

01072025 Bangalore AI Tour Keynote Satya Nadella

AI Tour Keynote
Satya Nadella, Chairman & CEO, Microsoft
Bangalore International Exhibition Center
Bangalore, India
Tuesday, January 7, 2025

SATYA NADELLA: Hello, everyone. (Cheers, applause.)

All right, it's fantastic to be back here in India and back here in Bengaluru. It's always unbelievable to come back and see the energy and the excitement, especially at a time like this. We are perhaps entering this next phase where we're going to go from talking about AI and admiring some of the new capabilities, whether they're in infrastructure or in models, to doing things with AI that are bold and big. That's to me, what's really in the air when I come here and I see people excited about what's happening and what you're doing.

I want to spend the next half hour or so just giving you at least what I see out there as the possibilities and what we at Microsoft are focused on in terms of building platforms because at some level, to me, Microsoft has always been fundamentally about two things. We are a platform company, and we are a partner company. Even in the age of AI, that's what's going to be true.

But whenever you talk about these platforms and platform shifts, you always need to be grounded in what's the foundational force that's driving the platform shift. In fact, when I look back even my 35 years in tech, it's been about one fundamental force, which has been Moore's law.

I always recount, Bill used to bring a bunch of us together every year, and he would just literally put the Moore's law, and then what's happening with memory and say, go fill it with software. That was just the single instruction to the entire company. It's very true even today, right? When you think about the scaling laws that are powering AI, and pretraining in particular, it's really Moore's Law at work again.

It started in 2010 with DNNs, and then obviously the GPUs, they inflected again, perhaps with transformers, just because of the efficiency of data parallelism with transformers, and what was happening, perhaps, of doubling of capacity every 18 months started to double every six months. That's really what the scaling laws were.

By the way, there's all this debate, what's happening with scaling laws for pretraining, will they continue, have we hit the wall? We fundamentally believe the scaling laws are absolutely still great and will work and continue to work, but they do become harder, right, as the sizes of data become higher, the parameter counts become higher, the

systems problems are bigger. These synchronous data parallel workloads are new workloads, and so therefore that will continue.

But the more interesting thing that you are now starting to see is another scaling law with the inference time or test time compute scaling law. At some level, pretraining had that sampling step. This is sort of always about using that sampling step more efficiently. It's about being able to scale with even inference time that I think is going to really take this to the next level.

We're very, very excited about it, o1 and think hard, like to me, Copilot, think harder has become the thing that I go to all the time, and that shows not just for sampling for additional pretraining, but to be able to use it during inference time to think harder and get you better results.

I think that these are capabilities that are just going to increase, ultimately, three things. One is this multimodal capability that we now have as the interface to all software. I recently set up, I guess there's this, I don't know what they call it, the action button, I think, on the iPhone. It's now set up to Copilot for me, and now I can sort of confidently speak in Hyderabad Urdu to it, and it's beautiful. It understands me. It's like speaking to my high school buddies. The fundamental idea that you can now have that familiar, simple interface with all of computing, I think it's going to change every software category.

Then you couple it with these planning and reasoning capabilities. Whenever you go to GitHub Copilot Workspace, that thing is invoking that ability to think about planning and executing that plan as a multistep process. It's the beginnings of what is truly agentic behavior.

Then the other side of it is to be able to stitch these things that are outside the system or outside of the model: memory, context, tools use and entitlements. In fact, if anything, in the next 12 months, every developer is going to be really focused on, how do I take my model and make it aware of the tools it can use, just beyond even function calling? How can I make sure it has the right understanding of its entitlements. How do I make sure it has memory and long term memory?

That, to me, I think, is what's going to truly help us create this rich tapestry of agents. When we think about agents, you stitch the multimodal capability, planning and reasoning, and memory and tools use in particular, plus entitlements, you can start building personal agents, team agents, enterprise-wide and cross-enterprise agents. That agent world is what we are looking forward to sort of all building.

Of course, at Microsoft, for us, it's never, quite frankly, about any of these technologies on their own. It's a means to the end, which is about empowering every person and every organization on the planet to achieve more. It's that sense of empowerment that this platform can provide, I think, is going to be next level, and that's what we're really, really focused on.

To that end, we're building three platforms: Copilot, Copilot and AI stack, as well as Copilot devices. What I want to do is just kind of give you the broad contours of what these three platforms all entail.

The best way to conceptualize Copilot is it's the UI for AI. One of the ways you think about, even in a very rich agentic world, remember, the AI will need to interface with us. That means the UI layer. That's why I think this organizing layer of Copilot becomes even more important in a world where there are many, many agents doing autonomous work, I think that's the best way to think about it.

Now, the approach we have taken is to build in the Copilot into the existing workflow. One of the best examples I've seen, quite frankly, is one of the high stakes things of knowledge work. There's a doctor. She's getting ready for a tumor board meeting. Think about a tumor board meeting. It's a high stakes meeting. That means you would have had to read all of the reports, know exactly how much time to spend on each one. That means the creation of the agenda is a reasoning task. It creates an agenda, which knows which one is the more complicated case that needs more time.

Then you go into a Teams meeting. All these doctors are now having a conversation on all the cases, they are able to focus on the case, as opposed to taking notes, because there's an AI that's taking detailed notes on all of it.

Then at the end of it, she's a teaching doctor, right? She wants to be able to take what happened in the tumor board, go to class. That means she's able to take the notes, putting it into Word, from Word to PowerPoint, go to the class.

That simple work flow that doctors everywhere are doing that has impact on life can be enabled with AI just being built into the workflow. That's just a good example of sort of how to think about AI being infused into current work flow.

Now let's take it to the next step. Now with Pages and chat, with web and work scope, I think it's sort of completely about thinking of new types of work flows. Fundamentally, I am now able to access information. Whether it's web information or it's information inside of my Microsoft 365 Graph, one query at a time, I can go get back all the data. I can then promote that data into this interactive, AI-first Canvas called Pages. Once I have it in Pages, I can then use Copilot in line in Pages to keep modifying it.

I sort of use this as a meta for where I am thinking with AI, and I'm working with my colleagues. Think about that. That's the new workflow where I think with AI, I promote things into Pages, I invite others, I collaborate with others, and by the way, AI is present even on that canvas.

Chat plus Pages is going to become this new AI hub for just like how Word, Excel, PowerPoint back in the day, changed how I worked, now chat plus Pages will be a new module that effectively enhances how you work with AI.

Now, we're not obviously stopping there. The next thing that we are thinking about is extensibility. How do you extend AI going forward? It starts with something called Copilot actions.

For those of you who are big users of Outlook rules, which I was for a long time, until the complexity was too much, think of it as rules for the AI age, but they don't work one app at a time. They work across the entire M365 system. That's the beauty of actions.

If you think about so much of our workflow is gathering information, distributing information, it's about connecting people to artifacts, other people to other artifacts. That's a lot of what knowledge work is all about. I can now set these up essentially as Copilot actions. That's the first extensibility.

Now, of course, you can build full agents, and we ourselves are building many of these agents that have scopes at group level, process level. You can have a project agent. We have agents that are working inside of Teams, like an interpreter, like a facilitator. This is just like having an additional team member that is helping you with your task, inside of your team.

SharePoint, in fact, every SharePoint now has an agent. Think of it as an intelligence layer on top of SharePoint that just built in.

I just want to roll a video to give you a flavor for these agents that are being built in M365.

(Video segment.)

(Applause.)

SATYA NADELLA: That's just an example of essentially agents that we have built into the system. But the real exciting thing is, of course, you all being able to build agents, and that's where Copilot Studio comes in.

Our vision with Copilot Studio is simple. This is the low code, no code, tool for building agents. Think of it like when you had Excel and you could build spreadsheets. To us, building agents should be as simple as building spreadsheets.

Copilot Studio is about helping every one of us to have real agency, to shape and reshape the workflows around what we are doing as knowledge work. That's essentially what we want to be. It's the swarm of agents around what we do that is helping us get work done, create more flow, less drudgery.

Copilot Studio, like take something like field service. It's as simple as first giving it a prompt, giving it the instructions on what the agent is all about. Then it's about grounding it in knowledge sources. Like in this case, it's about pointing it to the right

SharePoint source. Then once you do that, it just creates an agent out of the box for you. That simplicity of being able to create agents that are essentially no code, no code programmable is what we are doing with Copilot Studio.

Now that you have this Copilot UI for AI, you have the ability to extend it with actions, you have the ability to use agents that are built in and build your own agents, you have a complete system.

Now the question is ROI and measurement, because that's the other question, which is, okay, because one of the fundamental things we want to also ensure is that there's real motivation for change. After all, what am I doing that is better, that is helping improve not only my own productivity, but my organization's outcome?

That's where this measurement comes in, and we are building out these Copilot Analytics so that every individual, it's not just a top-down thing, right? A sales territory manager can now go in and take an output metric, something like increase sales, increase yield, and correlate it back to specific usage of all of these Copilot features.

To me, that's another one that drives the adoption cycle, so you're not waiting, but you're able to see in real time how, with increased usage, you're able to drive business results. That's really the first platform we're building, Copilot as the UI for AI, with extensibility and measurement.

Now, we're already seeing fantastic results. Inside of Microsoft, basically we are baking in double-digit strong improvements to productivity across every business process: customer service, HR self-service, IT apps, finance, supply chain marketing.

Think about marketing where there is sophistication in buying, but there are a lot of places where there's a very lot of inefficiency around content creation. There's a massive amount of leverage and operating leverage we get there. Significant use cases across the length and breadth of our own company.

Of course, when I come even to India, and I get a chance to meet with everyone here, the diffusion is so fast. It's no longer about waiting multiple years before it becomes mainstream. I learn a lot by watching many of you deploy this at scale. In fact, this morning, I had a chance to talk to folks at Cognizant. They were telling me about how they've deployed it across the entire sort of employee base.

The knowledge turns, one of the things that Andy Grove way back in the '90s used to talk about was knowledge turns, which was about being able to create knowledge fast and really be able to then diffuse that knowledge. That's just like supply chain terms, people talk about them inside of retail. This is about knowledge turns in any knowledge industry.

Another example was what Persistent is doing. They built a contract management, essentially agent. In fact, you can just go address it at Persistent, whatever, Contract AI inside of Copilot, you can get access to it, and that agent is available to you throughout

the entire lifecycle of a contract management, where even a single change can have significant impact. These are just two couple of examples of people already deploying this Copilot system at enterprise scale.

Now I want to talk about the next platform, which is the Copilot Stack and the AI platform.

Now to us, we've always conceptualized and built Azure as the world's computer. We continue to be super committed to it, because one of the fundamental realizations is AI doesn't sit on its own. It requires the entire compute stack. We are building that out at a worldwide level. We have 60-plus regions. We have 300-plus data centers around the world.

In India, we are excited about all the regions we have. We have Central India, South India, West India. Then we also have the capacity that we built up with Geo. We have a lot of regional expansion happening.

I'm really, really excited today to announce the single largest expansion we have ever done in India, by putting 3 billion additional dollars to expand our Azure capacity.

(Applause.)

I had a chance to meet the Prime Minister Modi yesterday. It was fantastic. It was great to listen to all his examples, his vision around how he wants to drive through AI mission. But it's that combination of really the Yojana he has, the India Stack, the entrepreneurial energy in this country, and the demographics, both on the consumer and the business side, that is all getting into a virtual cycle. That's why we feel fantastic about bringing core compute capability for the next generation AI.

Now, with infrastructure, there is in some sense a new formula, quite frankly, for any country or for any company. I think of that formula as tokens per dollar per watt. That's it, tokens per dollar for watt. Two years from now, five years from now, 10 years from now, we will be talking about the correlation, quite frankly, between GDP growth in any community, in any country, in any industry, or even in a company level, fundamentally their own growth, on how efficiently are they able to drive that equation.

To that end, that means infrastructure, infrastructure, infrastructure needs to be the highest priority, and we are innovating in every layer of it. Think about it, at the data center level, like these data centers are just everything from how we think about even the construction of a data center that is optimized for liquid cooled AI accelerators is a new engineering feat.

That's what we are doing, everything from how do we then work with upstream from us with renewable energy folks to get onto the grid base load that comes to our data centers, that then has the right cooling infrastructure with zero waste and zero water usage, how do we really make sure all of that gets built into as a system level.

Then, of course, silicon innovation. We are innovating with Nvidia. In fact, today I think we have our first GB-200 clusters up live at our data centers. Very excited about what that would mean. We are then also working with AMD.

We are building out our infrastructure. We're building our own silicon with Maia. In fact, Maia today is taking a lot of the customer service traffic in Microsoft.com. We are building, I would say, world-class AI accelerator infrastructure. But the entire system stack above it, optimized for training, optimizing for the kernels for inferencing, that is significant investment and innovation that's happening through us, our partners, and so we're very excited. In fact, I think of this as the golden age for systems when it comes to innovation.

Now, if you have the infrastructure, the next big, big, big consideration, in fact, today I met many, many partners, many customers. The first thing everybody wants to talk about is, it's fantastic we're talking about AI, but how do I get my data in shape? Data is the only way to create AI. It's not just for the pretraining. We know for RAG, you need data. You need data for post training. You need data for doing sampling, for doing inference time compute to improve pretraining. Data pipelines and data is everything.

The first thing to do is to rendezvous the data with the cloud, and that's where we are building out our data estate, such that you can bring all of your data. You can bring whether it's Snowflake, whether it's Databricks, whether it's Oracle, whether it's our own SQL, what have you, bring it to the cloud. We have fantastic operational stores that are being plumbed for AI. Whether it's Cosmos DB or whether it is SQL Hyperscale or Fabric for analytic workloads, all of these are AI ready.

In fact, if you look at even ChatGPT, they are some of the biggest users of Cosmos DB, because they're stateful applications. Where is all the user state of ChatGPT? It's in Cosmos.

To me, the data layer is a super important layer, and we are doing everything to make sure that we can help get the data in shape for you to be able to then use in conjunction with these models and to build models. Those are the two things. There is models that are being trained on the data, but you are also doing things like retrieval augmented generation using data. That's sort of why the proximity on the data gravity is huge. Locality of data will matter.

Now, once you have the infrastructure and data, the third thing you have to do is to have the AI app server. In fact, if you look back, when the web happened, what did we do? We built IAS as the app server. When the cloud happened, what did we do? We started building the cloud app, native app server. Same thing with mobile. Every generation has required an app server, and that's what we're doing with Foundry.

With Foundry it starts with models. Obviously with OpenAI and OpenAI's innovation now with o1, we're excited about what's going to come with o3, with obviously 4o, all of

that's available, plus all the open source models, whether it's from Llama Mistral. We have industry specific models, models that are being built out of India for Indic languages, for India's vertical specific needs. It's fantastic to see the amount of innovation that people are doing around models all over the world.

We want to have the richest model catalog, and some of the more popular models will be available even as models of service. There will just be underneath a facade of an API that you can then go access.

Now, once you have the models, you want to deploy these models, you want to be able to fine tune these models. You want to be able to distill these models. You want to be able to do evaluations on these models. You want to be able to do groundedness tests. You want to be able to do safety. All of that, instead of building them all separately, we're building them all into the app server.

In fact, evals are going to be the most important thing. Even for me, the guidance idea I give our teams is simple, which is, stay on the frontier of the new model, and then make sure that you have agility in the app server layer, so that you can keep moving with models. You will use the first, like the latest sample. Then you will cost optimize it, latency optimize it, and you will start fine tuning it for your specific use cases, for evals. That's sort of the loop that you are constantly going through. That's the idea behind Foundry is to just streamline all of that.

Fantastic momentum, again, in India. When I look at customers who are already deploying this, using it, in fact, lots of good feedback I'm gaining from lots of people. People are pushing on even these multi-agent type of deployment. We're learning a lot, quite frankly, from some of the ambitious things that you all are doing, and we are definitely very, very grounded in how we are going to progress on that roadmap.

I think next year we'll not be talking as much about models, but we're going to be talking about model and model orchestration, model evals, and how you're able to deploy these model-forward applications. That, I think, is going to be the big shift across the industry.

Today, when I saw the folks at Bank of Baroda, they showed me three things. They showed three agents. They built a self-service agent. For the new customers they built it, basically a relationship manager agent. They also then built an agent for their own employees.

I had a chance to see this fantastic start up, ClearTax. I guess I should be able to do my taxes on that, which is simple. On WhatsApp, you go submit your receipts, I guess, and then you get refunds. I love that part in particular.

Then I had a chance to see the ICICI Lombard folks. One thing I had not realized in health care in India, for example, that you don't have standardized claims form. Each one is a novel, different entry, and so someone has to sort of go read it. That's where I think you can improve. You think about any time you improve healthcare efficiency, that

improves the economy, because you then have assurance of your insurance being taken care of.

MakeMyTrip, I had a chance to meet with the team, and they're doing some phenomenal, high ambition work of being able to take one of the most sophisticated industries and verticals around travel and really, whether it's hotels or whether it's air or whether it's other transportation, how do you really have a multi-agent framework that really they can deploy?

But one of the other things is it's not just about the broad, big companies or startups that are doing it, but it's sort of the diffusion rate of this technology in India is what's exciting.

To that end, I just wanted you to see the video from the cooperative (Bar Amity?). Go ahead and roll out the video.

(Video segment.)

(Applause.)

SATYA NADELLA: To me, it sort of really connects all the dots. It sort of connects the dots between even all the technologies we have been building, from Azure IoT, the data connectivity back to a data plane, and then to be able to use something like Azure AI, but ultimately to empower a farmer to be able to do their farming with higher yields. That, I think, in some sense, speaks to the power of this technology and what we can do with it.

Now, the last layer I want to talk about, if you have infra, you have data, you have the AI app server, is tools.

Now, I always go back, Microsoft started as a tools company. We continue to be super passionate about our tooling. With GitHub it's fantastic to see what's happening. We have now 17 million members of GitHub in India. (Applause.) It's the second largest community, next to the United States. In fact, it's projected to be the largest, I guess I forget the year, I think 2028. I can't wait. Three years from now, 2028 is when the crossover will happen, where India will have more developers on GitHub. It's just exciting to see that.

Now, we also have contributions from India to AI projects that are just second to the United States. It's fantastic to see the active involvement of the developer community out of India in making progress on all the open source projects on GitHub when it comes to AI. It just speaks to, again, the talent there is and the energy there is in this community.

Now, we're continuing to make great progress on GitHub Copilot. One of the features that I was looking forward to we now have, which is the multifile edits. First, we started with continuations. We then had chat. We brought continuations plus chat together, and then brought it to multifile. That's fantastic, so that I can do repo level edits.

Now, the other thing we are also doing is bringing a free tier. In fact, we just launched it in December. I think in India, it's sort of like the place where it's really taken off. We're very, very excited about bringing to VS Code, GitHub Copilot free tier, and seeing that now broadly getting distributed.

The feature or the product area back, I forget now, 2019 now, or maybe not 2019, 2020, I think, is when I first saw GitHub Copilot. That's when my own conviction on what LLMs can do completely changed, when I started seeing it.

Similarly, when I first saw GitHub Copilot Workspace is when I felt like time had come for us to make the next leap beyond chat to real agents, because that's what it is. GitHub Copilot Workspace is the first agentic sort of piece where you now can take a GitHub issue, create a spec, which you can edit. You can create a plan. You can then edit the plan, and then you can see it execute across the full repo.

To show all of this, I wanted to invite up my colleague, Karin (ph), up on stage. Karin?

DEMOER: Thanks, Satya.

While Copilot offers contextualized AI assistance in your developer environment, Copilot Workspace is the next evolution, an agentic, AI native developer platform that helps you turn your ideas into code using natural language. Let's see it in action with our computer.

All right, we have an application here that lets us sell sporting equipment online, including some cool cricket gear. However, there's no way to add new products through an admin page. Let's see if we can build one using Copilot Workspace.

I have an issue here that describes my requirement, which is an admin page to add new products. Usually, I create a new branch and start thinking about how to implement it myself. But with Copilot Workspace, I can get started right here from an issue. Copilot Workspace gets to work right away on how to solve this issue. I am in control, so I can make edits to the plan as needed.

With the new Brainstorm agent, you can use Copilot Workspace iteratively. Based on the context, Brainstorm shares a few suggested questions to get started. Or I can even ask my own.

On the admin page we are creating I'd also like to add a requirement to upload product images. Let's ask Brainstorm with Copilot Workspace and see if it can make it happen. This time, I'll ask it in Hindi. There it is. Just like that, it's listed a few ideas for me, use a file input element, a drag and drop area, or a third party library.

Since I already have an existing component, I'm going to use that. I can click Add to Task and add it to my context right over here. Looks good, so I will click Generate Plan

to move forward. Workspace then generates a list of what changes to make in which files. When I'm ready, I can ask Copilot Workspace to implement the plan and make the necessary changes.

There it is. That's nice. Copilot Workspace has made changes to multiple files, including writing a whole new page.

With Copilot Workspace, you can now keep iterating on your plan and implementation by asking for a revision. Let's ask Copilot for a revision. I'd like to ask it to show me a preview of the image I upload.

This time, I'm going to ask it in Kannada. I'll say, [in Kannada], and there it is. Awesome, right? Copilot Workspace understood what I meant. It updated the plan and also implemented all the necessary changes. That's great, right?

All right, now before we proceed further, let's run some tests first to make sure my code doesn't turn into a horror story in production. We made some great enhancements to the commands you can perform in Copilot Workspace. Build, test and run are pretty common scenarios, and I have the commands configured here. I can execute the test command, which will run the test for me right here within Copilot Workspace.

Looks like my test failed. My bad. I thought my code didn't need test. But hey, it's Copilot to the rescue. I can ask Copilot Workspace to fix this for me. The Build and Repair agent within Copilot Workspace comes up with a solution on how to fix this error, which I can apply, and then it goes back and makes the necessary edits, in this case, modifying the test. Let me go back and then rerun the test, hoping that I am as good as the test where I think I might be. And there it all works. Great.

Now I think my feature might be ready, and I want to see a preview of my app. I can execute the run command right here within Copilot Workspace, which will bring up a development server that's accessible right over here, and I can open up a preview of it, and there's my admin page, along with the image upload functionality that I needed, completely written using Copilot Workspace. (Applause.)

In the past few minutes, we went from an idea described on an issue, brainstormed with Copilot Workspace, implemented code and also fixed errors, all while working in natural language. This is the AI native developer environment, an agentic workflow that moves as fast as your creativity.

Thanks, everyone. Back to Satya. (Applause.)

SATYA NADELLA: It's really exciting to see the progress in the dev toolchain. In fact, as of today, there is no more of a wait list for Copilot Workspace. We're very excited about that.

To me, even for me personally, perhaps the biggest game changes were Windows 365, where I have my dev desktop, plus GitHub Copilot and Copilot Workspace, plus Code Spaces. You put those things together, put me anywhere in the world, I'm a happy person. This has just been a wonderful sort of change in terms of dev productivity everywhere.

The last thing I want to talk about is the Copilot devices. We talk so much about all this innovation and infrastructure, starting with silicon in the cloud. Now that's coming to the edge.

We are very excited about the work we're doing with Qualcomm, AMD, Intel, when it comes to the NPUs. In fact, today, Jensen even talked about some of the next generation GPUs that are coming to regular PCs, which will be able to really run the entire Nvidia stack locally. We are very excited about what's happening with these Copilot and Copilot PCs broadly and even traditional PCs with GPUs.

But we're also excited about the fundamentals. When I use my Copilot PC, having the battery life last all through the day, having now these new AI features that are built in, like Copilot being built in, third party developers are beginning to start using it, whether it's Adobe or Capcut or others, it's a real beginning of a new platform on the edge that's going to be as exciting as what's happening in the cloud.

In fact, we don't think of this as the old client-server. This is not about disconnected, local model. It's about hybrid AI, the idea that you will now build applications where you will be able to do a bunch of the stuff offload on the local NPU as helpers, classify in here, calling LLMs in the cloud, so any application will truly be a hybrid application. It's not about running locally or running all in the cloud, and that, I think, is what we have looked forward to.

Let's play the video to just give you a flavor for everything that's happening with Copilot devices.

(Video segment.)

(Applause.)

SATYA NADELLA: Now, in order to really ensure that these three platforms get broadly distributed, broadly used, the key consideration is trust, trust around security, privacy and AI safety.

We have a set of principles, but more importantly, these principles and initiatives are grounded in making real engineering progress so that we can effectively ensure trust all along the way.

Take something like security. How do you protect against adversarial attacks, like say, prompt injection? That's the key thing that we are building in.

Take privacy and AI. What does it mean to think about confidential computing, not just for the CPUs, but in GPUs? That's what we now have with everyone, whether it's Nvidia, AMD or Intel.

Or when it comes to AI safety, one of the big things people talk about is hallucination. How do we ensure groundedness? A groundedness service with eval support is one way for us to make real progress on AI safety.

We're really taking trust as a first-class engineering consideration, having a set of principles, but more importantly, translating those principles into essentially the toolchain and runtimes that allow us as developers to build more trustworthy AI.

Now, I want to close where I started, which is our mission to empower everybody. But before I go there, I think the thing that I want to talk about is all this is about doing AI business transformation. Ultimately, it's about changing customer service, or whether it's about changing your marketing or sales or your internal operation. At the end of the day, it's about business results.

The three considerations I would submit are around Copilot as the UI for AI. It is about being able to sort of make sure you're thinking about the app server as the platform around which you build your AI applications, which is for us, Foundry, and then your data in Fabric. These three, perhaps, are the three key design decisions that need to get made, more so than any given model, because models will come and go every year, every month you'll have some new model. But the three foundational design choices is really, what's the UI layer and how do your agents really interface with that UI layer, how do you think about your data, and how do you think about that app server that, in fact, gives you agility on top of models? Those are the three considerations. That's what hopefully you get through what we have done throughout the platform.

Our mission to empower every person and every organization right here in India is what drives us, and to that end, ultimately, it is about being able to ensure that the human capital of this country is able to continue to scale to take advantage of the immense opportunity and potential this technology has.

That's why I'm very excited to announce today that our commitment, which we have always had, now is to train 10 million people in India around AI skills by 2030.

(Applause.)

To me, the thing that is always most important is to not think of the skilling in the abstract, but to see the skills translated into impact right here, one community at a time, one industry at a time.

I leave you with a video of already the impact all of the skilling is happening here in India.

Thank you all so very much, and thanks for all the great work you're doing on all the platform. Thank you.

(Video segment.)

END