Infosys Limited
Investor AI Day 2026

February 17, 2026

CORPORATE PARTICIPANT:

Satish H.C.
Chief Delivery Officer

Satish H.C.

Let me set the context before we deep dive into our Infosys AI Playbook. Every tech shift, whether it is PCs, internet, cloud, digital, each one of them led to rewiring of enterprise work and workflows. And AI is the next rewrite. Here is a typical enterprise landscape - sounds complex, this is one of the simpler ones. A typical enterprise will be far more complex because of the scale, because of the fragmentation, internal and external, the heterogeneity and the divergence of the operating model and the regulation with which it operates in, which is by design by the way, and of course, the technology debt.

So, AI in enterprise is not just about low-hanging fruits like localized efficiency or user productivity. Integrating AI in an enterprise is not just a software upgrade or a plug-in. If it was that simple in such a huge complexity, no surprise that AI projects fail. Then what is it about? It is about harnessing the full potential. It is about re-imagination of systems, processes and deeply embedded ways of working and rewiring legacy power structures. So, Infosys has done this enterprise rewrite with each of the past tech shifts, and we are now doing the next enterprise rewrite with AI.

So let us look at how is the enterprise tech stack changing. Typically, it consists of the systems of record, which deliver deterministic, programmatic capabilities, which are codified, structured processes and enforce policies and this delivers governance, accountability and compliance. We have systems of intelligence which facilitate engagement, collaboration, transactions, but usually, it is the humans who engage with data and the workflows.

Above the enterprise stack, lies a vast non-deterministic or non-programmed flows, these are unique, unstructured, they need novel problem solving. It needs experience, gut feel and is usually handled by humans. This is the layer, which is under served today, and this is ripe for AI-led transformation. So, it is a myth that enterprises need just 2 layers, which is AI and data, it is enterprise and algorithm.

How should we harness AI then? AI is not the end game of tech transformations. It is just another one, but it is a giant leap in raw capability, and it is not system complete. So, we need a multilayer transformational approach and purposeful orchestration to harness the potential of AI. This is why AI diffusion in an enterprise lacks the rate of AI adoption, and it needs time.

So how is AI transforming the enterprise stack? What does this enterprise rewrite about? The core intelligence of humans plus machines will be seamlessly shared across the layers, so that every layer gets reinforced within the enterprise stack. The systems of record need acceleration, so that the business and operational processes can become more efficient.

The systems of intelligence needs a seamless integration of structured and unstructured data, so that intelligence can be wired into user journeys, business processes, transactions and decision-

making. It is estimated that 60% of effort in an AI project goes into doing this. The models will come in later and this requires deep industry knowledge and context of the enterprise.

Encoding intelligence and AI in the flows leads to more automation and autonomy. And this is leading to the development of a new layer within an enterprise stack, which is what we call as the systems of cognitive work within an enterprise.

As an outcome, humans will shift from acting on data directly to a governance and oversight role on flows and decisions. All the new AI tools that keep coming at us at fast pace will get plugged into the systems of new cognitive work, and it will accelerate the re-imagination of an enterprise or its flows and decisions. Infosys will unlock tech debt and complexity and harness the power of AI to be enterprise grade and bridge the evolution adoption gap and expand our addressable market.

So how do we monetize this? Our playbook reflects the structural changes necessary in our industry, so that we drive value at the intersection of intelligence, engineering and domain. With AI, capability is a commodity, because it is available to all.

Sustainable mode for an enterprise can be created by deep integration in specialized workflows and unlocking unique organization knowledge and context. So every enterprise has got unique data processes, risks and complexity, and AI will not unlock uniformly across the enterprises because of this variance. So our client intimacy and deep understanding of our client context will help us, mitigate this and unlock value and also drive culture and change.

So we have built an engineering approach in our delivery where we can codify enterprise context, which will help accelerating scaling of AI. And this also enables enterprises to retain and protect this unique enterprise context within their four walls of the enterprise, so that they can keep their competitive differentiation, and this is not diffused into the AI models. Our depth in engineering and frameworks on IP and patents will accelerate AI readiness and adoption with our Topaz Fabric and our specialized talent in the form of full stack and FDEs.

For a financial services client, they were looking for an AI partner. When we started talking to them, we realized that they have a very strong enterprise AI platform that they have built. But then what we realized was they had an adoption gap, and we pulled out our agent control framework, as we call it, which would address the quality of code that is being generated, which would address the AI slop, which led to a poor adoption within their organization. So now we are working with them on taking our framework and fortifying their enterprise AI platform, so that we can accelerate that journey.

We have embarked on a talent transformation journey to build an ambidextrous workforce, which is deep in engineering and creative in reimagining work and workflows from first principles. We see new opportunities with domain stack. We have over 25 years of industry-focused experience and

when core intelligence connects with agentic economy, the play of AI elevates from how work is done to new outcomes that are possible. So we are invested in building the domain stack powered by our depth in domain and knowledge of how to deeply integrate AI tools and plug-ins into the flow.

We have created strong differentiation in our services stack. AI is now integral to how we deliver every service. Our services stack is powered by Topaz Fabric. We now have an approach to productize and reimagine work and workflows that will lead to a human plus agent model. We are also seeing momentum with new deal archetypes, legacy modernization with reduced risk, higher predictability on cost and accelerated timeline, large deals with integrated ops tech and transformation wired in, organization transformation encompassing enterprise stack and people.

This includes the AI-first DCC approach, which we have pioneered in the market. We have also elevated our play to take end-to-end accountability from strategy to actionable roadmap to execution and eventually outcomes. We have expertise in delivering both above the line, which is business value and below the line, which is efficiency. Infosys is best equipped to deliver enterprise AI ambition with the power of our client intimacy and our AI playbook.

A quick example. Here is a client, a CPG, who had an ambition of clocking about growing their revenue to $7 bn. They came to us to bridge their ambition and deliver an AI operating model, so that we will have an actionable roadmap and execute to it with executive AI value office, along with managing risk, governance and assurance. We used our Infosys IP and built their unified data foundation. We build their enterprise Agentic AI platform with the requisite guardrails. This enables them to rapidly innovate and diffuse AI across their functions. Today, they have 10 Agentic AI products in the business across different functions from R&D, sales, marketing.

Above the line, with the agents that we developed in R&D for product formulation, they now have line of sight to $50 mn revenue, which they did not have before. Below the line, we have been able to unlock $25 mn cost savings through just optimizing operational work. And then beyond this, we were also able to deliver 40% of business productivity improvement in functions like procurement and marketing. Thank you.

Infosys Limited
Investor AI Day 2026

February 17, 2026

CORPORATE PARTICIPANT:

Dinesh Rao
Chief Delivery Officer

Dinesh Rao

Thank you, Simran, and good morning, ladies and gentlemen. Thank you very much for sparing your time here today. As we navigate the landscape changes with respect to AI, I think our priority in the last 6 months to 1 year has been how do we accelerate our customers from experimentation to really looking at an AI scalable at industry level.

Now AI as a technology is as good as what it can really understand from orchestration across systems of record, which Satish alluded to, ability to really understand the complex business processes and navigate. And most importantly, how does it even get to understand the deep-seated legacy systems that are there. And frankly, looking at all the complexities that we had with respect to some of the estates that we have been working as well as with the customers, we decided to codify our entire services across 6 strategic pillars. And these 6 strategic pillars are all integrated. In a way, a customer who wants to start the journey of AI adoption has to look at each one of them. Our intent is to take this strategic pillar to walk them through this entire journey of scaling the AI.

Now let me start with the pillar number 1, the AI strategy and engineering. Nandan alluded to it, so as Salil. Organizations are extremely complex. One needs to really look at the top down, looking at how do I really create an AI blueprint, a need of really setting up an AI transformation office. What would that mean? I need to first understand which business units, which business processes that I would be able to unlock the AI value. You cannot really be spraying across the AI in multiple different business units. Now that is a very key thing because that is where the value realization framework of Infosys comes to bear to look at how do I map the process to the value that it actually arrives with.

Now the second key thing is how do I change the ways of working. Because you have so many models, so many platforms that have come by, one needs to really put together a very technology-centric AI platform with the models that it needs to go with, the governance that it has to happen and how do we make sure that we diffuse this particular platform across multiple layers of the organization, so that we have one standard way of really looking at how do I scale the AI.

The third layer, obviously, is to really look at the entire governance model that you need to put together. In a way, this AI has gone very strongly. The last one, obviously, is to really look at a purposeful selection of an AI infrastructure where you also need to continuously keep an eye on what would happen to my AI Ops.

Very recently, we have been working with the CEO as well as the Board of Danske Bank, where we really started working with them to looking at how do I enable this bank towards the journey of being AI first and digital first. So we set up an AI office along with the Board, working along with the CEO to really put together a strategy all the way to engineering, augmented that with innovation labs as well as identification of core processes like KYC, the fraud detection, credit and what not. So in a

way, the journey is just not that I would start doing software development, but it needs to have a top-down view of really setting up the strategy for an organization.

Now the second pillar, my friends, data is everything that AI needs. Today, organizations have multiple pools of data sitting. It is just not structured data. It is, in fact, even unstructured data, multiple modes in terms of videos, in terms of speeches and what not to an extent of about 60% to 80%. And today, the most of the time is spent on data in all our projects. Data is one which accelerates your AI journey or potentially could also decelerate your entire velocity of how you are going to do it.

So our framework here today is to really look at help our customers try and harness the data, transform the data, bring them all under one uniform data fabric. And it just does not sit there. At the end of the day, data also have to drive intelligence. The intelligence can be driven only by connecting the semantics and the ontology on top of it. And that is a very elaborate process one needs to go through, because the processes are different by region, by business units, and it has been all codified in systems of records as well as in systems of experience. And that is a humongous job.

The third layer there is to really look at how do I govern the data, because the data fingerprinting is extremely critical as you really look at who would get the access of data, so that in some sense, there is a framework if you really look at. So in one of the large industrial manufacturers, there was 10 petabyte of data that we had to really bring in all together, harness it as well as create the semantic ontology today, help them actually drive the supply chain optimization by over 20% to 30%.

Now the third pillar, Nandan alluded to this, this is all of reimagining the entire business processes. Most of the large enterprises today have point-to-point solution. One need to really look at in the context of domain and reimagine the entire business processes. And these processes have to be reimagined with respect to how the human intuition works along with the AI and agents. And it is extremely critical that every workflow by persona has to be reimagined and has to be codified the way AI would actually come by.

Now this also has to be contextual to the business and the regions where you really work by. You just cannot take a sourcing and a procurement process and say that it is a domain that it could be applied to every industry. It has to really be contextualized to the industry and the business that you are really looking at. And most importantly, since it is going to change and touch every one of us, it has to also be looked at how do I operationalize the entire workflow, the technology and the change management, all put together to really help realize the end-to-end value.

We have been now with Toyota Motors Europe working through a supply chain transformation process, where we took our industry asset of automobile on top of it, the entire agentic playbook, really looked at by persona of a buyer, a planner or a customer service agent and really double-clicked to look at what does this transformation really mean. And just to look at one critical process

of drop ship, which is so critical for an automobile -- is to really bring in agent and orchestration there to really take away all the manual work and bring in much better inventory visibility.

Now the fourth pillar, obviously, is modernization and everybody has talked about it. It is the an Achilles' heel of our large organizations. Today, there is so much of tech debt sitting there. The code is obsolete. There is no written documentations. There are no availability of SMEs today. And it is not that the customers did not really try to do the legacy modernization, but the ROI and the time it took never stacked up in the past technologies. So what AI models as well as powered by our Topaz Fabric today enables is to make sure that we transition some of these large legacy, both on the data as well as on process side, into the most modern cloud-based microservices architecture. And today, it does stack up, and you would hear one of the examples very soon.

Now physical AI is something we believe is at the cusp of really accelerating the AI journey. Whatever intelligence that we all thought and built as a part of the digital workflow is finally now moving towards the actual physical objects. And this, friends, helps the acceleration of AI in a lot of the products that we really look at. Some of the key cases that we really look at is -- and in terms of data first, the new product introduction, wherein the entire process would be reimagined as well as infused with AI, the products have to be defined or designed with AI in the front. And with more and more products with the software bomb being larger than actually the physical and the mechanical, we have a huge play in terms of embedding the AI as we look ahead.

The second case is the intelligence today; the real-time intelligence is moving from cloud to the edge. Now this will help accelerate the decision-making at the edge, which means vehicles, the industry operations, running infrastructure, all of this would actually increase the advancement and usage of AI. And lastly, the autonomous systems today, the prevalence in a lot of industries as well as areas is continuously keeping on increasing. Towards that, we see that we accelerated the journey of actually infusing the AI in the physical.

And here, I also want to draw your attention to the two acquisitions that we did, one InSemi, which meant the silicon design as well as validation. The other one in-tech, on the automobile, directly fits exactly into this particular pillar where it would help us bring more context as well as acceleration in terms of enabling the physical AI.

The last, ladies and gentlemen, is not the least and the most important is the trust and the governance.

If I today ask everybody here, who has used the AI? I am sure that all of you would raise your hand, but how many of you really trust the output that came by. I am not sure whether everybody would raise their hand, because there are hallucinations, there are model breaches, and there are also governance issues with respect to the new AI Act and etc, coming in. So in a way, the trust has to

permeate through all the other 5 layers for us to really make sure that we have the output that comes in an enterprise, which is trustworthy. To me, the trustworthy AI in every enterprise would be a huge differentiator. So we want to build the trust to our customers, and we want to monetize the trust for our customers.

Now having looked at all these 6 important strategic pillars, right, I just want to dwell on one of the cases, Nova Chemicals. They are a large petrochemical manufacturer based out of Canada and U.S., and it is very asset intensive. As you know, the industrial operations is extremely complex. If an asset goes down, they would have an impact on the top line and the bottom line. In the current context, most of what they were doing with respect to maintenance was all manual, logs and etcetera. So we were invited to a program in smart maintenance, where we actually brought in the data across their machinery, their OEM manuals, their maintenance manual, the historical data, the log data and etcetera, to help a planner to really, with a simple NLP on a chatbot would be able to guide them on what part of the industrial operations have to really go through a maintenance.

And most importantly, we were also able to actually bring in orchestration with agent AI, where the OT systems and the IT systems come together. So seamlessly, we were able to really move towards actually creating the entire work order process, the planning process, which actually moves over from OT systems to the SAP or the ERP that we have. And we see the impact of bringing huge planning efficiency, asset utilization and etcetera. Of course, here, we partner with Microsoft, and we used all the stack of the entire Azure to really make it come to life.

The last case I want to really dwell is about Hertz. I am sure all of you know it is a very large mobility organization where we are today, as we speak, embarking on really modernizing their entire reservation, their fleet management, their pricing and the whole thing, which today is approximately 3 million lines of code actually sitting on a tandem computer. And I would like to here play the video, let us hear from the customer on what their experience has been and what we have done with respect to this journey.

So the important point to note is, models are there, I think workflows are there, but our context of Hertz in terms of what their processes are, how the code has been written, how the existing architecture is and what is the new modular architecture that we need to really help them migrate to is the context that this iLEAD brings to bear. And that is very critical as we really look at this entire legacy modernization.

So that, ladies and gentlemen, summarizes the six value pools that we are talking about, and thank you so much.

Infosys Limited
Investor AI Day 2026

February 17, 2026

CORPORATE PARTICIPANT:

Balakrishna D.R
Head of Global Services

Balakrishna D.R

So my colleagues, Dinesh and Satish, talked about AI-first services, which is new value pool that is being created out of AI. What I am going to be talking about is AI augmented services. What we mean by that is taking our traditional services and infusing AI in that, and we want to be a leader in this space as well, and I will talk about how we are doing that.

So we have taken each of our services that we actually traditionally provide, whether it is application development, testing, modernization, migration, engineering services, business operations. So 20 plus services that we actually provide, and we have created detailed playbooks of how we can use AI in transforming the way we deliver these services.

When we are doing this, we are partnering with the best of the technology that is out there. We are working with leading models from Anthropic, from OpenAI, from Gemini, from Amazon Nova. And in fact, we are also working with open-source models, like DeepSeek, Llama, and we have even created our own coding model in this space, perhaps the only GSI that has actually created a coding model. So we use a combination of these models that in our engagements.

So in the Hertz example that you actually saw, we, used two models. We used a Claude model that actually generates the code. And one of the things about LLMs are they are better in critiquing output than generating output. So we used an OpenAI model to critique that output, and that is how we actually improved the accuracy.

So one is the models. The second is the tools. And in the tools, we are working again with the leading tools. GitHub Copilot was the first, and it had almost 100% of the market share. 2 years back, we set up a GitHub COE that was inaugurated by the CEO of GitHub. And then we are still ranked the number 1 GSI in terms of GitHub adoption. Just a few months back, we got this award as the leading GSI for working on GitHub.

But just not GitHub, we are working with Anthropic, we are working with Gemini, we are working on, on the models from Anthropic CLI. We are working on the new age models, Devin from Cognition, also we are working with Cursor.ai that you would have recently seen. So we are working with the leading models. We are working with the leading tools.

But then, again, as my colleagues also talked about, a lot of these models do not understand the enterprise context. They do not understand the standards that are there in the enterprise. They do not know other libraries that are there, other programs that are there. So we have to do a lot to actually bring that context into the tools that we are actually using. So, we do that by creating MCP registries. We create a knowledge graph of the enterprise context and we combine that. In the recent release that we just did with Anthropic, you hear Dario talking about that. They need Infosys to bring the enterprise context onto the models, and that is what we actually do.

In addition to that, we have actually created agents specifically for each kind of services. So in application development, we have taken the life cycle and said we need agents for requirements, design, architecture, etc. And we have created 100 such agents that we are using in each of our services. In addition to that, we have to create other tools. You saw in the Hertz example, the LLM models cannot digest huge pieces of code at one time. Some of these enterprises that we work with have millions of lines of code. If you give that to the LLM, they start to hallucinate. So, you have to shrink the code, you have to create graphs, call graphs on how these are actually associated, and that is when you get better output from LLM. So we have created all of these assets that are part of our Topaz Fabric, which my colleague, Rafee will actually talk about in more detail in the next section.

In spite of doing all this, we need talent. And sometimes people ask me, if LLMs are generating code, why do you need people? Why do you need developers anymore? And I think they asked the same question to Boris Charney from Anthropic, because Anthropic is continuing to hire developers. So somebody on X asked him, if Claude code can actually generate code, why are you still hiring software developers? And his response was that engineering is changing, but great engineers are a requirement still and in fact, the most important requirement going ahead. And that is something that we also actually believe.

The way that we actually deliver code or the way we support applications may change, but still you need great talent. And so what we are doing is to take each of our developers and training them on AI. So we have 90% of our developers that have been trained on AI. It will never become 100% since we are always hiring new people into the stream. But our intent is to actually have everybody be able to use AI in their daily work.

In addition, we need specialized roles, like forward deployment engineers that will create the platforms that the teams will actually use. And then we have created COEs for each of the partnerships that we have. So it is a combination of all of these that actually helps us deliver our AI-induced services.

It is the blueprint, it is the technology, our own technology plus the leading technology that is out there and the people that we actually create. The way we are going to market is also, as you saw in the Hertz example, is not about PPTs anymore, it is about actual demos, and that is what we see creates the impact, and that is how we actually go to market in all of our large deals.

I will talk about a couple of examples of how we are using it in actual programs that we are executing. What better example than Microsoft, who is in the leading edge of this kind of technology adoption, creating the technologies and also adopting this technology. So in Microsoft, we have a 360-degrees partnership. What we mean by that is that we go to market with Microsoft, we are one of the big customers of Microsoft, Microsoft is our customer. We provided services to Microsoft and we do multiple engagements with them, but I will give you a couple of examples of what we are doing.

Microsoft themselves are going through a big transformation. From enterprise agreements, they are going to what they call MCAs, Master Customer Agreements. What this means is that through the master customer agreement, they want to eliminate all the paperwork that they had to deal with in the enterprise agreement. They also are talking about evergreening the licenses. So enterprise agreements had only 2 years' timeframe. This is perpetual agreement that you can use for multiple years. Enterprise agreements had a minimum seat count of 250. This actually has no minimum seat count. In addition, in MCAs, you are able to monitor your usage, adapt your usage, you are billed based on your usage, you get a very flexible billing. So multiple advantages for the customers of Microsoft by using this agreement.

And also for Microsoft itself, because you are not having papers and documents anymore, it eliminates and accelerates the way they go to market and also the operations that they have, On this engagement, they had to build the IT system to manage all of this. So Infosys is actually working with them to build that. We used all the technologies that I actually talked about, and we are getting 2x developer velocity and 35% improvement in the time to market.

The other engagement that we are working with Microsoft is on their Intelligent Cloud. So, as you know, this includes Azure and Microsoft Office. They carry a lot of mission-critical systems of Microsoft customers on these clouds. And it is important for Microsoft that for these mission-critical applications that there is no downtime and then it is actually trustworthy. So, Infosys is again providing support for Microsoft on this. And the way that we have used AI, is that AI agents today monitor the logs and predict issues before they actually happen. And they give all of this intelligence, which we call triaging, and then we route it to a specific support engineer with the information, so that they can actually work on it before the issue actually happens. So this is for both reactive and proactive issues. And so you can see that we get 40% improvement in the incident response and 10x improvement in the RCA turnaround.

So there are several such examples. The other example that I have here is Danske. Danske has a Forward '28 Strategy where they are looking at modernizing their entire landscape. They want to bring process efficiency and they want to completely create it as a digital bank. So you must have seen the press release where they have chosen Infosys for this transformation, and we are helping them on the AI strategy and transformation as well. And we are doing everything from AI strategy to implementation. We have set up an innovation lab for them, and we are creating multiple AI solution work streams.

So we are using Agentic AI in the code development, more than 2 million lines of code that we have generated. But then again, these lines of code that we generate has to be validated, and that is what our engineers have been trained to actually do. And 97% of the engineers are using that. In addition, we have created multiple AI solutions. They wanted to use ChatGPT, but they wanted the guardrails.

So we created enterprise ChatGPT, which has over 16,000 users. They created other solutions on risk and HR, which have been quite successful.

D. R. Balakrishna

With this, I will hand it back to Simran. Thank you.