary problem we're trying to address with the data fabric is really around how data comes together and how it expresses itself in other applications. Take the word "customer". Customer means one thing in CRM, but I have another application where customer is expressed differently, and the fields are sort of different. It may be capital "C U S T" and the other one may be sentence case "Customer" and the other one might be lowercase "customer". So the way data technology is looking at this stuff, there are different fields and there are different ways I express this notion of customer, of price, of SKU, of product, of organization, of partner. And until we semantically organize that data and give them one canonical model to come together, that data will never be able to talk to each other and build what we consider this data fabric.

The second part of the data fabric isn't just the semantic organization of the data fields and how we call them. It's the metadata attached to these fields. What does this mean? Why is this data important? How is it used? Who

uses it? For what purpose? And so there's a lot of actual metadata around that actual data that is super important.

And then the third thing is when we think about a data fabric is how do you relate those individual attributes we have captured together and how do the edges fit together, how do they relate? And that's what we think of as a knowledge graph. So a technology that organizes all these different data attributes, puts them together and figures out are they related, are they important, what is their meaning? And so it's a variety of different things that come together to have this layer of data that's actually got some intelligence within it.

So that goes beyond just the profile and managing it and knowing that Chris O'Hara is a travel intender, and a father of three, and lives in this zip code, and has this income. It's how I interact with all the attributes that surround me. How they relate to other profiles, how they relate to different applications in the application space, how they relate to non-human entities like things in the IOT world. So it's a bit more complex. That's really the approach.

SS Given all that's happened with AI in the last two years, what's the data management future look like three years from now?

CO We start to be more on the cusp of that agentic future. Agents will deploy really interesting multi-channel experiences to customers automatically and your data cloud will be hooked up to many, many different endpoints. The way you connect with these data clouds will be very important because they'll ultimately have the data that decides where marketing dollars are invested. And we may live in a world where we can see if those marketing dollars were spent wisely because there's finally some closed loop around how consumers reacted to our messages. But I think it'll take another seven years before that really starts to coalesce.

[1] AppNexus was a cloud-based software platform for programmatic online advertising, later sold to Microsoft and rebranded Xandr.

[2] KRUX was a data management platform (DMP) that was later acquired by Salesforce.

[3] Jon Suarez-Davis (jsd) is Chief Commercial Officer for super{set} which is a startup studio based in San Francisco that founds and builds data-driven software companies.

[4] SAP Datasphere is a data fabric solution that helps businesses connect and manage all their data assets.

Stephen Shaw is the chief strategy officer of Kenna, a marketing solutions provider specializing in customer experience management. He is also the host of a regular podcast called Customer First Thinking. Stephen can be reached via e-mail at sshaw@kenna.ca.