

Customer
first
thinking.

Mastering Agentic AI:

An Interview with Jean-Pierre Lacroix, Founder and President, Shikatani Lacroix Design

Martin Kihn:

Martin Kihn is the SVP Strategy for Marketing Cloud at Salesforce and author of the book "Agentforce".



Never before have so many people become as infatuated as quickly and as enthusiastically with an emerging technology as they have with AI.

According to one estimate, 66% of people now use AI regularly¹. It has become an indispensable utility in their everyday lives. People use it to shop. To plan their travel holidays. To consult on health concerns. To manage their finances. To help with repairs. To interpret technical jargon. To manage and respond to e-mails. To plan dinner. To organize their day. Even to serve as a confident and friend.

All of this adds up to an historic inflection point in the market. AI is influencing how people discover, compare, and buy products and services. It has already upended search marketing. It has disrupted the way brands are discovered. It has turned the marketing funnel upside down. It has drained away web site traffic. It has neutered the power of brands to influence consumer choice. It has made prices fully transparent. Most of all, it has raised the expectations of customers to new heights.

Suddenly the old marketing playbooks have become relics of a bygone era when brands were more or less in control of their destiny. Marketers are now faced with reinventing the way brands go to market. But it is not yet clear that very many brands – especially me-too brands most at risk – grasp the full extent of the threat. Instead marketers are caught up in merely using AI to compress campaign timelines and workflows. Understandably so: they spend more than half their time today on tedious and repetitive production work. CMOs who can automate routine execution tasks and trim headcount are always seen as heroes by their bosses.

Marketers can now push out more campaigns in less time with fewer resources across more channels with more precise aim. But they are falling short in harnessing the full power of AI to reimagine the customer experience. Which is why so many AI pilot projects fail to scale. AI is just plugged into current ways of doing things. Marketers just want to keep doing what they've always done – only faster – when what they really need to do is rethink processes from scratch. Instead of seeing AI as a way to free their time for more strategic visioning, they think of it as scab labour. So there is an enthusiasm gap that is holding back progress.

The other major stumbling block is a longstanding one that marketers can't solve on their own. They are often stuck working with legacy systems and a mishmash of loosely connected marketing tools that are simply underpowered to cope with the demands of generative AI. Hyperpersonalization is an alluring concept – until it runs headlong into the sobering reality of stove piped systems and dirty data. What marketers really need – what they've always needed – is a unified customer data infrastructure and knowledge base – machine readable by AI – that they can trust.

Which is exactly what Salesforce has brought to market: an all-in-one enterprise solution called Agentforce that connects all of the requisite components in an integrated, AI-ready architecture. By implementing Agentforce marketers can jump start the deployment of agentic AI, sidestepping the perilous path of going it alone. There is always a lot of heavy lifting involved in taking a do-it-yourself custom-built approach, and often it does not end well, mainly because of messy, conflicting data and disconnected systems. Too many moving parts, too many chances for things to go sideways. Whereas with an interoperable, purpose built, tightly connected solution, especially one from a known vendor like Salesforce, the on-ramp is less steep.

Since its inception almost 30 years ago, when it pioneered the concept of cloud-based SFA software as a service, Salesforce has continued to push the boundaries of what's possible. Its longtime market dominance in CRM gives it instant credibility and peace of mind with enterprise-size businesses looking to build a central operating system for agentic AI automation.

In his role as strategy lead for Salesforce's Marketing Cloud, Martin Kihn is charged with explaining the advantages of using their platform. His gift is that he is able to translate technical jargon into plain language that marketers can easily understand. His two prior books pulled back the curtain on customer data infrastructure, building an aspirational picture for marketers of what their technology foundation should look like, from front to back end. In his newest book Agentforce, Martin Kihn lays out an evolutionary path for marketers to follow if they hope to harness the full power of agentic AI.

SS

Stephen Shaw (SS): Do you see yourself as the “Chief Explainer” for Salesforce?

MK

Martin Kihn (MK): I think that's exactly what I am. If I can get to a point where I understand something technical - which takes work - then I can explain it. Richard Feynman² said that if you can explain something complex to an eighth grader, then you really know it. I feel almost anything can be understood, even AI, if it's explained correctly. So that is my role.

SS

Your book starts off quoting Mark Benioff calling this “a moment in time like we've never seen”. A lot of people have likened this period to the emergence of the Internet in the 90s. You were around in those formative years. How would you compare the two eras?

MK

The Internet was absolutely revolutionary, but it took time. And the real impact didn't happen in the 90s. It happened when we got the iPhone, and then social networks. In the very beginning, companies came and went. Like Pets.com. They did a lot of really funny ads, you may remember. And that's exactly the same business as Chewy.com today. Chewy is very successful - Pets.com was not. It bombed. And it's just because the tech wasn't there yet and people weren't comfortable giving their credit card. The Internet was too slow. It took probably 10, 15 years to catch up to the hype. With AI, it's happening more quickly. I think there is hype. There's an over expectation. But AI is evolving quickly. These frontier models can actually improve their own code. So the release cycles have shortened. But the real revolution hasn't happened yet.

SS

In the book you concede there will be disruption. Marketing is certainly acknowledged as being right in the crosshairs. Does marketing need to reinvent itself?

MK

I think there's going to be a lot fewer people focused on individual channels and tools. Like the Salesforce “Admin” who knows email really well - there won't be any pointing and clicking anymore - a lot of that will be “agentified”. But then that person will be managing a team of agents. Marketing is more complicated than customer service and B2B sales and it's just because it's campaign related. Customer service is a lot more predictable. Eighty per cent of calls to a retailer are either “Where's my order?” or “Can I return this order?”. You can identify and automate a lot of those processes more easily. And our Salesforce sales team is all in on AI - marketing not so much. And the reason is because marketing has a lot of pieces that have to fit together. You have to create the campaign, you have to know the audience, you have to come up with the idea, you have to do the creative. Then you need a media plan. So it's not just one agent, it's many, many, agents. We're getting better at it, but I think it's happening more slowly than people expected.

SS

Marketing's traditional role has been hitched to the attention economy. As we transition now to an Agentic Commerce economy, where the consumer is using AI as a mediator, all of a sudden marketing's role becomes a little murky, am I right?

MK

I think marketing itself is going to become more accountable for revenue. In the past, in my agency days, we looked a lot at brand awareness just because that's the only thing we could really track. That's not the same as "Did we sell anything?". But it was the best we could do. Now there's Agentic Commerce. There's a way to attach your payment to the agent experience. You can close the attribution loop really easily: they were in GPT and they did this and then they made a purchase. Our clients are much more focused on ROI now and real attribution, not theoretical attribution. That trend is clear.

SS

Everything I've read suggests that two thirds of companies still are AI laggards and way behind where they need to be, certainly far behind the early adopters. What's the biggest mind shift companies have to make to be successful, and where are most companies today going wrong with their AI deployments, just from what you've observed?

MK

AI is a tool. It's supposed to make our lives easier. It's supposed to make us more efficient, more effective. But it's not something that's a cure all. It's not something where I can say, all right, I need AI to improve all my campaign workflows and just improve my marketing, and it'll be like magic. It doesn't work that way. It is a tool. It requires education. It requires implementation. It requires a lot more work than people think. So I would say just approach it the way you would any new technology. What can I do incrementally with AI? People are too focused now on the fear factor: what is it going to do to my life?

SS

My understanding is that companies are failing to get out of pilot mode and scale because they're trying to fit AI into an existing workflow rather than reinventing the work. Is that your sense of things?

MK

It is a different workflow. Also AI agents require access to a different kind of data. They like unstructured data, call center logs or the content that's on your site, like your videos. But to make that data accessible to any system requires a sophisticated foundation. A lot of customers are stumbling over that still. It was always a problem but with agents it becomes a really big problem if they don't have their customer profile data unified. It's not as real time as possible.

SS

What kinds of real transformative changes are you beginning to see amongst the more advanced AI users within the Salesforce base,?

MK

The majority of the use of Agentforce is around service. What's happened is they add agents who are handling all of the work that didn't get done in the past. Like lead capture on the web site. Most leads are scored and then they're ignored. So there's just a lot of work that the companies couldn't do because they just didn't have enough people.

SS

So that's customer service – but what about on the sales and marketing fronts?

MK

There was a study I saw recently where AI generated copy was not only more persuasive than human generated copy, it was generated more quickly. So people are accepting AI generated sales and marketing content in a way that we didn't predict. Also in imagery. Video is less mature - it's not production ready.

SS

But by transformative I mean have they made things better and not just more efficient and less costly?

MK

Yeah, buyers focus on cost savings first. I think most of our them do.

SS

Take me through a bit of the history of AI at Salesforce. You introduced Einstein - machine learning - way back in 2014. How has it evolved from there?

MK

So it's more than 10 years ago that Salesforce opened the office of AI. We had machine learning built into the products as soon as we could. But I think when people talk about AI today they're referring to frontier models. It was a little less than three years ago Agentforce was launched. Nobody was really talking about agents at the time. The whole concept of an agent was kind of confusing. What is an agent exactly? And we were trying to say, well it's a way to get work done for the enterprise. What we wanted to do is "enterpriseify AI".

There's a big difference between using it as a consumer - come up with a recipe, do a workout plan – if you make a mistake, doesn't really matter. If you're a company and you're creating content for customers, that's a whole different world. I mean you could be on the hook for a big lawsuit or whatever, it's a big deal. There's all kinds of guardrails and layers of protection that have

to be built. There is the trust layer - toxicity detection, bias detection, making sure no personal identifiable information is left, tokenization - all of that needed to be constructed. It didn't exist off the shelf.

And then the other thing was to make it easy to make agents. So create a templated framework where you can describe a topic and guardrails and instructions and then have test agents in a sandbox, see how they work, improve them, and then monitor them once they're out in the wild. The vision for Agentforce was always to be an agent factory - a way that your team could come up with ideas and test and deploy them pretty easily without being coders, without being AI experts.

SS

Benioff describes Agentforce as a form of digital labour. Is it a form of "enterprise ready" AI versus "DIYing AI"?

MK

Yeah, an agent is something that gets work done.

Usually it has a very narrow scope - the more narrow the better actually. Ideally it would be semi-autonomous because otherwise they're not efficient. They are in a way analogous to a human worker. Humans aren't fully autonomous either. I can go to Agentforce and it'll guide me through the process and build an agent.

"DIYing AI" is possible. It's the same way you could build your own CRM if you had enough engineers, you had enough time, you had enough patience, nothing's going to stop you. But it's just a lot easier to acquire something that's already been built and they'll maintain it for you and they'll assume the headache and you'll just pay them a fee every month.

SS

Can you provide an overview of the Agentforce architecture?

MK

Agentforce itself is built on "Hyperforce"³ - so that's Salesforce on the public cloud. It has the trust layer that I alluded to before. That is the interface between the model and the Salesforce world. And the trust layer is just the way to make it trustworthy.

Above that is Data Cloud. Data Cloud is formerly Salesforce CDP Data Cloud - Data 360 we call it now. And Data Cloud is a requirement. You can't really use Agentforce unless you're using Data 360. They're linked. And the reason is because the data that Agentforce uses is accessed through Data 360. Data 360 is a way to harmonize data. So you take it from many different

sources. And increasingly now you don't actually import data into it. You can just point to information that's sitting in Snowflake or Databricks without copying it, which is very powerful. Data 360 is the unified profile layer. And that also has the unstructured data component I was talking about earlier. On top of that is Agentforce itself, the platform. And that has different modules to it, but one of which is templated agent building. And then there's also like a testing center, a way to test agents in a safe environment. And then through Mulesoft a way to monitor agents that are both Agentforce agents and agents you build elsewhere. You can monitor them all in one pane of glass. So you have a visibility into your agentic workflow.

SS

Do you have to take Salesforce CRM along with Agentforce and Data Cloud?

MK

No, Agentforce is reliant on Data 360 but it's not reliant on anything else. You don't have to have CRM or Marketing Cloud. We definitely have customers who only use Data 360 and Agentforce. So you can have Klaviyo or Marketo or whatever accessing the data through Data 360. The platform itself has changed over the past four or five years and became more composable. The idea is that all the products run on top of this platform, including Marketing Cloud, Commerce Cloud. It's supposed to be a system that works together easily.

SS

There's what you call the Einstein trust layer - why is that so important?

MK

The general idea behind the trust layer is you have your own enterprise data, your customer information, your account information. This is very precious to you. This is something that you have that nobody else has. It's the lifeblood of your company. How can we protect that but still make it available to Claude or make it available to GPT? If you provide access to your own data then it becomes much more powerful. You could do personalization and all kinds of amazing things. Allowing that to happen is what the trust layer does. It uses RAG - retrieval augmented generation - tokenized PII. We have agreements with the frontier models that they will not use the data for training - and there's all kinds of checks and balances about what's coming back.

SS

So it's a technical way to prevent ...

MK Data leakage. And that's a big deal.

SS Especially when people are skeptical of LLM supplier ethics. Another big component - you reference it as the brain of Agentforce - is the Atlas Reasoning Engine. How does that work?

MK Atlas Reasoning Engine is the place where the agents apply their intelligence. When you think about the way these models work, you send them a prompt and then you get a response back and that's the loop. Prompt and response. But an agent doesn't just do that once. It'll do it once - get a response - maybe run another kind of subquery on that to get even better, back and forth, like a conversation. The Atlas Reasoning Engine is the part making the decision on how the prompt is changed - how to optimize the workflow to achieve the outcome. It's a hive of agents swarming on a particular problem, making decisions in real time. It's the thinking part on the Salesforce side.

SS Take me through the steps of actually building your own agent using the Agentforce platform.

MK There are different components to it. The headline would be the topic. So you describe the agent in simple terms and then below that you give it instructions. So instructions are explicit things that it has to do, as well as things that it can't do. For instance, make sure that you get the email of the customer and if you don't have access to it, ask them for it. And then there's guardrails. So things that you can't do. And some of that is actually baked in - like toxic language - all that is filtered out behind the scenes. And then how are you going to measure success for the agent? What counts as a successful outcome? That's the basic framework. It's topics and instructions. The next step is testing it and then improving it and so on. All of that can be handled in Agentforce as well.

SS So it's easy for the end user to step through the whole process.

MK And there's guidance. We created an agent to fill in the template. So you can just describe your project and it will do the topic and the instructions for you - give you a draft. And that has helped with adoption.

SS Let's move on to how companies actually deploy Agentforce. Where to even start. And you make the point

that it should be work that's routine and crying out for automation. But you also make the point that it should augment the work rather than simply replace it. How do you prioritize?

MK Any decisions that need to be made at scale. For instance, if you need to decide how to personalize the message to each of your customers, that's something that could work nicely for an agent. Anything where you're just held back by the capacity of your human team. In my agency days, the bottleneck was always the creative team. You had to wait for them to finish the image ... wait for them to write the copy. So wherever you're just constrained by the size of your team, that's a good place to look. And then the other one is, what work is routine? This is more around efficiency. What work is repetitive? And I think you just ask people which parts of your job are repetitive. So you can put an agent in the loop there. In marketing there's a lot of cutting and pasting and putting things in Excel.

SS Let's talk about the other issue that companies are facing, which is the escalating, in some cases out of control, costs of AI. How does marketing prove the investment is paying off?

MK This comes up all the time. We had the issue too at Salesforce. We have a marketing pipeline. A lot of channels, a lot of leads. We couldn't really figure out if the event drove the sale or the sales call or the email, webinar or whatever. But we use Data 360 to organize the data story and then you can put AI on top of that and build a nice algorithm model and come up with something that's 80% accurate. Building the case for marketing is pretty much the same problem it was 10 years ago. Get your data story and then apply the analytics.

The cost issue is significant. Once you start using AI, you realize you can burn through your budget in a month very easily. There are ways to manage it, and I know Anthropic is trying to help people be aware of their token consumption. But that's an art and a science. It's consumption based pricing now. You pay for what you use, which is great. But then in the beginning of the year you need to set your budget as a customer. So how do I set my budget? Predicting that is very difficult.

SS

Do you have partners who are helping your companies accelerate their adoption transformation? Or are you helping companies work through a lot of these more challenging issues.

MK

Yeah, we noticed customers were trying to implement it and then there would be an issue and a lot of them just stopped. So we developed a whole new function called Forward Deployed Engineers. So once Agentforce has been sold, they are embedded in the customer to do exactly what you're talking about. Figure out the right use case, figure out how to build the agents, how to monitor the agents, just to help get them started. If you push through a little bit, you can make it work. It's not as easy as it looks.

SS

Just how far up the decision ladder will AI creep? Will it decimate the senior ranks of marketers eventually?

MK

No, I don't think so. There'll always be a CMO and a Chief Governance Officer. And then underneath them there'll be the Brand Czar, setting the tone for the brand and trying to just make sure all the agents are obeying the rules. Then there'll be the governance police and then there'll be the highly technical people who understand how the models work. Below that you will have new roles. Some of them are emerging now like agent orchestrators, in charge of trying to come up with the right use of AI. What does go away are entry level jobs, more single channel focused, single tool focused, jobs that are repetitive. Those things are quite obviously going to change quite a bit and those people will be

doing different things in the future.

SS

Are we a few years away from a transformed future or are we a decade away?

MK

It'll take longer than people think. I think 10 years is closer to it. There'll be obviously big changes. But the problem is never the technology. There's amazing technology that exists. But adoption in the real world is another thing. And there's just a lot of inertia, there's legacy systems, there's processes, technologies, workflow. All of this stuff. So I think it's a much slower process than we may think.

¹. Synthesia, "AI Statistics 2026: Adoption, Usage & Workforce Trends".

². Richard Feynman (1918–1988) was a renowned American physicist (he helped build the atomic bomb) who said that if you cannot explain a hard idea in simple words, you do not know it well yet.

³. Hyperforce allows the Salesforce platform to run on public cloud providers (such as AWS).



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.