



Coforge Webinar **Operationalizing Enterprises** **Autonomy**

Thursday, September 17, 2026

MANAGEMENT:

Mr. Sudhir Singh

Chief Executive Officer and Executive Director

Mr. Anup Nair

Chief AI Commercial Officer

Mr. Ashish Kumar

Executive Vice President – Global Head of Cloud & AI Infrastructure

Mr. Deepak Manjarekar

Head of Data & AI

Mr. Lalit Wadhwa

Chief Technology Officer

Ms. Erika Moore

Chief Officer for Strategy & Growth

Mr. Niraj Patel

Senior Vice President

Moderator:

Thank you for standing by and welcome to today's executive briefing from Coforge. My name is Inba and I will be your host for today. Today's topic is operationalizing enterprise autonomy and the Coforge leadership team will be discussing the challenges that enterprises face when operationalizing AI and what it takes to bridge the AI execution gap at scale. First, a brief note on the format. Today's session will be 90 minutes long and we have allocated time at the end of the session for Q&A. We will be taking questions live, so please hold on to them until we open the Q&A session. To ask a question, please click on the Ask a Question button below the media player, enter your details and wait for your turn to be announced. After your question is answered, you may click leave to return to the webcast page. Participants are requested to keep your questions relevant to today's meeting's topic. Now to kick things off, I would like to introduce our first speaker, Mr. Sudhir Singh, CEO and Executive Director, Coforge. Sudhir, over to you!

Sudhir Singh:

As always, thanks a lot, Inba. Thanks for setting the context and a very good morning, very good evening to you across the world, ladies and gentlemen. At the outset, thank you for your time, thank you for the interest and we do plan to be exceptionally good custodians of the next 90 odd minutes that you have granted us to talk about something that truly excites us, the end destination of enterprise AI, which is enterprise autonomy, so welcome to the session. All of us are really excited to make sure that we sit down with you here, virtually at least, and talk about what it is that we are seeing in the market, what it is more importantly that we think is the end destination that we need to get our clients across the finish line they need to cross and I always have the interesting task of just setting the context and then stepping away so that you can really hear from the folks, who actually bring all of this to life. You do not need me to tell you this, you know this. Our industry is at an inflection point and I do not mean that euphemistically, I do not mean that if there is a word called journalistically, I mean that structurally. This industry, our industry is at an inflection point. Again, you know this, raw intelligence is commoditized. Open-weight models have accelerated this dramatically, of course. So, world-class AI with all of that is now downloadable, it is fine-tunable, it is approaching zero cost, that model name that we talk about is becoming a utility there is absolutely no denying it. This in turn, is creating something that we do not want to talk about, but we must, and we do. It is creating an AI deflationary squeeze across tech services. Intelligence is cheaper. Software creation is cheaper. Knowledge work is cheaper. Effort-based services are therefore facing this structural issue. But here is what most people miss, and hopefully some pieces of it we shall try and address over the next 90 minutes, that deflation is also funding the next wave of value creation. You read more reports than we have, analyst reports, and just about every report, the sum and summary of it is that every organization is trying to invest, is

becoming increasingly serious about AI, but very few have been able to truly scale AI and that gap between adoption and operationalization is enormous. We feel it, we sense it, and we come across it in every conversation that we have. I am going to be in New York later this evening, and I suspect it is going to hit us again out there and it is the single biggest opportunity in our industry, ladies and gentlemen, right now. That is our submission, as Coforge to all of you.

The competitive question has shifted to who can embed intelligence into real enterprise operations with real data, with real governance, addressing the real complexity, and delivering real measurable outcomes. That destination, as I said at the outset that enterprises need to go to is what we call the end goal, which is enterprise autonomy, the ability of an organization to continuously make and execute intelligent decisions across its operations with manual intervention. You are going to hear those two words a lot, enterprise autonomy, and hopefully we will try to make them real as we take this along. Today, we are going to show you how Coforge is operationalizing this across the data layer, across the agent operation or the agentic operations layer, across our delivery model, and in live industry implementations. So with that preamble, Inba, can we go to the next slide, please?

Great. Look at the slide. Let me give you a picture of what is driving our direction. The industry, we talk about it internally, at least in our own language. We say it has been shaped by four forces simultaneously. The first big piece is capability gap between LLMs is shrinking. Open-weights have made and are making, as we speak, intelligence utility, and that is already here. When every enterprise has access to the same intelligence, I said this about a month back as well, the enterprise's own context is what becomes a true scarce asset. Second, AI-generated code is creating technical debt at machine speed. Code is cheap to create now, but it still needs to be secured, it still needs to be governed, it still needs to be maintained, it still needs to be integrated and that remediation and governance layer is itself a massive recurring services opportunity for Coforge and for every player that operates across the tech services sandbox. Third, effort-based pricing is under structural pressure, and we think thankfully it is under structural pressure. Clients are increasingly unwilling to pay for hours and if they do want to pay for those hours, they want to pay significantly less for those hours when they can see AI doing more of the work. We live this, we breathe this, we see this on a daily basis, on a weekly basis, and on a monthly basis. The commercial model has to evolve for tech service players like us towards outcomes and that is a big shift and you have heard me say last quarterly call that 6% to 7% of our revenues for the last quarter were outcome-based and therefore, this whole adoption to scale gap that I talked about remains fine and as I said also at the outset, almost every enterprise we work with is using AI across one or more functions, but very few of them are able to

scale up. So then, where is value actually migrating as a result of all these forces, and you see this out here. It is moving from model access to enterprise context. Again, the same point, when everyone has the same model, the differentiator becomes your knowledge of the client's business, their data, their rules, their exceptions, and their institutional judgment that is where it is going to. Again, it is moving from, you can see this, from pilots to production execution. The ability to take something from experiment to operating at scale, operating reliably, operating within the governance guardrails. Three, it is moving from effort-based delivery to outcome ownership. I talked about this. Clients want to pay for results, and it is moving from horizontal AI tools to industry-specific autonomy and this, we think, is important. Erika will talk about this. Neeraj will talk about this. An airline's path to autonomy looks completely different from a bank's path to autonomy. It has to be built from the inside, industry-by-industry and therefore, how are we responding? We built three layers for this moment. Coforge Nuuron, which is our enterprise autonomy enablement suite and it helps organizations operationalize AI at scale you will hear about this. Second, Mod Squads, AI-native delivery pods that combine human expertise and AI agents around measurable business outcomes. That is how you industrialize, that is how you scale, that is how we industrialize, and that is how we scale. Finally, Momentum blue. Lalit Wadhwa, and please do watch out for that session, is going to talk about our Forward Deployed Engineer Unit. FDE is going to the client environment, they are the ones who go capture the context, they are the ones who prove value and they are the ones who create the foundation for everything else that follows for a tech services plan. Each speaker today is going to go deeper into these three areas that you see on the slide.

Inba can we go to my final slide and then I fade out. We have got six speakers, six colleagues of mine, incredible professionals and each of them is going to create, cover a specific part of the story that I was trying to hopefully weave through and hopefully been able to weave. Leading the charge is going to be Anup Nair, Chief AI Commercial Officer, Coforge and he is going to set the stage for what I talked about in enterprise autonomy with live proof points. He is going to be followed, Anup Nair, Executive Vice President, he is going to be followed by Ashish Kumar, Executive Vice President, Global Head of Cloud and AI Infrastructure for Coforge. He is going to cover agentic ops. What does it actually mean? What does it take to run agents at enterprise scale? And that is Ashish watch out for that session. Watch out for every session. He is going to be followed by Deepak Manjarekar, head of data and AI, covers data enablement and he does that because autonomy finally starts with data readiness watch out for that session too. Lalit Wadhwa, Chief Technology Officer, Coforge, is going to be speaker number four. He is going to walk you through Momentum blue, the FDE unit that I talked about and the forward deployed engineer model, what makes it real? What

makes an MDE an MDE and not another AI architect or AI engineer with a fancy monitor and fancy wrapper around it? How do we activate it? How do we scale it? How are we activating it? How are we scaling it? And then it is going to be Erika Moore, Chief Officer for Strategy and Growth for our travel, transportation and hospitality business globally. She shows you enterprise autonomy as we are enabling it, making it real inside airline retailing and finally, as always, one of the best for the end, Niraj Patel, Senior Vice President will close with a vertical you have not heard us talk about in the past, but I suspect we have a very compelling story now around automotive. He is going to talk about autonomy in vehicle remarketing and running in production. So that is my piece, we are excited. We hope you are. We look forward to covering all these six trends and with that, Preamble, Anup, over to you.

Anup Nair:

Thank you, Sudhir. Good morning, good afternoon, and good evening, everyone. Inba, can you please move to the next slide? So you just heard Sudhir talking about enterprise autonomy and how intelligence is becoming a commodity. Intelligence, as represented by large language models and small language models, is becoming a commodity. I am going to bring some examples to life proving the point, and I am also going to primarily focus on three messages. Inba, can you please move to the next slide? So, the first message I want to leave you with is that the inversion is already live, it is real and it is live, so what do I mean by that? What I mean is the traditional model where humans did all the work assisted by AI is reversed. We have seen time and again AI doing the heavy lift and humans validating judgments and outcomes that is what I mean by the inversion is already real and live. We know it because we are actually delivering it. Here are some examples. These are real examples of how we are working with our customers. The first one is vehicle inspection, so, we mentioned a little bit about it and Niraj is going to go into much deeper on this one. I am just going to give you a summary of this so that I can cover this background around it. So, vehicle inspection, this is for a large auto remarketer, so, this area is a very human intensive in nature, the reason why it is very human intensive is because it requires tremendous amount of knowledge and judgment, people with 15 to 20 years of experience they do the vehicle inspection and there is a significant amount of dependency on these people. So this customer approached us and the way we are solving this and we have kind of made the impact here, is by encoding the entire knowledge that the people, who are able to make those decisions into an intelligence layer and that intelligence layer is being used by AI agents to do all the work. This reduces the dependency on people, and you see the impact on the cost per transaction has gone down significantly. Another example, this is a very interesting example as well, this is about title operations. It is a very complex 41-step process operation. We are doing some of this work already. So, the customer who we are serving to do this came to us and said,

hey, we want to increase your workload, move from 75 people to 750 people. Now, in a traditional way, you would take a year to 18 months to really ramp up that people and train them and manage them and manage attrition. We took a slightly different approach. We took an agentic way to do this thing. We created an autonomous process where AI agents did the 41-step process, not only they executed it, they actually also did QC around it, but the best part is that humans handled only exceptions, judgment, and underwriting. This is a very complex activity that heavy load is lifted by AI agents.

The third one is an even more interesting example, mortgage tax operations. In this case, what we did was we built AI agents that actually made a call to the county office, asked for tax records, validated the information, documented the tax information, and humans just reviewed it this tripled the capacity for us. So, as I said again, the real heavy lift is now happening with AI agents. We are already doing it at multiple places. Therefore, inversion is real and it is life. I would like to go to the next message. Please go to the next slide, Inba.

This is the next message I want to leave you with. Economics do not just work, they flip. What do I mean by that? It means that things are getting cheaper. The real question in the room is, is deflation real? See, everything on the left-hand side is getting cheaper. Intelligence is getting cheaper, Sudhir talked about this. Software is getting cheaper. Knowledge is getting cheaper. Effort-based services are getting structurally challenged. We know this is happening, but people miss certain things. I am going to walk you through some examples showing what people miss.

Here is an example of a large global travel tech company where we are really reimagining the mainframe estate that they have. This is a massive program we have undertaken, we have reimagined already about 2 million lines of code and you see the impact that we are delivering. The interesting bit about this is not just we are just relearning and reengineering, we are actually accessing business logic from code and making that available for AI agents to run. This is very transformative way of modernization.

Here is another example. This is for a large brokerage firm where we have taken their entire SDLC and reimagined the SDLC with AI. Now, you would say that AI agents exist even before what is the big deal. It is not easy to have a multi-AI agent architecture where agents talk to each other in a seamless fashion. It is extremely complex to do that. That is possible only when you implement it in a certain design pattern. We are doing that for this customer. You see the results with 94% automation in testing that is significant.

I would like to take one more example from here, that is the global investment firm that you say. We have reimagined and brought an autonomy into the entire talent operations, agentic talent operations is what we call it. 110 AI agents

deployed and the accuracy is improving as we speak. The beauty of all these implementations is that the fifth project is cheaper than the first project that is an implementation that we do. So, for example, I am modernizing a Capability 1. The Capability 5 is going to be cheaper than the Capability 1 for the customer this is what people miss. The deflation on the left-hand side is actually funding this transformation. That is the second message I wanted to leave you with. Inba, could you please move to the next slide?

The third message is the system gets smarter and this is very interesting and you know that the capability gap between the first ranked LLM and the 10th ranked LLM has shrunk significantly and open weights and open source is actually challenging this as well. So everything on the left-hand side, again, models, agents, compute, code, all of that is commodity.

What is really different? What will differentiate is context. What is context? Context is how an enterprise operates. It is spread across multiple places, it is spread across your data, it is spread across people, it is spread across documents, it is decision patterns that have happened in the past, it is rules in the code, it is industry context, all of this together is context. The enterprise that brings it all together at speed will differentiate and not just that, this context, if designed properly, has the ability to make the enterprise smarter over time. That is key and it is a massive engineering lift. It is such a big engineering lift that people often underestimate it. Who builds this? Who does all of this stuff? Inba, please can you move to the next slide?

We built it. We have invested into this area. Some of this is public data. We have done significant amount of investments into AI platforms, into assets. We have done significant amount of investment into domain encoding. Domain encoding is fundamental for context generation. Data is one of the most important elements. We have done significant amount of investment in it. Sudhir mentioned about the enablement suite. We built Coforge Nuuron, which is the enterprise enablement suite. It brings all of these things together. Momentum blue, a group of top-class forward deployed engineers and this is key because, as I mentioned, it is a massive engineering lift. The right engineers are extremely important. Mod Squads for scaling. We have done all of this stuff. We have done a significant amount of investment in this area and this is working like a charm. Let us do a little bit of double click on this. Inba, could you move to the next slide, please?

Here is a picture of the whole vehicle, vehicle for driving enterprise autonomy. We call it Coforge Nuuron. It is meant for operationalizing enterprise autonomy. I want to point a few things here which are really key to the mix. On the left-hand side, it has an intelligence system, this is very key, without this the enterprise does not become smarter over time. As I said, the objective is to make the

Company, make the firm smarter over time. It has the ability to generate context, it has the ability to compound, that is a key aspect, the compounding aspect, where you start becoming smarter, smarter, and your marginal cost keeps going down, and it is composable. Any enterprise can take it and compose from the various capabilities catalog that we have to deliver that kind of intelligence, but intelligence is only as good as how you can make it real. We think that there is no point in having a great intelligent framework and not being able to activate it. Therefore, the second big element of Coforge Nuuron is this vehicle for activation. Momentum blue, which is a top-class Forward Deployed Engineer, Lalit Wadhwa is going to talk about that, and Mod Squads for scaling because proving value is one aspect but scaling it at a large level from an enterprise perspective is not an easy thing to do. Now, all of this is great. Without trust in the world of AI, it is of no use. There is no point in doing any of this if you do not build a trustworthy system. There is governance that is extremely important, and that means operations is one of the topmost things in this entire journey. Agentic operations are very key. Ashish, who is going to go just after me, is going to talk much more deeper into the agentic operations space. What is it, how important it is, is something that you will realize as we go through the journey today. Now, all of this is useless without data. It kind of becomes the foundation for the whole thing. Data is, without the right data, the quality of data, AI, and this whole intelligence system is of no use. The objective is to go to an enterprise and enable their data and make it ready for autonomy, that hence becomes the foundation for the whole thing. Deepak, who is going to talk about Data Cosmos and our data enablement for enterprise autonomy. As I mentioned before domain expertise, domain encoding, intelligence, activation, trust, and data, the big elements of Coforge Nuuron. Let me invite Ashish to start off with agentic operations as the next topic to talk about. Thank you.

Ashish Kumar:

Thanks, Anup. Good morning, good afternoon, and good evening to all of you. While the role of intelligence layer and the importance of data capability is very well understood, I would argue that degree of enterprise autonomy really hinges on agentic operations foundations. We will deep dive in the next eight to nine minutes. We will give a preview of what we do with 160-plus clients, how we are solving the agentic ops problem for them. Next slide, please.

When we look at our clients and prospects, we have been talking to about roughly 200 to 300 customers about their agentic journey. Two-thirds of them are at some stage of agentic operations. They have implemented AI tools. When you double click on that, most of the agentic capabilities on agentic ops space are really enriching the dashboards, the telemetry, and the enterprise AI hardware. Only about 21% of them have been able to build the context and intelligence to gain insight, and less than 5% can today orchestrate unsupervised actions autonomously, so why this big gap? When the investment

is real, intelligence has become commodity, and data importance has been understood the real gap in my mind, in our experience, has been the three fundamental problems that agentic ops need to overcome before it becomes a reality. One is a resiliency gap, second is security gap, and third is control gap, and we will talk about how we are solving this for clients. Let us go to the next slide, please.

So, if we think about resiliency gap that is the real hindrance in our world, but if we fast forward ourselves and see if we are operating in an agentic enterprise. In the agentic enterprise, we expect the change velocity to move from hundreds per month to thousands per day, and if you have a resiliency gap, that gets magnified. That will essentially break today's enterprise. In the post-Mythos era, the discovery of vulnerability has become a commodity. Enterprises are discovering not hundreds, not thousands, they are discovering millions of vulnerabilities in no time. What is even more important is how they are remediating those vulnerabilities. What used to be an expected norm in our industry to have a 45-days cycle for remediation has become a 48-hours cycle. Finally, with the usage of AI, the tokens have become a real challenge. Usage of what used to be millions per day has become billions per hour, and with that there have been questions about how do we control the entire enterprise AI hardness and that is the problem that we solve today for 160 plus clients. Roughly around 37 of our clients have been able to achieve 98.2% accuracy and when we talk about agentic operations, it is not the number of agents and it is not the number of AI workloads that is deployed, it is the accuracy that is the most significant measurement of how your agents are able to successfully deliver 100% of the time and today we support about 392 mission critical businesses for clients to establish the agentic operations. Let me go and give you a very specific example, if you move to the next slide.

Now, the resiliency problem. Resiliency problem essentially manifests itself in either the containment of an issue. If you have dependencies, whether on third-party dependencies, whether you have dependency on application that is deployed in the cloud and cloud sees an outage or you have a dependency on software supply chain that you might have, that is the context and that is the containment where we see the impact that we create maximum by having a solution called EvolveOps. It is a combination of about 28 agents and again, as I mentioned earlier, it is not the count of agent that matters, it is the accuracy with which they deliver the mission critical system and operations is what matters the most. So, we have been able to create 28 agents that manage the entire agentic operation lifecycle from sensing to detection to app and the results are very obvious for us to see.

For example, one of our Latin America Telco client, we have been able to de-risk, reduce their tech debt by migrating not hundreds, not thousands, 10,000

workloads to OpenShift and AWS Outpost in less than six months. That is the power of resiliency that we are able to create. For a global Ethiopian Bank, we have been able to move all of their infrastructures to have an automated way of managing and across the seven markets we are delivering that resiliency at a scale. As a result, this is something that I am really, really proud of. Some of our real clients who always had a Monday morning 7 am call schedule to review the weekend changes. As a result of our EvolveOps AI implementation, they have canceled their recurring call to review the post-implementation change reviews call. Many of our clients have been able to go to the games or go on with their vacation trips because we were autonomously managing one of the biggest changes that they have been implementing in their network and in their cloud infrastructure that is the power of resiliency that we are delivering today. We are not only able to contain in case if there is an outage, we are able to recover and restore much, much faster that is the problem that we are solving today for many of the clients. If I go to the next slide, please.

The second problem that we are solving, as I mentioned, in the post-Mythos era, discovery has become a commodity. Most of the LLMs are able to detect vulnerabilities and weaknesses in the software chain, in the IT systems very, very easily. What we have really created is a sensor mesh with Zscaler, which protects wherever your AI workloads are deployed, whether they are deployed on the edge or whether they are deployed on the cloud. We have created a solution that protects along all the vulnerability that can be visible to LLMs that is the number one problem that we have solved for clients.

Number two problem, most of our clients have different security platforms managing their security for workloads in the cloud, workload in data centers, and workload in third parties. They have different ecosystems created for managing security for AI agents, for human and non-human identities. We have been able to unify all of that to provide them the intelligence. Many clients have, in the past, been using time-to-time audit and assessment to detect what the weaknesses exist in the environment. We have been able to move them away from point-in-time sample-based testing to a continuous detection and continuous decision engine. As a result, we have been able to help clients mitigate in minutes not in days. The philosophy with which we operate for each of these clients is that we have implemented zero trust for clients, we have zero delay when it comes to getting intelligence and zero delay when it comes to mitigation. That is the power of capabilities that we bring to bring the remediation at a machine speed. Inba, go to the next slide, please.

We are obviously into the metered era of AI. Every client that we know have burned their annual budget of tokens in months if not days and that has obviously resulted into every single client looking at having some kind of ceiling, some kind of metering implemented in the environment. What we believe in this

space is just having a meter does not really solve the problem we have to go beyond and start looking at sovereign AI as a capability that drives the nonlinearity, that bends the curve, allows you to have a mixture of models and finally it allows an enterprise to scale nonlinearly when they adopt AI at a scale.

If I go to the next slide, I just want to sum it up by saying that we are seeing a tremendous tailwind. Working with 160 clients, we have seen 37 clients have experience and solved one of the three problems. The true power of AI and enterprise autonomy will really be realized when clients are able to solve the problem of resiliency at scale, remediation at speed, and control at source. We have quite a good proof points of the first two, and I believe the third point: the sovereign AI, the adoption of AI, launchpad that we have for Open-weights, OpenShift, and TokenOps are the three solutions that we are creating to drive the control at source, and that is where I think most of the growth will happen in 2027 for us. With that, I will turn our attention to Data Enablement for Autonomous Enterprise. Let us welcome Deepak Manjarekar, Coforge Head of Data and AI. Deepak, over to you!

**Deepak
Manjarekar:**

Thank you, Ashish. So good morning, good afternoon, and good evening, ladies and gentlemen. Ashish just talked about agentic ops and how his team and his unit run the agents, keep them safe, and also run them under the budget. You heard Anup talking about Nuuron. Nuuron is our enterprise autonomy enablement suite, and he talked about models, agents, compute, code, and everything getting commoditized. The only thing he talked about is context and context is becoming key. You heard Sudhir talking about context as well. Every organization gets the access to Mythos, access to Astra, access to DeepSeek, whatever model are out there, so the intelligence is commodity, it is accessible. What makes it unique is the context that Sudhir was talking about, their data, their processes, and their secret sauce.

In the next few minutes, ladies and gentlemen, I will talk about three things about building this context. I will talk about the challenges that customers have in building this context and why it does not get built properly. I will talk about how we at within the Data HBU use our Data Cosmos innovation platform and deploy some of our assets under the Nuuron an enablement suite, to build that context and how that creates tremendous value for our clients as well as for us, Coforge and our investors. Can you go to the next slide, please?

So here is the autonomy paradox, and I am using a metaphor of an iceberg. Everything above the line and there are a few things above the line, which are very shiny. Everybody, the execs, the boardrooms, are spending on models and building agentic systems, but they are not looking at the critical stuff that makes it work. That is data, and everybody talked about it before me I am just going to make it real for you. Now you see a few numbers here. Again, these numbers,

these reports, are created by some of you guys, your colleagues, so I do not have to substantiate them. I am just using them here to prove my point. While above the waterline everybody is spending money, 64 billion and counting, below the waterline there are problems.

There are projects that are getting cancelled, many AI projects remain in the AI purgatory, there are projects that see the light at the end of the tunnel, do not deliver the ROI and then Gartner is already predicting that 40% of agentic AI projects probably will get canceled by 2027. Primarily, it is because of the stuff that stays below the waterline. People do not do anything about it because they do not make headlines, they do not give you promotions and two reasons one is they do not get any promotions nobody gets promoted saying I fixed quality for last year, data quality. Second, this stuff, the data, stays underneath the agentic system underneath the UI/UX, so it is not seen and for autonomy is as good as the data and nobody is fixing that data, so autonomy is an elusive finish line. Can you go to the next slide?

So how we are actually solving, you saw Anup talking about Nuuron and the entire context layer. We are deploying the five different assets within the Data Cosmos platform and we have about 130 plus assets within the platform, 68 agents and we have deployed these things about 50 plus clients. We use tools like DataFlux, Agentic, Scriptless, Dynamic, Data Ingestion Platform. Agents actually manage that platform. Agentic DQ Resolver, data quality management, DG-Nexus is a governance framework, Data4AI, AI for data is the data activation layer and all of these tools are fungible. If somebody has Collibra or somebody has Atlan, we do not have to take the DG-Nexus we simply work with that and even if somebody does not have, let us say Nuuron, and they do not want to do Nuuron because they already started on the journey, all of this will be plugged into their Nuuron like enterprise autonomy structure, but when a customer takes Nuuron this whole thing is extremely seamless. That is how we build the context. That is how we bring the data from the system that you see on the left side and put it right into the enterprise autonomy layer, whether it is Nuuron or otherwise and that drives the autonomy, the intelligence layer, and the decision models. Can you go to the next slide, please?

So how does that help us and people would say so what? So, it actually creates disproportionate value for our clients and for us. Three ways: one is it gives the right start and once you get the data layers straightened out, everything just works like a charm, as Anup said and we focus on only a few industries where we go miles deep. We understand the ontology, we understand the data model, we understand the relationships of the data, and we just build that through our own secret sauce. The second mode that we have is our own IP on Data Cosmos and as part of Nuuron. It creates the IP arbitrage for us. It shifts everything to the left, speed up the entire execution, brings predictability in delivery and reduces the

risk and then the last one is essentially with the Data Cosmos flywheel, we compound the value over time and Anup talked about it. Every second, third, fourth engagement that you would do will be marginally lower cost and it just keeps sliding down and that is the value that we will give and with this, we basically create disproportionate value not only for our clients but also for our investors and shareholders and that was my pitch, ladies and gentlemen. I would invite Lalit to talk about Momentum blue and everything great stuff that he is doing in that area. Lalit, over to you!

Lalit Wadhwa:

Deepak, thank you very much. Hello everyone. There are two key pieces that over the next few minutes I am going to go ahead and discuss and hopefully it would provide everybody with a view into how some real complex problems are going to be solved.

So if we move to the next slide please, the number one thing that at least I will try to go ahead and communicate with you is this paradox that probably each one of us has read enough number of times and the paradox is every single firm today wants to at least have a roadmap towards autonomous enterprise, being an autonomous enterprise and then you look at the data, which is on the left hand side of the slide, which says fundamentally that the success rate in deriving ROI out of everything that most enterprises are doing at this point in time, that success rate is extremely low.

How do you solve for this paradox is the number one thing that I am going to go ahead and share with you. What is the approach that we have taken? What are the kinds of results we are seeing? And why are we confident of the path that we have gone ahead and taken at this point in time? I think the second important thing that at least I want to go ahead and share with you is building a solution frankly is not the toughest thing. How do you go ahead and scale it over a period of time to go ahead and meet the emerging need and to meet the emerging demand, I think is going to be very critical and hopefully, by the time I go ahead and at least in my session you will have a view into it.

Let us dive into the first piece of it, which is enterprises wish to go ahead and build up an intelligent fabric. They want to have a data layer, which is smart. They want to have an enterprise memory, where every single decision of the enterprise that has been taken in the past is stored. Agents can use that decision to go ahead and plan their actions. They can go ahead and take some actions based on that planning and then they can actually write back into that data layer. We also heard Ashish talk about the agentic ops. If autonomy, and autonomy being defined as the act of agents taking on more and more of the human labor, that has to see the light of the day for a vast majority of our clients. How do you manage those agents? How do you operate those agents? How do

you capture that telemetry? All of that is going to be important. The solution is right in the middle and most of us have heard of the forward deployed engineering piece over the past few years.

Momentum blue is that operating unit within Coforge where all our FDE capabilities sit. Now, it is very important just to go ahead and spend 30 seconds on why a forward deployed engineer is not a software engineer because I see this getting confused again and again. I do hear lots of companies making claims about 10,000, frankly some companies making claims about 100,000 forward deployed engineers and I just want to go ahead and provide the audience with a view to what we believe a forward deployed engineer's job is. He or she has to be an exceptional engineer, number one, has to be a change agent, has to be an expert in discovery, has to be an expert in guiding the customer towards adoption, real change management within the client, and has to be AI native by mindset.

What does AI native mean? AI native, the definition of AI native is right in the middle box. An FDE does not see an AI agent as another tool, it sees the AI agent as a first-class citizen, which is going to assume certain responsibilities, which in the past, would actually have been taken on by certain humans and now AI agents are taking on those responsibilities. So first-class citizen of the pod and not a tool to be used, that is what AI native means.

These pods, the FDE pods within Momentum blue are designed with some very specific principles that we have crafted. Number one, the ownership of the outcome resides with the senior FDE. This is not time and material. Everything is priced by outcome, it is not priced by time and material, it is not priced by the number of hours. So, the risk for us, and of course, the risk for the client is proportionately lower. Risk actually shifts a little bit to the delivery side, which is onto the Coforge side. Most importantly, and I think this is poorly understood generally, there is a designed exit plan.

The plan in the past used to be, if you get into a client's project, you go ahead and continue to manage the project for a very long time. In a forward deployed engineering pod, you actually have a methodology to go ahead, assume the role with humans taking on a lot of work, agents taking on a little bit of work, but very quickly flip the shift as was defined by Anup early on in this entire conversation and so the economics start working. Every next capability costs much lower, is delivered much faster for the clients if we move along.

On this slide, I am going to go ahead and talk about what does it look like when you look at the workflow that an FDE team goes ahead and performs. Remember, I said, this is an outcome-based engagement or an outcome-based engagement model. You can only design an outcome-based engagement model if you know what the current status. So, number one, and you probably would

find it surprising, but in many outcome-based engagement in fact, at least a majority of them the number one task that the FDE pod or the Momentum blue pod has to do is figure out the baseline. You figure out the baseline, you go ahead and embed the FDE pod and the agents into the client, you go ahead and start improving the outcomes, you reach the end state, and that is when you go ahead and plan out your exit, you transfer the capability to the agents that you have put into place, and you handle with those capabilities to the client, you move on to the next set of problems that have to be tackled over a period of time.

There are three small examples that at least I want to go ahead and share with you. These are active either in production or just to go into the production over the next few weeks. First of all, on the left-hand side, a global consumer goods company. Look at what the FDE pod over there is building out, an entire knowledge architecture for that enterprise. Every single decision, every single e-mail, every single drawing, every single video chat that is there in that enterprise over the last X number of years is being put into the enterprise memory, on top of which agents will be able to reason, agents will be able to learn, agents will be able to plan their action, agents will be able to go ahead and deliver the next set of actions and this is at stage number three and three basically is moving towards production with human in the loop.

Look at the next example, which is in the middle box. This is an HR technology platform Company, which provides services to millions of small business owners. They completely want to reinvent how they design software and this is not building along with the agents this is actually agentic software development with a majority of the software development actually including a majority of software testing is going to be done by agents and not by humans. Humans, of course, are going to be there to go ahead and govern and make sure that the agents design, produce software, and test software in line with what the humans are looking for.

Look at the right-hand side box which is about a significant asset management and a wealth management company designing their entire AI platform, which is going to serve their entire enterprise and what does an AI platform mean? It means you have a governance layer, every single call to the external world, the APIs, where the models are/is governed, is compliant. There is telemetry that is emerging out of it, hundreds of models, and of course, a model router, which goes ahead and figures out which model is the best to go ahead and respond to the query, complete governance on top of it, an enterprise fabric, intelligence fabric, where every data point is stored, hundreds of agents on top of it, some of them for software engineering, some of them for their horizontal workflows frankly, and some of them running the workflows with their end-clients on the outside. These are the kinds of things the FDE pods at Coforge at this point in

time are building.

Let us move to the next slide and here is the second message that I want to go ahead and share with the audience today. This is about how do you go ahead and scale this? Primarily, there are two scaling mechanisms that we have gone ahead and invested in. This is what we think is going to go ahead and create the fastest increase in the overall scaling that we require at this point in time. The number one thing is we are going to build our own FDE talent engine. This is something that has been in place over the last five months at this point in time. We have our own academy, this academy has a 12-week program, it takes some of the brightest, the exceptional AI native engineers and transforms them into FDEs. So, remember what I shared with you, many layers have to be created over an exceptional software engineer before that software engineer becomes a forward deployed engineer. All of that is done over here, right here in our Princeton Headquarters, the US headquarters and these FDEs then go ahead and shadow with some of our senior FDEs. They understand how the customer sees the world, they understand how to go ahead and respond to the customer's view of the world and every 45 days, every six weeks, we go ahead and put together a brand-new cohort. A cohort generally is about 20 to 25 people for us, and this is how we are scaling very rapidly.

I think the second important thing to recognize the scaling methodology that we have gone ahead and used is the flywheel effect that we keep talking about. So, if you look in the middle box, you go ahead and deploy an FDE pod with a client during work at that particular client. There are things that you learn. We could have improved the assets that we have, and we could have improved the processes that we have. There are new learnings that we have out of that deployment. How do you capture it and we capture it using agents. How do you capture it? How do you provide that feedback into the team, which is building out our assets, as a result of which the fine-tune existing assets or build new assets if required and so the next implementation that this pod does or any other pod does is superior, cheaper, faster, simultaneously, as far as the client is concerned. That is what leads to the right-hand side, which is you begin with human labor and a small piece of agentic labor and over a few weeks you completely flip the script.

With agents taking on a lot more work and humans decreasing their work at that point in time, the FDE pods hand over the work back to the client and move on to the next problem that has to be solved. So hopefully over the last few minutes, both the questions have been addressed. Now to see how this actually works in practice with our clients' two case studies. I would like to go ahead and invite two of our senior leaders, Erika Moore and Niraj Patel, to provide the audience with a view how is this getting implemented at our clients. With that, over to you Erika!

Erika Moore:

Thank you, Lalit. Alright. Well, hello, everyone. In the interest of time, I am going to really zone in on what we are doing in the travel industry and hopefully bring everything that my colleagues have shared into life, so if we can go on to the next slide, please.

So, we are going to focus on travel and specifically we are going to go into the airline ecosystem. Now, I am really excited to share this with you because I believe that travel, unlike many other industries, is actually very emotional. It touches every single aspect of our lives and once you leave your home, you are completely at the mercy of the Gods of travel. So, we are going to zone in on one of the largest, most significant transformation points that the industry has ever reached and it begins with the airline ecosystem and I call it the airline ecosystem because while we focus on the airlines it will create a ripple effect onto airports, cargo, travel management companies, and all sorts of experiences that will be shifted for the traveler. Now, throughout the course of the next three couple of slides, we are going to focus on what is this transformation, what does AI actually enable, and what will be the implications of AI in this transformation, how it all comes together, and why Coforge has a unique advantage and is taking the lead in making this transformation real for the airline ecosystem because we are actually redefining the travel experience. So on to the next slide, please.

Alright. So let us talk about this. What is this transformation? This transformation is called modern airline retailing, and it is an inflection point that is actually being driven by the end traveler, by Gen Z, by Gen Alpha, who will forever consume and experience the world in a very different way than any of us had ever experienced it before. What is enabling the transformation to actually be delivered at a space that had never been experimented for the travel industry before is AI.

Traditionally, transformations for the travel industry take between 20 to 30 years to complete. In the case of modern airline retailing, we are seeing this get accelerated and Coforge is at the forefront of being able to enable this. So, what is happening? What we are seeing is that with AI, technology is no longer the issue. Technology can deliver all of this transformation. What is required, that is the real challenge, which is the change in the mindset, the scalability, and the adoption of this transformation at an enterprise-wide scalability. So, what is happening? We are seeing AI pilots deployed everywhere within the airline, within airports, within cargo management companies, and within travel ecosystems. I am sure all of you have heard here and there experimentation taking place of agents to agents and agents managing travel and agents managing the experience, but there has not been yet a complete solution which orchestrates all of it and brings it together. Why? Because not only is the airline

industry regulated, it is also highly complex and the isolated ecosystems and fragmented data pods where things exist make it extremely complicated. Further to that, you have this challenge a 50-year-old technology that needs to be ripped apart and you have this in places where it is not documented anymore and the people that actually had the knowledge to stitch it all together sometimes are no longer in existence within that particular enterprise.

So, what is happening and what is required, what is the must for this to come together? It is required that we rewire the entire ecosystem. The keyboard is going away. It will be agents to agents only, but the challenge is in the orchestration, that orchestration that will allow autonomous agents to continuously interact and make real time decisions that drive those personalized experiences that are now being expected by the end traveler. That is the challenge, and that is where Coforge is deploying its entire 30-year knowledge of everything that we have ever done for the travel ecosystem, bringing it together with everything that you have seen to enable it. So if we go on to the next slide.

What we are seeing, discovery is going to be fundamentally different. You will no longer call your travel agent. You will no longer take your phone and start zooming through and tabbing through different things and looking at different options and looking at different web pages that try to compare things for you. You will have your own personalized agent and what you will expect is that your own agent transacts with the agents of all of the suppliers in the ecosystems, bringing those collated experiences that match your needs. Then your journey, the expectation will be that the entire ecosystem tailors and adapts and adjusts to your journey and any disruption that occurs within that journey. That is easier said than done, but that is absolutely the future of travel. It is what is being mentioned as the enablement of AI to actually touch that deep emotion that comes to life when you are traveling, that generates anxiety when your trip is disrupted, that generates that ability to dream and to imagine and to become inspired around your next holiday and to actually be able to contextualize whether you are going for business, for pleasure, traveling alone or with your family or meeting up with groups of friends. All of that must be orchestrated, but in order to do that, what I mentioned in the slide before is that entire technology that can actually come together to drive this is not there yet.

So, what does this mean for Coforge? It will be our role to help shape that IT transformation, to help drive the strategy of how an airline, an airport actually come together to help drive that transformation, the execution. This is critical and as all of you know, this is what we are most known for, our ability to execute, but now it is our ability to execute and to orchestrate, to orchestrate these siloed ecosystems, these siloed databases, and to bring it all together and help drive that transformation and that inspiration for the next era of travel. How will we do

this? If we go on to the next slide.

This is just a graphic, very simplified version of how we are enabling it and the secret to this driving and having the unique advantage in the modern airline transformation for the industry is Aeronova. If you want to think of it in a very simplified way, what Aeronova does is it comes together driving what Anup shared about Nuuron and what Lalit shared about Momentum blue, whereby Aeronova is how we deploy and how we help drive that transformation for the airline ecosystem. It is where those agents help orchestrate and change the entire operational system for an airline, but more importantly, what is actually really changing within the airline. It is the way it envisions its own business processes. It is the way it will forever operate and interact with travelers and with the other aspects of the traveler journey. Momentum blue is the way that we execute Aeronova. So, this is how it all comes together for travel.

Now, if we go on to the next slide, we will focus on Aeronova, which is exactly the platform that is helping drive this transformation for the airline ecosystem. Now, what happens? And I mentioned that IT is no longer the main obstacle or the challenge. The real challenge, and I always say this when I have the opportunity to speak to customers, IT is no longer the challenge because if you have the right tools, if you have the right platforms, and if you partner with the right company, we will take care of that for you. The real challenge are actually humans. This change in an airline's operation is deeply emotional. Deeply, it will forever transform the roles, the interaction, and how that airline operates. If you can imagine the conversations at the board level and at the C-suite and being able to embrace this is a mixture of we must with a great deal of fear of not understanding exactly the extent in which everything will get disrupted.

So, what do we do by enabling Aeronova within the airline's ecosystem? We de-risk that massive transformation. We accelerate it. We help drive revenue faster into what this means for the airline. Now, it is expected that with modern airline retailing, there will be an uptick of roughly \$7 per passenger for the airline.

Can you imagine what this actually means for an airline that is an early adopter and can start reaping the benefits of this transformation sooner rather than later? It is absolutely huge. The challenge is Airlines are not sure on how to navigate this. They need visibility into understanding exactly what business processes are going to get affected. What are the implications? How will they need to transform their organization? And then all the streamlined data that goes on to airports and to regulatory agencies to be able to drive this. This is what Aeronova supports and helps them navigate.

So what you see there are the different pillars that are enabled by the adoption of the platform. It is Modern Airline Retailing Business Suite, which actually talks about all of the business processes that get affected while driving this

transformation. NDC, which is the “New Distribution Capability”, actually starts to address how the content for the airline will have to shift in the way it gets distributed. How does it actually start to navigate the journey into being discoverable by agents, Offer Order Test Suite. All that is, is that while this transformation is taking place, you need to have every single point in time for an airline, be able to know where the old system and the new system of doing things are coexisting so that nothing falls through the cracks.

The Knowledge Academy is absolutely new and we come together and identify where the airline will need to be retrained or untrained in some aspects of the way that they did things before and how they need to move forward to adopt a new way of doing things and last but not least, the most important piece is that trust that my colleagues have talked about, that trust and that governance to make sure that while AI and the orchestration of the agents that come together to drive this enterprise transformation is taking place, it is always in the context of that rule and the differentiator for the airline.

One of the biggest challenges in this transformation, and as we have talked about, is that fear that with AI, because knowledge does become commoditized in some way, that airlines will not have the ability to be able to differentiate themselves from others, that they will lose contact of the actual customer because it will be all agents to agents. This is a great fear. And what we are working on with the travel ecosystem is to ensure that while they drive this transformation, and yes, there will be many aspects that become commoditized, they are able to continue to drive that differentiator aspect that makes one airline different from another and makes them identify with the end traveler at any point in time. We through Aeronova and through the suite of solutions and products and deployment execution arm of Momentum blue are early on in this stage. Just to give you a little bit of data points, this transformation journey has not been begun by over 80% of the travel ecosystem. It is an early stage, and it is a transformation that we believe is going to take approximately 10 years to take on. So by putting all of this together and being here at this point in time, we are actually helping drive the definition of how this all will come together and that definitely places us at an advantage. With that, I will pass it on to Niraj to speak a little bit more about auto industry and what we are doing there.

Niraj Patel:

Thank you, Erika. You just heard how AI is changing the travel industry. I want to show you something different, not a future vision, not a pilot, and not a proof of concept. This is a production system already operating at scale in automotive.

If you want to understand enterprise autonomy, do not start with AI. Start with a business problem. Imagine when a car or a vehicle is returned today to a manufacturer or to a finance company. It could be a lease, it could be a rental,

and many different ways these cars come back. At that moment, thousands of dollars of value is trapped. Trapped in uncertainty. Can it be certified? Should it be repaired? Should it go to auction? All these decisions have to be there, but every day, nobody trusts the condition, the asset is worth less money. The traditional response to this business problem would have been, improve the inspection. Take more pictures. Document more damage. Produce better reports, but the business does not create value from the inspection itself. The business creates value from the decision that follows. The inspection is not the product. The decision is the product.

The clients need faster, more consistent decisions and what is required is not another dashboard and not another workflow. It requires trusted intelligence in that asset. The trusted intelligence asset is what we build for the manufacturer. Next slide, please.

Most inspections begin when the vehicle arrives. We asked a totally different question. Why wait till it arrives? For the vehicle, how much information exists? Before anyone touches the vehicle today, information already exists in organizations, in multiple systems, diverse systems, and siloed systems. The vehicle identification number, mileage, maintenance history, or in warranty systems, they are all over the place. So now, before the inspection, AI agents, the system assembles the context. It puts it all together and gets it all ready for the inspector. So, the inspector is not wasting time, what I call on screen, putting together information for the inspection. They are in the field. The inspector is inspecting the vehicle with a guided video of the car versus pictures of the car versus different ways to do this. And all this, is what the inspector gets to do to inspect the vehicle, not search for history. History is already there. How do we assemble that and make it work? So, what does this mean? The inspector should inspect the vehicle and not the history. Everything is ready. This live production example year to date, has processed over 66,000 vehicles as they come through the process. The deeper value is the years of institutional judgment that became repeatable. The process is consistent. Evidence is auditable, and each interaction creates a better data set for future decisions. Next slide.

This is the slide I want you to remember. Most people see condition record as an output. It is not. Condition record is the input to every commercial decision that follows. The inspection happens once, the decision happens many times. Repairer asks, what should economically be done? Do I move it to a different channel? Do I send it to auction? Do I have a dealer ready to go? How will these patterns help the entire ecosystem as Erika described in the travel world. It is not just about where the information comes, the decision happens. The key piece is once that decision is trusted, the information is trusted, it has impacts throughout the entire ecosystem of that vehicle.

Again, I want to make sure the patterns of this travel well beyond automotive. The same principles happen in mortgage processing as you heard Anup talk about. Claims, travel operations, and financial services. Key item here is inspect once, decide many times. If you have one takeaway from this is the value is not in task automation. The value is in encoding institutional knowledge, so every future decision can become faster, more consistent, and more scalable. You heard Anup talk about how much value this brought to the table from a transaction point of view.

I am going to give you one more layer as that compounding effect happens. As this trusted data is available, let us say a dealer was going to buy that vehicle. Today, the process is they do not trust the vehicle, and when the vehicle arrives on the lot, they take more inspection and more pictures, but because now it is a trusted asset, that information is electronically identified and sent to the dealer even before the car shows up at the dealer's lot, and they have access to that information, saving time, value, and making that asset available for sale immediately instead of having to wait for physical arrival. I hope that helps explain how autonomy and how we think about autonomy in a live use case. I am sure a lot of you have cars and you have not seen the back end to this process, but as Erika says, a lot of these industries are highly complex, highly fragmented, but you cannot have dollars tied up in a vehicle because trust is not there and this process, this autonomous process is bringing trust to the industry. With that, I will hand it back to Lalit.

Lalit Wadhwa:

Erika and Niraj, thank you very much. So hopefully this was good use of the audience's time. Of course, we are going to go ahead and do a few questions. But here are some of the key messages that, at least from our end, we would love for the audience to go ahead and take from over the last 70 to 75 minutes or so.

I think message number one is that the concept of enterprise autonomy is real. It is something that some of the largest firms and some of our largest clients are pursuing, and they are pursuing it with their dollars, they are pursuing it with their investments, they are pursuing it with change management, they are pursuing it with technology shifts within their own organizations, and they are changing not only themselves, but they are actually changing the industries and the ecosystems that they participate in.

The second message is the inversion is real. It is live and what does inversion mean? We defined inversion as the ability to go ahead and incrementally hand over more work to agents, smart agents, intelligent agents who work with real data and consume the human labor to a lower level increasingly and incrementally. We discussed the concept of context. Every single company does stuff differently. Their workflows are different, even within the same industry.

Their value chains are run differently. They come to outcomes in a very different manner. All of that is context. Capturing that context, building out the enterprise memory using that context, building out the data layer, these are critical tasks for any enterprise. That is where a lot of investment is going to go and AI-ready data is going to be one of the primary investment areas for many of our clients and we are beginning to see this day in and day out with some of our largest clients at this point in time. When you go ahead and hand over operations to agents, those agents has to be monitored, even when they are autonomous, you have to monitor them. You have to monitor them for speed, for accuracy, for the kind of decisions that are being taken, which basically means telemetry is important. Without telemetry, you cannot go ahead and attain autonomy. You cannot go ahead and attain the secure environment within which agents should go ahead and deliver the results and of course, it becomes very difficult without those two to go ahead and make sure that the decisions are being taken in the right manner and you are not ending up with incorrect decisions.

Now, to go ahead and operationalize all of this, we discussed the concept of the forward deployed engineers. We discussed why a forward deployed engineer is absolutely not a software engineer. We discussed the time and material model is being dismantled, frankly, far faster than anybody in the industry thought that it would be and ownership of the outcomes is becoming the primary driver and in an FDE pod, the ownership for the outcome lies with the senior FDE. That is what they are hired for and then finally, using the two case studies and use cases that Erika and Niraj shared with you, at least we demonstrated to you how industry by industry challenges exist, but those challenges are being dismantled incrementally again. Domains are going to play an important role in attaining the autonomous enterprise state, because every single domain has a very different meaning and a different path to getting to that autonomous state. With that, I think from our perspective, we close the session, but I believe we do have a little bit of time to go ahead and take on questions. So over to you, Inba.

Moderator:

Thank you. Ladies and gentlemen, we will now begin the question and answer session. As a reminder, to ask a question, please click the “ask a question” button below the media player. Enter your details and wait for your turn to be announced. After your questions are answered, you may click the “leave” button to return to the webcast page. If you have a follow-up question, you may submit the same way or use the “raise hand” button if you remain on the panelist platform. We also request our participants to ask questions related to the theme and topic of the meeting discussed today. If the question is not directed to a specific speaker, Paresh or Radhika will direct it to the right panelist to answer. With that, I will call upon our first participant that is Lim Jake Warren of Aberdeen. Could you please unmute your microphone and go ahead with your

question, Lim? We have sent you a prompt to unmute your microphone. You may go ahead, please.

Lim Jake Warren:

Sorry, I think there is a mistake. I did not raise my hand.

Moderator:

We will now unmute Prashant Kothari from Big Tech. Prashant, you may go ahead and unmute your microphone, and you may ask your question.

Prashant Kothari:

This is Prashant from Big Tech. I just wanted to know a bit more about the FDE setup that you have got, how many FDEs you have got now, and what is the kind of cost per FDE, revenue per FDE kind of model versus the typical software engineering model that we are more familiar with?

Lalit Wadhwa:

I will take that, Prashant. We have a stated goal of 150 FDEs by end of March. We are well on the way. We are currently somewhere in the middle of that road, somewhere close to about 70 to 75. To the second point of the question that you had, which is, what does the revenue model look like? The revenue model is very different from the revenue model for a software engineer. We discussed the fact that the skillset of an FDE is very different. That is the reason the skillset is in such short supply and that is the reason the revenue model or the billing rates that are attached to an FDE are completely different. Now, of course, that also means that the cost model of the FDE is going to be very different. You cannot expect an FDE with all of these traits to cost the same as what a software engineer would, so very different cost model, very different revenue model.

At the same time, Prashant, I absolutely want to be very clear. When we hear announcements of 10,000 FDEs and 20,000 FDEs, we look around and try to figure out where did this announcement come from? You look at Gartner, you look at several other analysts in the US market, everybody tells you the total number of FDEs, current FDEs in the US market is somewhere between 5,000 to 7,000. That is the reason you see such a vast disparity between supply and demand for FDEs and that is the reason we truly believe that the 10,000 and 100,000 numbers that are being branded about, these are software engineers and not the real FDEs, at least the way we define the FDEs.

Moderator:

Prashant, do you have any more questions?

Prashant

That is it. Thank you very much.

Kothari:

Moderator:

Thank you. We'll take our next question from Sarthak Batra of Aditya Birla Sun Life. Please go ahead. Sarthak, you may go ahead with your question. Well, there is no response from this connection. In the meanwhile, we will take our next question from Satish Kumar S of Kotak. Satish, you may go ahead, please.

Satish Kumar S:

I have a question on IP protection. So, let us say if an agent runs in a client environment and it continuously adapts and learns based on, let us say, by doing tasks, then how is the new IP generated? How is it partitioned between, let us say, the client and Coforge? I mean, clients would probably want to keep the IP to themselves, whereas Coforge would require the IP to continuously keep on improving its Nuuron platform?

Anup Nair:

Let me answer that question, Satish. Data is customer IP. Agent assets that build out the agents is our IP. We are very clear about that. Any data that is needed to run the agent is customer IP. So that demarcation is handled as part of our contractual discussions, and we are very clear on what remains ours and what remains the customers. I do not know if I answered that question, Satish.

Satish Kumar S:

Got that, Anup. That is helpful. Thank you.

Moderator:

Thank you. We have our next question from Abhishek Mundra from HFS Research. Please go ahead. Mr. Mundra, your connection is unmuted. Please go ahead. There is no response from this connection. We will now unmute Tushar Dhonde of ITOM. Please go ahead, Tushar. Could you please unmute your mic and go ahead?

Tushar Dhonde:

This is Tushar from Shanghvi Family office. So my question was directed to Ashish or maybe Lalit even can take it. So, we address that there exists a resiliency gap and then there is a vulnerability gap in the enterprise environment and AI is helping to address that? Now, there is again a different school of thought. So, the recent comments by the Anthropic's CEO saying there is higher chances of AI agents may be getting rogue, so how do we prevent that in our Momentum blue model where we have a senior FDE, associate FDE and AI agents executing the actual work? What are the additional steps that we might need to take so that these AI agents do not get rogue and it does not hamper the

client environment?

Ashish Kumar:

Let me answer that and Lalit you can add that right. I think when we talk about trust and control with agentic operations right, we talked about essentially not just having mitigation strategy or detection strategy. The first order of the day is to have a secure S2 Cloud environment, so that you have a total control on environment in which your agents are running. So protection strategy is the first order of the day and that is where the zero trust, zero dwell, and zero delay strategy is what we have gone ahead with and that allows us not essentially to react before an agent goes rogue. It is actually protecting before agent goes rogue, just like you would do in human identity as well. And Lalit, on the FDE side, if you can please add.

Lalit Wadhwa:

Yeah, Tushar. Great question. I must also admit that there are no perfect answers to the question that you asked and that is the reason I think that this is a great question. This is something that every single company, every single industry is pondering with and Ashish actually covered this during the course of his conversation, where he said, I can identify vulnerabilities very quickly. I can produce a patch for that very, very quickly, but should I deploy it or not still requires some degree of human judgment, right. So, the human judgment piece is still required. The way you go ahead and build out agents in a client's environment is cover them with guardrails, ensure that compliance, the regulatory framework of the industry, regulatory framework and the policy framework of the company, and the right alignment with the actual goals that the client wants to achieve, that is taken care of. Now, does that guarantee 100% that no agent will ever go rogue? It does not.

Humans do not go ahead and give that guarantee when they go ahead and work with companies. Similarly, everything is done to go ahead and protect the agent from going rogue. Guardrails, compliance, policies, the structure in which they operate, air gapping in certain cases with inability to go ahead and reach the outside world or limit the connection to very specific components and systems within a client's environment. Those are the things that you would do with humans. Those are the things that you go ahead and do with agents, but it is a subject that absolutely is of great interest, great discussion at this point in time. So hopefully between Ashish and I, Tushar we have addressed your question.

Tushar Dhonde:

Yes, absolutely. Thanks. Thank you for addressing. All the best.

- Moderator:** Thank you. We will take our next question from Apurva Prasad from Franklin Templeton. Please go ahead.
- Apurva Prasad:** Hi, could you give some metrics around at organization level, what would be the token consumption? How would that be trending? Or perhaps maybe agent work units or any metrics?
- Sudhir Singh:** Lalit, I am not sure if we have those metrics, but do you want to take a stab at it?
- Lalit Wadhwa:** Absolutely, absolutely. Apurva, great question. But I also want to go ahead and shift the conversation a little bit. Measuring total token cost, according to us, is an incorrect metric. The right metric is the cost per unit task that is leveraged. So, if to go ahead and operate a part of the value chain or a workflow, I am today expending \$1 human labor or human and automation labor, and I am able to bring it down to \$0.50, it does not matter whether I am using anthropic tokens or I am using tokens from a non-frontier lab, right. So, the right measure is, are we able to go ahead and measure the cost per unit of work that is being delivered.
- I can also admit to you, most companies do not today at this point in time understand the unit cost, per workflow right. I admit that, but that is exactly how FDEs go in. They go in, they measure what the current cost is and then they go ahead and put together, along with the client, they co-ideate with the client what the future looking cost has to be. And then they go ahead and chart that path between the current state and the future state. So, I think trying to measure raw token cost, I do not believe it is the right measure. It is the cost per unit of work that you have to go ahead and deliver in any value chain any workflow that is the right metric and that is something, which frankly is not too difficult to go ahead and control at this point in time, given the variety of models, the capabilities of the models, the onset of the open-weight models, and the abilities and the capabilities of the onset of the open-weight models, and the rapidly declining difference between a frontier model and an open-weights model for 90% of the work that any enterprise does today. So, hopefully, Apurva, your question has been addressed.
- Apurva Prasad:** Yes, Lalit, I was not asking the token cost. It was more on the consumption units, which is what I was asking.
- Lalit Wadhwa:** Right. But does my answer fit that question too around consumption units, right?

Apurva Prasad: Yeah.

Moderator: Thank you. We will move to our next participant. That is Srini Vadapalli of HFS Research. Please go ahead.

Srini Vadapalli: Thanks for the insightful session. I just wanted to understand from your experience what kind of conversation which you are having with the enterprises with respect to GCCs? So, are they talking about agentification, modernization, or automation? So just wanted to understand from your conversations with the clients or enterprises at the headquarters.

Sudhir Singh: Lalit, Anup feel free to jump in. Conversations with GCC are not materially different from the conversations with clients. So maybe you want to add some color there.

Anup Nair: So let me take a stab at it. So, all GCC conversations we are having includes automation and driving enterprise autonomy and every enterprise, every conversation, every deal we are doing fundamentally is based on the fact that we are bringing AI agents to accelerate the whole process of onboarding or exercise the process of delivering from a GCC center. So, like Sudhir mentioned, it is not much different. For us, it is quite similar. The same frameworks apply. The same solutions apply and what we have built over the last years is basically what we used to deliver that kind of an outcome, even for GCCs.

Lalit Wadhwa: And Srini, just to add to what Anup mentioned, I am not talking of today, but I am trying to project into the future, say two, three, four, five years into the future. If, as we discussed, a majority of the labor would be agentic labor and not human labor, does it matter whether it is being done on this planet, another planet, or this country, or another country? That is a question I do not believe any of us can answer in a very deterministic manner at this point in time, but it is a question that internally within Coforge we ponder about. If the labor is going to shift to agents, and humans are going to go ahead and take very significant strategic judgment calls around how to go ahead and build the next era of the enterprise, autonomous enterprise. Some of these current models could possibly also open up to innovation. So don't have a perfect answer to your question, but if we try to project into the future, things could be very different. Our current conversations

are exactly how Anup went ahead and projected to you, which is they are no different than what we have with any client based either in Europe or in the US.

Srini Vadapalli:

Got it. Thank you, Anup.

Moderator:

Thank you. Our next question is from Sourabh Agarwala from Pictet. Please go ahead.

**Sourabh
Agarwala:**

Thank you for your time. Just I have two questions and which are related to one another. This is related to the inversion and flip which we spoke about. So, what happens after two years and when you hand over the project, like what is Coforge role and responsibility with regards to that project and how does the revenue pool, impact for the agent and for the human after two years? One is that and second which is related to this is at client side, are there multiple companies which are deploying their platforms and who would be responsible to manage multiple agents which have been deployed by maybe multiple vendors? Yes, those were two questions. Thank you.

Anup Nair:

I will address that, Lalit, if you want to add to that right. So, the point that inversion is real, that is your first part of your first question. There will be a point in time, I mean, two years, one year, three years is a question that we can all ponder on, but the fact of the matter is that there will be a time where there will be significant amount of agents working with each other to deliver that outcome, right. We are not seeing the agents doing validating judgments. Agents will validate judgment, but humans would need to be involved even to make sure that the outcomes are achieved.

So, for a company like Coforge, right, we do not see this as a problem because at the end of the day, we are working on outcomes. We tie ourselves to the outcomes. We make sure that we have skin in the game, and it is about making sure the outcomes are not just, outcome is not just delivered once, but it is consistent. There is a significant amount of work by the way. If you look at it, agents do deliver an outcome today does not mean that it is going to keep on delivering the same outcome on an ongoing basis. You need to manage a lot of operations to make that happen and Ashish alluded to some of that during his presentation today. So that is the first part of your question and so you want to add something on to that, Lalit?

Lalit Wadhwa:

Sure, Anup. So, Sourabh, again, great question. I think it is such a massive opportunity that is probably very, very poorly understood. So, you go ahead and put together hundreds and thousands of agents into a client environment. All is beautiful. All processes are running. The backend models will continue to get more intelligent. Agents will drift because they are going to see different data over a period of time. They have to be retrained. Reinforcement learning has to come in. Reinforcement gyms have to be built. Just like humans have to be upskilled, they have to be trained. You have got to go ahead and retrain those agents too, in reinforcement learning gyms. All of those are opportunities for the future that we are working on at this point in time. So as the population of agents within companies or within firms will continue to increase, and frankly, they will go into thousands. The ability or the knowledge to go ahead and train and retrain them, upskill them, remove certain agents and inject new agents because workflows with change over a period of time, they never remain consistent. Customer demands, customer-facing agents will go through change because customer demands are going to continue to evolve. They are going to continue to change. All of those are significant opportunities for any enterprise which fully understands what the inversion looks like and where the inversion is happening and at what speed it is happening.

So, this is the thinking within Coforge. This is the path that we are on at this point in time. The opportunities are there. They will emerge, they will be incremental, but they will be very significant. It is not build and exit. It is build, exit from the build phase, but then move on to the next phase, which is all about reinforcement learning, which is about new workflows, which is about meeting new customer experience demands, so on and so forth. Hopefully Sourabh that gives you additional insight into at least how we think about this opportunity.

Anup Nair:

I will address the second part of the question. I mean, I guess Sourabh your second part was around building out. Of course, it is going to this is there is a race going on right in terms of who builds out the intelligence first. There are multiple parties that bring platforms. We are very conscious and we have made sure that we are not projecting it as a platform. We are bringing it as an enablement suite, which is a collection of things that will not rip and replace customer investments. We are going to play along with the customer to make things happen.

Look at it this way, right? At the end of the day, what are we promising? We are saying that we will build out an intelligence fabric for an enterprise that will keep on compounding and make the company better and smarter over time. And whichever company starts building out this intelligence first has the moat. Now it is on both sides of the whether it is a vendor or it is even the customer. I mean, if

an insurance company starts building on this intelligence layer sooner, than others, they will have a moat because they would have built out the intelligence and it is richer than someone else. So that is how we should look at it as opposed to platforms, because this is not about a platform play, right. This is a play of making, enabling enterprises to create a digital version of themselves not just from a data perspective, but from a decision perspective. That is how we should look at it. I do not know whether I addressed your second question, but if you have follow up, please feel free to ask.

Sourabh Agarwala:

Yes, that addresses my question and just the inference would be the client wallet will consolidate, right, if that is how the environment will look like.

Anup Nair:

Absolutely.

Sourabh Agarwala:

Okay, thanks.

Moderator:

Thank you. Our next question is from Rahul Maheshwari of Ambit Investment Advisors. Please go ahead. Please go ahead with your question. We are not able to hear you, Rahul, if you are speaking. Could you take off your headphone?

Well, there is no response from this connection. We will move to our next participant, that is Prateek Maheshwari of HSBC. Please go ahead.

Prateek Maheshwari:

Thank you so much for the opportunity and thank you for the session. So my question was for Anup and Lalit. So today I just want to understand what kind of, how many projects are going on in the enterprise autonomy, like just trying to understand from the company scale perspective and then I wanted to understand today, what would be the scaling problem? Is it the client's willingness problem or is on the FDE side where probably today you have to scale it to 150 and then probably larger cohorts, right? So just wanted to understand where we are, what the constraints are to the growth on this side?

Anup Nair:

Let me start. Okay, Lalit, please go ahead.

Lalit Wadhwa:

Thank you. Prateek, great question. When I look across the Coforge landscape and I look at the top 50 clients, every single client has a roadmap. Now, they may be at different places on that roadmap. Some may be figuring out how to go

ahead and finance this. Some may be deep into implementation. Some may be, have implemented two or three different use cases and now may have gone back and said how do we go ahead and solve the scaling problem around becoming an autonomous enterprise? But in top 50 clients, every single client completely understands what their vision looks like when it comes to autonomy. They have different versions of autonomy, again, because they belong to different domains, different verticals, and different industries.

The version of autonomy is a little different, of course, in very regulated industries. Autonomy will continue to keep human judgment in the loop for a longer period of time for deeper decisioning, and so on and so forth and in industries which are lightly regulated, the move is going to be much faster and that reflects in what the clients show, or we see, at least in the clients. There are some clients where, you know, figuring out how to go ahead and scale this out is a major conversation, but I do not see a client at this point in time who says, I do not know how to go ahead and finance this over a longer period of time. I think from a scaling perspective within Coforge, the most important thing that we will absolutely continue to focus upon is, this is not the job for a traditional software engineer. I think it is still not very well understood. An FDE skill set, the way the FDE produces outcomes, the move to the outcome, and I do admit every customer is not comfortable or at least ready to go ahead and move to an outcome at this point in time, but they will absolutely move to an outcome-based model. I think those are the things which are constraints that have to be, I would not call them demolished, but those are kind of constraints that we will run into and we will have to go ahead and find ways around those constraints. But if you ask me, are there any challenges which seem insurmountable at this point in time in terms of scaling or in terms of moving to an outcome-based model or in terms of building out a very differentiated AI delivery model, I would say that there are no insurmountable challenges. Anup, if you want to add anything else.

Anup Nair:

Sure. So just to add to that right, so the scaling problem right, there are a few big elements into it. One is of course, as I mentioned in my presentation, and I think multiple people mentioned this thing that doing this is a massive engineering lift, right? And it cannot be done by regular engineers, which is why FDEs. FDEs, the faster you were able to do, bring FDEs to market, the better it is and the right FDEs, by the way, not general FDEs. That is a big, big element. The other big elements are data, right because you see, none of this is going to happen without proper data and what I mean by data is that people have built data platforms, data lakes, and data warehouses. All of that does not really matter.

What really matters is, if that data is available for this intelligence to do something about it, right? That has been one of the stopping factors and we have

to, I mean, again, the way Coforge thinks is that we have to be able to unblock that right and you might be wondering the two U's in Nuuron. One 'U' is understand and second 'U' stands for unlock right. Our main concern is how do you unlock intelligence? Everything we do is around that. That is the data and the third you would expect it to be another tech problem. It is not. It is really change, adoption. There is no point in building any number of agents if people are not going to adapt it into their workflows. So that is the third big element. So that is how I would address the scaling question you had. As far as number of projects, either we are not giving any public numbers around number of projects we are doing however, let me just give you a certain perspective here. There are two angles. It is not just new projects. We are talking about existing projects as well. Any implementation we do has to eventually end up being building out this intelligence fabric. It could be a legacy modernization. It could be a AI engineering. It could be a KYC for that matter. All of that goes towards a single goal and that is enterprise autonomy.

Sudhir Singh:

I think just to add on to what Lalit said and what Anup said, the question around number of projects, percentage of revenue is a number that we are asked about often. We suspect that is a question that is not unlike the question around what percentage of your revenue is digital which used to play back about 10 years back. You have got the folks here who fundamentally run the company between Lalit, Anup, Deepak who runs data, Ashish who runs cloud and AI infrastructure, Erika, travel, Niraj, who is driving everything. There is nothing that moves.

There is no account, Lalit talked about it, top 50 accounts, where we will allow things to keep operating on a BAU model without AI based intervention. So, it is very difficult. It would be almost unfair to say what is the portion of revenues that can be ascribed to AI. We cannot think of any that is not and therefore when pushed we say 86% is from data, cloud and AI engineering but honestly very, very difficult for us to say that the BPS revenue is not anymore. One of the biggest mandates that Lalit carries is to agentify BPS ops right so that is the last layer that I wanted to add on to the answer to that question. Deepak and Ashish, I am not sure if you guys need to add something here or later, but that is your call.

Ashish Kumar:

Yes, Sudhir, I think especially from a cloud and AI infra perspective, all 160 clients are on the agentic ops journey, because that is the foundation and the goal here is not to take an AI project to production. The goal is to really, what Lalit talked about, to make it live in the production and AI agents will age, just like humans do age, and they need trainings. So I think most of our clients, in fact, all 160 clients whom we manage their cloud and AI infrastructure are on

different stages of agentic operations journey. I just wanted to add to that point.

Sudhir Singh:

Yes, Anup talked about data and maybe Deepak you want to add some more color there.

Deepak Manjarekar:

Absolutely, absolutely. So, I have 147 active clients and we run about 350 plus projects with these clients. Almost 100 plus are on their AI journey. Some are in extremely advanced stages and every time they bring a use case for agentification, workflow that needs agents and humans. After a few layers, it all boils down to data. Do you have your data right? That actually creates a larger program, which is data and, by the amount of questions that were asked on agents and, FDEs and all that, and not a single question on to data, I kind of wonder that all of that still remains below the waterline in an oblivion, but that is the stuff that makes the agent work, that makes this agentic system take and deliver their value. And client realizes that and they then say, yes, I need to talk about agents, but you need to make that data right for me. So that agent work, I get my promotion, I get my blurb in the media, but, you need to fix that and we are fixing all of those for each of our existing clients, depending upon where they are in their journey and most of them adopting agents and we are kind of creating RAC pipelines, CAG pipelines, CART models, MCP structures, whichever way they want to make these agents read their data makes it extremely contextual for their business operations.

And I think there was a question about IP. The agent is just the wiring and just the software code. We put them there behind the perimeter of client. We train them. We feed that context layer into them and then that agent becomes a personified agent within the corporation and behind the firewalls. That is their IP. The simple software is our IP, right? So yes, there is too much work going on right now in data and all of that is data but fueling towards the AI motion and adoption.

Sudhir Singh:

Thanks Deepak. Thank you. Inba, I know we are terribly overtime, but maybe a last question or two.

Moderator:

Sure. We will take our last question from Vibhor Singhal of Nuvama Equities. Vibhor, you may unmute your microphone and go ahead with your question, please.

Vibhor Singhal:

Thank you to the management team for taking this session for us and educating about the entire AI stack that we are operating. I will not hold you back, guys, for long. We have already done that quite a while. So, one basic question. I think

Ashish's presentation had a very interesting statement when he said that your AI bill is basically a design decision. So, I think I just probably wanted to dwell a bit more on that. We know that basically at this point of time today, the token costs, as you mentioned, are basically going through the roof, and that is also incubating a lot of application designs which clients are asking for. Now, looking at any technology development as it happens, the normal expectations would be that these token costs are going to come down over a period of time, as it happens with every technology development. So how do you see the impact of that on enterprise adoption of these applications? Does it mean that they will be able to free up more capital and maybe do more work, or would it mean that a lot of work that was being done by us in terms of application design, that can now be done by the LLMs itself, because now they can use them more with a lower token cost and the second question, since I have you all guys, maybe an out-of-context question?

We recently had the statement by the Anthropic CEO and seconded by Sam Altman also about slowing down of the AI progress. Now, how does that statement sit in context again with the enterprise adoption? Let us say, for example, an enterprise was looking at legacy modernization. So does it also slow down its legacy modernization decision waiting for when AI reaches that kind of a point or does this just snap and say that, okay, hell with it. I will just go ahead with whatever we have and ask the system integrators to help me with whatever is available in the system. I know a bit of hypothetical questions, but these are the ones which kind of are running through the minds of our investors and if you could help us with both of them, it would be really great.

Sudhir Singh:

Ashish, why don't you take first and then Lalit, if you can jump in.

Ashish Kumar:

Sure, let me attempt that. Thanks for your very, very intriguing question. So I think, we are seeing a tipping point. right? As I mentioned, 2026 is for the first time when every enterprise starts thinking about metering token uses because agents have become extremely capable, right? The biggest shift that we are seeing, two shifts really, the trend from most of our clients who are ahead in the AI adoption journey.

One is, you know, they are going not towards a model, a mixture of models. That is a big shift that we are seeing and therefore unlocking and having the choice and the flexibility that they need to implement. Second thing that we are seeing a biggest change is with our AI launch pad that we have implemented and this is early stage of working with few clients. The open-weights is a reality now and that is something is abstracting the engine, the deployment and the access

layer. That is where I think the industry is heading towards and that will drive the cost down.

In fact, we have in our research, we have looked at 7 billion token consumption is the tipping point when clients have to start looking at consuming and owning the entire stack versus subscribing to a frontier model. That is the tipping point in our experience and that number will go down further and further as open-weight and the sovereign AI becomes much more real. The cost is drastically down. So I think, that is probably, I would say, from the first part of the question and second part of the question, I will defer it to Lalit to answer. He probably has more experience.

Lalit Wadhwa:

Sure. Vibhor, thank you for that question. Again, this is a question where the answers are going to be imperfect. There are frankly no perfect answers. Anybody says that there are perfect answers for that, at least my belief is, they do not know what they are talking about. But the questions are great.

See, model intelligence continues to double every four months. That is the rate at which model intelligence is doubling in 2026. It used to be about six months last year. It is shortened to four months. Model intelligence is doubling every four months. So, over the course of the next 12 months, model intelligence will be eight times what it is today, right? Two raise to power three, that is eight times. Along with that, if you look at the last seven-month data, the unit cost of tokens has declined by about 100x for the same intelligence, but the consumption has grown by about 1000x. That is the reason you continue to see that the Frontier Labs, their revenues continue to go ahead and climb.

Anthropic is now running probably close to about 70 billion a year. Open AI is almost at a similar run-rate and so on and so forth. If you go ahead and look at OpenRouter or if you look at the data that is available on Hugging Face, you are going to find that open-weights token consumption now forms about 35% of all consumption of tokens. It was around 11% or 12% 16 months ago. Early 2025, that number was 11% or 12%. It is somewhere close to about 35% at this point in time. This is a direct outcome of the fact that companies have understood that as they go ahead and start consuming a lot more tokens, it is important for them to go ahead with open-weights strategy.

Frankly, Vibhor it is no different from the cloud strategy that most large firms have today. Most large firms today have a multi-cloud strategy. Most large firms will end up with a multi-model strategy and that model will be Frontier Labs plus XYZ open-weights models at this point in time. To the second question that you had I think we have to dig a little bit deeper into the motivations behind asking for the industry to slow down. Now, if there is somebody who truly believes at

this point in time that the industry is going to slow down, that the pace of AI development is going to go ahead and slow down over any period of time, I do not think so it is going to happen. I think the pressure from the top four or five Frontier Labs and the pressure from the open-weights, and these open-weights are not only the Chinese models, there are some amazing open-weights that are being built over here in the US.

You look at the Nemotron from NVIDIA, you look at Muse Spark out of Meta. If you go and look at the poolside models, these are all world-class models. I think the pressure from the open-weights side and of course the other Frontier Labs is going to be too much. Frankly, if you look at the administration's point of view, to what Anthropic and OpenAI had to say, the administration's point of view is, nope, we are not going to go ahead and slow down. This is a race. So, I do not have a perfect answer because there is a little bit of geopolitics involved over here. But these are the statements that have been made at this point in time. I do not know how this will end up, but at least what I do know is that many frontier labs and open-weights, they continue to see this as an opportunity to close the gap rather than slow down.

**Deepak
Manjarekar:**

I want to add a little bit perspective from the data side, Vibhor. There is a human cost that gets compared with the agent cost and the FDE cost, because that is where the new model is moving to. If these costs balance out, then there is zero economic value for this entire equation. So, what we need is on the right-hand side with the token and agent, FDE cost that has to be lower.

The way we are lowering the agent cost is the token economics and everything that Anup, Lalit, and Ashish talked about. The FDE cost again is a Momentum blue challenge and we are taking it head on. I just did an estimation review yesterday morning for an ANZ client. It is a bank. We are deploying the Data Cosmos assets there that reduce the efforts, that reduce the agent efforts as well as the human efforts and typically the work that we would have done with 30 people, we are essentially doing it with six people in half the time. That is leaving tremendous value on the table for the client, at the same time making very respectable gross margins for us. That is where the economic value, that was my third slide actually and that is the moat that we are creating. So yes, the entire economics is kind of sliding down but if we do it right it will still create a lot of economic value and a lot of new programs that people who want to take on that they do not have money for.

Vibhor Singhal:

Got it, got it. Great. Thank you so much for taking my questions and thank you guys for this wonderful session. Thanks a lot.

Moderator:

Thank you. That brings us to the end of the question session for today. Back to you Mr. Sudhir Singh, for closing remarks.

Sudhir Singh:

There's not much to say. I think Lalit did a fantastic job of summing up as he always does. Ladies and gentlemen, Sourav, Radhika and I are very keen that instead of hearing from three of us, you hear from the folks who actually run the firm. They are all in the market. They are not thinking of the market. They are in the market. The only challenge we have is how do we get time from them because they are on a flight most of the time, running from client to client and a lot of the stuff, hopefully, that you picked up from them is real, right? These are client conversations. Deepak was talking about the estimation he was doing yesterday for an ANZ client. We wanted you to hear from them. Please reach out to Radhika if you would like to have more in-depth conversations with any one of them. And if we do manage to find them between the flights that they are taking, we'll make sure that we set up time with them. Thank you so much for your time and your interest. We look forward to the next session around this. Thank you. Bye, bye.

Moderator:

Thank you. Ladies and gentlemen, on behalf of everyone at Coforge, thank you very much for attending. As a reminder, we will be making a replay of this session available on our website at www.coforge.com. Thank you again for attending and that concludes today's session. Goodbye.