

SERIES 11: EXTREME PERSONALISATION

Agentic AI: From automated to autonomous

GFX:

[AWS LOGO] EXECLEADERS

OLIVIER KLEIN: Hi my name is Olivier Klein, and I am the AWS chief technologist here in Asia Pacific and in all my day to day conversations it's always about Agentic AI lately.

GFX:

OLIVIER KLEIN

CHIEF TECHNOLOGIST, AWS, APJ

First, what actually are AI agents? AI agents are autonomous software systems that leverage AI to reason, plan, and then complete tasks on behalf of humans or other systems. They can fulfil actions end to end, like performing code reviews, compiling research, processing claims, planning a trip or maybe even managing complete enterprise applications. And the main point is they can act autonomously, to some extent, which brings me to the first thing to be aware of.

Generative AI has been a great advancement that allows us to be more creative in both the visual and linguistic space. It is however output oriented and generally requires you to prompt something against the AI. In this next evolution of Agentic AI that we're building AI agents can act completely autonomously and they are goal oriented. They basically break down the tasks into smaller steps effectively mimicking human logic and reasoning, and then either way execute those step by step tasks by themselves or hand it off to other agents to complete that job. And this by the way is already a reality and it allows us to have AI take actions on our behalf.

Now this brings me to my second point. In the future, and to some extent already today, our end customers might be using agents to complete their day to day tasks. And one thing that made that possible is that AI now has visual understanding. This probably started with Anthropic's feature called computer use which allows an AI model to control a computer screen and reason from it. Just like that video that you see playing here the model autonomously navigates through a Google sheet within a web browser. It then opens up a web form and extracts the right and relevant information from it. It then eventually fills that into another vendor request form. And all of that within a standard web browser interface which was actually meant for humans to fill out, not artificial intelligence. And again everything you see here is clicked and typed in automatically using an AI. And that effectively allows AI to use any system that wasn't even programmed or meant to use prior from an AI.

And we also recently launched a new Amazon Nova capability called Amazon Nova Act. With Act you can give an AI model intelligent control of human interfaces. Just like what you see here, Nova understands the interface and tries to then think about using that interface to complete the task it was asked to do. In this scenario it's using Google Maps to find the biking distance and time between two locations.

Again allowing AI to think and act on a graphical user interface that was designed for humans to actually use. And it obviously doesn't stop there. We've already started seeing some of our customers implementing this new way of interacting with websites and apps. Just like we said before and predicted, a good chunk of traffic will come from AI agents and not humans anymore against your applications.

What you see here is Perplexity, an AI powered search engine, but more specifically you see the Perplexity Assistant which can become your default assistant on your android phone. What is really interesting here is that it is multimodal, meaning it can take in text, images, videos, but more importantly beyond just referencing its own knowledge that Perplexity has it can interface with other apps on your phone. This means you could say I want to go home for example and it automatically would open up your Uber app, book you a taxi via that third party app, without integration to happen just to the app that was interfaced by the AI here. All that to say, the future engagement channel of your customer might not be your application but rather an assistant that they use that is talking to you.

This goes back to the old saying of the promise of AI is no UI. And now my third point is about extending AI capabilities into many systems.

First, this means it's not one model that's going to rule them all. It also means that we start seeing multi-agent collaborations, specialised AI agents that are specifically trained to do one task.

Such as for example what you've seen earlier of using a graphical user interface. And as part of this shift towards autonomous systems we're also seeing a rise in multi-agent systems that can reason and collaborate with other AI agents and humans to accomplish more complex tasks. But apart from just training agents together though we also want to integrate agents into other SAS platforms.

This is where another interesting new protocol comes in which Anthropic recently proposed, the model context protocol or MCP in short. The idea is that we want AI agents to talk to our SAS platforms that we are using. Maybe your Slack channels, your Salesforce records, or even your Google Drive. Now instead of building integration for every agent into all of these platforms we can use MCP, a standardised way for AI models to connect and interact with external tools, data sources, and services, in a simple and unified manner.

I'd say think of it like a USB-C port but for your AI. Just as USB-C allows to plug in many devices and connect them easily to your computer with one plug, MCP allows AI systems to access various external resources to a single and consistent protocol without needing to custom code for each of these connections. Pretty cool hey. Okay time for my fourth point now.

I believe that Agentic AI will allow us to be proactive rather than reactive. Agents should anticipate needs before they even arise.

So let's take an example, Alexa+. Alexa+ is proactive, and just like a real live assistant it improves and becomes more personalised over time. You can even delegate tasks and it will take action on your behalf. But look instead of me talking through it let's have a quick look at how it works.

And this is a wonderful segue into my fifth and last takeaway about Agentic AI.

Agents can become like assistants that you hand work off to do. But the final decision might still be done by a human. And we start seeing many companies build onto this idea of having an AI complete tasks in the background for oneself but with traceability and visibility of what's actually happening. Take for example Manus. Manus AI is an advanced autonomous artificial intelligent agent launched in early 2025 by a Chinese start-up. Unlike traditional AI, that needs constant human input, Manus can independently understand, plan, and then complete complex tasks from start to finish.

It acts like a digital executive managing a team of specialised sub-agents to break down and handle those multi-step workflows such as hiring, such as data analysis or apartment hunting. Manus integrates then multiple of those powerful AI models to enhance its reasoning and decision making, works asynchronously in the Cloud, and notifies the end-user, you, only when the task is done. And as you see in the video that's playing here it allows us to visually track what the AI is doing. It evaluates every step, if we want to remove for example that black box effect by seeing what's going on of using that AI, and we now fully understand how it comes to those conclusions, and obviously to the final output eventually. And this is great because now we can treat AI agents as our new assistants whilst having control over what they are actually doing. So those are my five things to be aware of right now when it comes to Agentic AI.

GFX:

[AWS LOGO] EXECLEADERS