

Kush-Gupta-Senior-Vice-President-and-Global-Head-Professional-Services-AI-day-Session-10

Hi, my name is Kush. I take care of the professional services business for Hexaware. I've been with Hexaware for 25-plus years.

So what we're going to cover today is how we are creating non-linearity in professional services. See, this is an industry where it has gone past demos, experimentation, and AI is actually very much in practice, right? I'll start by giving you examples of areas where we have started creating non-linearity. So a typical audit process, let's say for a large enterprise client like Coca-Cola, which has offices in more than 200 countries, it takes anywhere between 20 weeks to 24 weeks and you need like a bunch or an army of auditors, 250-plus, to complete that audit process.

Now we are infusing audit AI in all the audit processes and what we are seeing is that such complex audit process, the timelines are shrinking from 20 weeks to somewhere between 8-10 weeks and the number of auditors that are required is also shrinking to somewhere between less than 150 or so. And plus the quality of audit is also getting better and better. Typically in an audit process you do a sampling of all the financial statements.

Now you can do a full population testing with AI because that physical boundaries of timelines and people going through that manually has gone away. This is creating a big opportunity for all the big audit firms to go into the SME market which they have never looked at, right? So it's creating a kind of non-linearity for them. Similarly in taxes, right? Hexaware has created one of the world's largest tax chart platform.

Now what this platform does is it kind of ingests data from multiple sources and it's able to pre-fill significant amount of information that is required in a tax form. As a result, more than 40% of the time is getting reduced in just filling up your taxes. The third example and this is in continuation to what earlier Sanjay had talked about.

For a legal firm, a typical challenge in an e-discovery process is that most of the client data is on-prem or on their private cloud. So your traditional LLMs are not the optimal model to use. So we created a legal SLM and as a result of that, the case intelligence and the summaries that were created, we were able to get a significantly higher precision, more than 77%.

And the amount of manual reviews that was required was less than 5%, right? Now I'll talk about the scale at which we are playing. Keech already mentioned that more than 50% of our revenue currently is all AI influenced. The core AI work that we do, which is agentic in nature, our revenue is at 17% currently, okay? Our developers today are consuming more than 75 billion tokens every month.

They are utilizing these tokens to do code development, automation testing, deployment and building all the agentic applications. Hexaware is building one of the world's largest enterprise AI platform. Currently, this platform is already used by 1.8 million business users across 180 plus countries.

And it is processing close to 250 million transactions per day. This platform is also the largest consumer of Azure tokens in the world today. We already have more than 150 AI agents in production.

And these are complex AI agents that I'm talking about. So typical examples include evidence extraction agent for audit, tax interpretation agent, your M&A due diligence agent or content creation agent. So all these agents are already in production.

There can be no scale if you don't have proper execution. So as a unit, we are hyper-focused on making sure that we do all the execution right. More than 50% of our AI projects are all outcome-driven.

So clients are no longer buying just billable hours from Hexaware, but they are buying billable outcomes. These outcomes could be legacy modernized applications, zero tech debt or agentic applications. Currently, we have more than 100 plus AI projects getting executed.

Now, I believe this is a great differentiator for us because it's helping us build the necessary delivery muscle. It's helping us build the reusable components, the accelerators, and bringing a lot of predictability in our delivery. Now, just going by the zero theme that my friends had earlier said, I believe Scrum is going to be zero.

People will not execute any project in Scrum anymore. We are already seeing more than 70% of our projects, all AI projects are being done in Kanban, and in next quarter or so, it will be 100%. See, all the AI projects are very non-linear.

Requirements keep changing very frequently, and as a result, the refinement, planning has to be done on a daily basis. So to manage this flow, Kanban is the right model. Otherwise, you will always be restricted by the sprint that you have to follow in a Scrum model.

Our workforce are already AI ready, most of them are certified, but it's not just about technology and tools that you need to train your people. It's also important that your people understand how an AI pod works because, hey, AI burnouts are real, right? Mistakes are happening because of the pace of change, which is extremely high. So an integration lead becomes a very important role in this AI pod, right? So training your people to know what an integration lead does is extremely important.

And what we are also seeing is you cannot have more than, sorry, more than three developers in a pod. We are hyper-focused on our partnership score. We do believe we have to tie up with the best partners in each service line that we operate if we have to take a leadership position in professional services.

So for Databricks, we are their number one partner in professional services. We plan to publish at least one solution in Databricks marketplace every quarter. And our approach to Microsoft is also similar.

Microsoft is also the number one hyperscaler in professional services. We are their number one partner today for Big Four. We have recently partnered with Factory.

There was an announcement that was made a month ago. We evaluated Factory on the basis of three key requirements. The first one was, can a non-developer develop agentic AI

application? Second was, what is the quality of code that comes out after you give proper context? The third parameter was, you know, does it provide the necessary flexibility to choose the models of your choice based on, you know, the kind of clients, what their preferences would be, right? So that's a view of what we are doing in professional services.

I know we are short of time. So I just want to say thank you and invite Jit, who will take the session around innovation.