

Infosys Limited Investor AI Day 2026

February 17, 2026

CORPORATE PARTICIPANT:

Mohammed Rafee Tarafdar
Chief Technology Officer

Mohammed Rafee Tarafdar

Over the next 10 minutes, I will cover how we are going to power AI-first services and AI augmented services using our IP and platforms. Now earlier during the day, we heard about the complexity with enterprise landscapes. And when I think about enterprise landscapes, I think about city maps. Every city is different. The map for each of these cities are very different.

Now if I bring any AI model, any tool any platform, the only way to accelerate is by creating runways within an enterprise that can help us accelerate AI adoption from pilots to hundreds of projects, and that is where Infosys Topaz comes in. And we do it in 5 different ways.

First, we have created a rapid experimentation and innovation infrastructure where our teams working with our clients evaluate the latest developments that are happening in the AI space. They look at all the noise that is happening and identify tech that is relevant. They then build proof of values that are very relevant for their business and then they demonstrate the art of possible in a very rapid manner. And today, there are about 39 such innovation labs that we are running with our clients across the globe.

Second, we take a very value-centric approach in how we look at the end-to-end process, because over the last few years, we have realized that use cases cannot deliver significant value. And this is where we are bringing the 25-plus industry blueprints we have in order, to come with already reimagined business and IT workflows that we can use with our clients to accelerate. So today, when our consulting teams, what they do is they sit with our clients, they understand the problems.

Then they use the product discovery and vibing tool from our Topaz fabric to very quickly identify good solutions, build a prototype, and then using exponential engineering, they actually create a production scale application by end of the day. And then using this, they are able to demonstrate how we can reimagine the complete workflow.

The third, Nandan talked about creating an architecture that is very evolvable. Now what we have done is, the way we have designed our IP and platforms is to make sure that we give optionality to our clients. They can pick any model they want. They can pick any agent framework they want. They can run on any AI platform. They can run on any AI cloud, and we can integrate with any AI-native tool that they have partnered with. And today, in most of our production deployments, we have a number of varieties that today we are already supporting.

The fourth runway is to build that enterprise context. Now here, we are doing 2 things. One, based on the work that we are already doing with most of our clients, we have built an enterprise context. Think of it like a map. Whenever I want to navigate in a city, you need a map which tells you how to go faster. So we have built this enterprise context or map that tells me how the systems work, what

the infrastructure looks like, where the apps are running, where the data resides, and how they all connect.

And on top of it, we are building an industry context. The industry context tells us what happens within a retail, within a bank, within a CPG context. And these are the models that we are bringing out of the box through a graph technology, and we are building these enterprise twins. And this is what will enable us to accelerate. Eventually, in most enterprises, to drive projects at scale, they need multispeed IT governance projects, so that they can onboard these AI tools at speed. They need to put these guardrails, and that is why we are building a lot of tooling that enables them to deliver these AI solutions in a very trusted manner. And all of this comes through Infosys Topaz.

Now in the IP and platforms that we have been building at Infosys, we have always kept our customer needs in mind. So if a customer comes and says, look I want an end-to-end vertically integrated AI and agentic platform, then we use AI Next as a platform to accelerate value. Or if the customer says, look, I have already made some investments, I want a composable modular agentic and AI platform which can help accelerate my own AI journey at speed, then we essentially bring the Infosys Topaz Fabric. So with the combination of these, we are able to meet most of the demands of what our enterprise customers have.

Now let me talk about Topaz Fabric itself because over the last few months, we are starting to integrate all the different IP that we have at Infosys, into one common way through which we can deliver our services. Topaz Fabric enables 5 key capabilities. The first is this builds on our customers' existing investments. So this is not about replacing what they have. So this works above their model layer, above their platforms, and above their enterprise systems. And Fabric can integrate with any model, any framework, anything that they have. So that is the abstraction that we have already built within this.

Second, it provides close to 600 agents, which have been purpose-built for different AI-first services, AI augmented services and also for industry-specific flows. So this is something that we bring out of the box to accelerate the journey for our customers. Third, what we have also done is, we said we will create out-of-box integration. So we have out-of-box integration with most of the coding tools, . We have out-of-box integration with different models. We have out-of-box integration, with business platforms like SAP, Oracle, Salesforce. We have out-of-box integration, with data platforms like Snowflake and Databricks and with enterprise platforms with ServiceNow.

Now what this enables is it enables us to deliver value to our customers quickly. Now in all of these, it is also about bringing the enterprise context and hybrid intelligence. So the way we are doing is we are starting to build a number of different ontologies and models that comes pre-packaged as part of our Topaz Fabric, and that is something that we bring out of the box. And as we deploy it, we

learn from the data, we learn from the processes. And this is how we create a closed feedback loop where the context keeps improving as it gets used over a period of time.

Now lot of our clients also want to use a lot more predictability in the way AI is deployed. This is where we bring a lot of deterministic rules, couple it with AI models, and we also bring our own small language models in order to create a right value proposition from cost as well as from a time-to-market standpoint. Now all of these is backed by a lot of deep research and patents that we have filed over the years.

Now while we are doing a lot of innovation internally inside within Infosys, we also acknowledge that there is a lot of innovation happening outside. So we are working with AI native partners in 3 different ways. One, if they have a platform that is really good at doing something, then we are leveraging it for the tasks that are relevant to enterprise. For example, Nandan talked about brownfield. So our Cognition partnership is largely to use Devin for a lot of brownfield engineering, because we find that, that is really good at it.

Second, we are building embedded agents that can work within our partner tooling. So today, we have built agents in Fabric that can run within Claude Code, that can run within GitHub Copilot, that can run within ServiceNow. So whichever platform the client has, these agents work within that environment.

And third, we have integrated with their tooling, so that we can cover the end-to-end value chain that is required to accelerate the journey. The next is also we are focused on where the industry is heading on AI, and this is where we are working with universities to do joint research. We do today on agentic technology, on also scaling, and on trust with Cambridge, with Columbia and Cornell. And we are also doing this a lot more with the research centers that we have set within Infosys. This is to make sure that we continue to build on what will come next.

Now let me bring all of this to life with an example where for one of our logistics client, they were finding that the customers were able to process the orders and bookings in a lot more accelerated manner, and this was creating an issue for them. And they said, we want to be able to process these customer care services in a much more accelerated manner. This is where we took AI Next as a platform, because they had already tried with multiple platforms and they did not work. And we said, we will use AI Next first to uncover the existing knowledge, because they had rules that are very specific to each customer. So we pulled out close to about 8,000 different rules that exist there.

The second innovation that we brought here is that we automated the entire workflow. So when we started, the extent of automation was about 0% to 10%. We took them to about 70% automation in their entire workflow. What this meant is the turnaround time reduced from 24 hours to about 30 minutes. The next is, they also were very concerned with the sovereignty of the stack. So what we

did in this case is we used the Mistral and Pixtral models to make sure that it addresses the sovereignty needs that they had. And this today supports orders or bookings from 116 different countries, and it supports 15 different languages. That is the power of what this could do.

And eventually, as we started scaling, cost became an important driver. So we had to bring in a lot of optimizations to reduce the cost significantly for that customer. And this is something that we did over the last 1 year, and today, this is live. Now you can see a lot of these IP platforms as you come to our living labs, and I encourage all of you to please spend some time and experience these technologies that we are talking about. Thank you very much.