

AI Tour FY26 - Mumbai – Keynote
Satya Nadella, Chairman & CEO, Microsoft
Karan MV, International Developer Relations, GitHub
Zoe Hawtof, Senior Technical Advisor, Microsoft
Puneet Chandok, President Microsoft India
Jio World Convention Center
G Block, Bandra Kurla Complex
Mumbai, Mumbai 400098
Friday, December 12, 2025

(Applause.)

SATYA NADELLA: Good morning. It's wonderful to be back in India and here in Mumbai, and have a chance this morning to talk about AI, and mostly, to celebrate all the AI and the AI solutions that are being built by people in this country, in this state and everywhere. And it's just exciting to be able to come back and see the progress, even year over year, that I see in action.

To us, it's always about empowering people and organizations. When we talk about even new technology platforms, it comes down to, how is it being really used to empower an institution, an organization. I'm going to talk about, in fact, Crime OS that the Chief Minister was here, really sponsored and made happen in this state so that even every citizen of this state can benefit. And that's the type of innovation we want to see in any place.

But the key thing for us is to be grounded in the fact that the frontier of what is possible is changing. I mean, that's the exciting part for any of us, is when we talk about technology, it allows us to reimagine what we can now do, whether it's in our state, whether it is in our organization, whether it's in our lives, the ability to move the frontier.

We always start there, right with the customer experience. Now, you can think about and reimagine it, whether it's the employee experience, whether it's the operational efficiency, whether it's the rate of innovation. If you're a drug company, can you really get through clinical trials faster? Can we get more new drugs out there? That, to me, is ultimately moving of the frontier that we are all chasing.

And to do that, though, you have to embrace a new approach, even, what worked previously. In some sense, you have to do this hard job of unlearning and learning what's possible. Sometimes we pattern match too much from what worked in the previous era may not necessarily work. In fact, there's a new frontier.

When I think about that, it starts with that mindset, freeing ourselves, that there's a new approach. For example, when it comes to AI, if you look even at building of an AI solution, you don't start with a spec. You actually start with the test. I mean, when people

– there’s fancy new words. People talk about evaluations. They create these rubrics to do eval on a particular outcome. Then you create a learning system. Then you will climb on that learning system.

Guess what? That’s literally inverting the process of building an information system. You start from the outcome, you create a robust eval, and then you spec versus going left to right.

But the key thing is that means you’ve got to pick up new skills, all of us. And in order to pick up the new skill, you need a new toolchain that you are using every day and to be able to then have all of your data that is being brought to harness this new toolchain, this new skill set, with the new mindset. That’s really the way we want to approach it.

To facilitate this, this can’t just be an abstract framework. This has got to be an everyday practice for all of us. It’s an information worker today or a software developer. All of us are going to be using essentially, the same approach.

To realize that, we built a tech stack. And I want to just literally walk up this stack or walk down this tech stack, starting with the experience layer, where we have taken this idea that, let’s bring this AI Copilot where you work. You may be in Outlook. You may be in Word. You may be in Excel. You may be in Teams. Wherever you are doing your work, how do we bring the Copilot there?

And think of Copilot, though, as not just a request response. That’s what we’ve been doing now nearly for two years. We know now that, hey, we can ask it any question, it’ll come back with a better answer than anything that we have previously experienced in, let’s call it, an information retrieval search type experience, but now we have agents.

In fact, think of Copilot as your browser for the agentic Web. In fact, right inside of Copilot, you have very powerful agents built in. One such agent is your Researcher. And Researcher with computer use is a real breakthrough.

It’s literally like having 24 by seven, someone very sophisticated, who is your research assistant, who can go take any topic and then, because they have a computer. They just are not – think of it as not a single shot, I go and bring back information after thinking. The Researcher can use a computer to do research. It’s not only you have a computer, your researcher has a computer. That’s the way to conceptualize.

Same thing with Analyst. I can load a whole bunch of Excel spreadsheets and give it to an analyst, and it’s like giving it to your data science department or the data scientist. It’s, again, having a 24 by seven data analyst who can then go give you insights.

And lastly, it’s like one of the things that I’m very, very excited about is what happened in software development is going to happen with Excel and Excel Agent. I, in fact, had a chance to participate recently in the Excel World Championship. By the way, ESPME

even broadcasts Excel championships. We have an honest to god, great way to celebrate the Excel promise.

And so, I was able to participate in it, that digital challenge using this Excel Agent. It is so cool to be able to see, and I've worked on Excel all my life, essentially. And to be able to see how the Agent Mode understands deeply the entire semantics of a spreadsheet, it's not just a one shot, I created a great model, but it's the ability to iterate on a pretty sophisticated model continuously using this Agent Mode, just like how we do it in software engineering.

It's just exciting to see this, and also it's really, to me, at the end of the day, all these tools are like the new Word, Excel, PowerPoint, Office. That's as simple as that. And they're in the hands of everyone here, and I see it being used like that. No one is praying at the altar of these tools. They're just using them.

And to me, to see the proliferation, the diffusion, is probably faster than anything I've seen. I've come to back to India, come back to Mumbai, throughout every generational shift in technology, starting in the early '90s, but when I see now the rate of pace of change, there's no difference, quite frankly. I could be in any city anywhere in the world, and to be able to see all of this applications just being broadly diffused.

To show you some of this in action, let me invite up on stage my colleague, Zoe.

Zoe, take it away.

(Applause.)

ZOE HAWTOF: Thanks, Satya. I'm Zoe, a product marketer on the Microsoft 365 Copilot Team. And like all of you, I start my day commuting in and preparing for the day ahead. For that, I'm going to use Copilot on my mobile.

Hey, Copilot. What's on my agenda for today?

COPILOT AGENT: You've got a few things lined up today. First, there is an AI Tour India meeting that started earlier this morning. Then at 11am, there's a Contoso sales call. Later in the afternoon, you've got a Teams Think at 2pm.

ZOE HAWTOF: Thanks, Copilot. And has my manager sent me any emails?

COPILOT AGENT: Yes, you've got a few emails from Nicole. Just yesterday, she reminded the team about a Friday mixer happening in the common area. She also mentioned that finalizing the monthly business review should be a high priority. She also

—

ZOE HAWTOF: Thanks Copilot. Copilot doesn't just have my work context. It knows who I am, what I worked on, and who I work with. It's why I use Copilot. It's a true AI assistant.

Now, I want to take a look at my manager's ask around that MDR updates from my V Team. And you'll see I'm actually using GPT 5.2, which we released yesterday, and it's immediately available to all Microsoft 365 customers.

(Applause.)

Yeah, it's exciting, I know!

What I love about GPT-5.2 is it's allowing Copilot to be even better at finding the most relevant resources. My responses are very contextually relevant, and I'm able to take quick action, and it makes a very visual output. Now, only Copilot is able to combine this constant model evasion with your work data to provide the most relevant response.

Now, I want to go and update my existing monthly business review document, but I want to change it with my November data. For that, I'm going to use Agent Mode in Word.

Agent Mode allows the agent to edit, to draft, to format the document, all with natural language, but it still keeps me in complete control. And I can see it's starting to look across all of that work data, my emails, my meetings, my team chats, my files.

Now, that normally takes me hours to do. I have to find the right information. I have just sift through it, but now you'll see the agent is actually, actively applying it on the lefthand side. And you'll see that it's actually leaving parts alone that haven't changed month over month, like our mission, but rather using this to update everything and keep me on par with exactly the November update.

Now, last, I want to make sure I understand our team budget for the next fiscal year. I'm going to again go to Excel, and I'm going to turn on Agent Mode and ask it to get me a little bit more aware of what's been spent, what's uncommitted, what's at risk, and to build a visual dashboard so it's easy for me to understand my data.

Agent Mode in Excel natively speaks Excel, just like Satya said. It's not going to create a random number that I don't know where it came from. It will actually build formulas right in my spreadsheet, right where my budget already exists. And it will even create native objects, like charts and tables, things that I can continue to manipulate.

It also walks through everything it's doing on the righthand side. That way, say my manager, Nicole, comes to me and asks, how did I come to my conclusion? Well, Copilot that actually helps me understand what it did. That way, I can continue to work with it.

I can also ask Agent Mode to do simple things like conditional formatting, which you'll see, highlighting things in red that maybe are at risk, all of these things to continue to work with Copilot and Agent Mode in order to get a better understanding.

And you can see it's created this dashboard live. If I click in, I can see these rich formulas it's actually built to give me an answer, the conditional formatting to highlight what's at risk in red, and a native graph that actually, I'm able to manipulate even further, should I want.

Now, Copilot, it hasn't just made me better and faster at my job. It's helped me make more impact, and I can't imagine working without it.

Back to you, Satya.

(Applause.)

SATYA NADELLA: Thank you so much, Zoe. Whenever I see that Excel demo, I like to think about that when I had hair, I loved Excel, and when I don't have hair, I love Excel. (Laughter.) I tell you, it's one of those generational things that just gets better and better in spite of my follicle challenges.

To me, if you take that experience layer, one of the things that's super important in the age of AI is to think about your data as one of the most strategic assets, but you've got to bring that data contextually to the AI. In fact, people talk about it as context engineering, and I think it's a good term for it, because, in fact, a lot of what needs to be done, both at this engineering level and even, I would say, business decision making level, is to think about not the silos of the previous generation, but the contextual way you can feed intelligence with data.

To that end, for example, underneath Zoe's demo is perhaps the most important database in any company that uses Microsoft 365 today. If you think about it, the tacit knowledge inside any corporation, any organization, is the knowledge of people, their relationships with other people, their work artifacts, whether they're spreadsheets, documents, their conversations, in email, meetings, in Teams, to be able to relate all that, because that's where the tacit knowledge of an organization is. And it's in that database.

And so, fundamentally, what we have transformed that is into Work IQ. When we say Work IQ, it's the stateful part underneath M365 that now can be brought to bear in the context of any AI solution you want to build, not just Copilot. Any agent you build out there, you can now use Work IQ. Think about that.

The second thing we have done is, in fact, one of the other ubiquitous tools all over is Excel and Power BI. That means you've taken real, great pain over decades building these semantic models to do analysis. And now, with Fabric and Fabric IQ, you can take all of that data, along with all the operational data, and make it available to your AI.

And then the last one really makes it all come together. What we have done with Foundry IQ is said what people describe as retrieval augmented generation. One of the biggest things that we have done in the last couple of years as we build these AI solutions is to do retrieval of data.

And we have now built the next generation retrieval system, essentially. It does more than retrieval. It, in fact, plans and executes and reasons over that plan on how to retrieve relevant data from Work IQ, Fabric IQ, and any other blob storage, unstructured storage, or what have you, and bring it in the context of your AI application.

That's what this IQ layer is. In fact, it's probably one of the more salient, yet important pieces on how anyone who builds any AI solution, going forward, will be most successful. In fact, direct correlation between the evaluations or the performance of your models will be, how good is your data being brought, just because, in some sense, models are becoming, quite frankly, a commodity.

There is lots and lots of models out there which are very, very capable. The question is, how do you marshal that capability? I think it'll come down to having a great experience layer where people can discover the agents you're building, and having a fantastic capability around context engineering.

The next layer for us is really about, now that we have an experience layer, we have the data set up, let's build agents. And the beauty of building agents, quite frankly, is like building spreadsheets or writing a document. It's not more mystical than that. In fact, agent builder that's built right into Copilot is like saying, hey, doc new, or spreadsheet new.

And so, we start that. In fact, now we have an app builder. In fact, Bill's big thing always for ever since I joined Microsoft in the early '90s was, why the heck do I have a difference between a document, a website and an application? His concept was, there should just be one thing, and I should be able to transform.

And guess what? We finally have it, which is App Builder, you can go in and just specify the app you want to build, and it'll build it. It's kind of like a document. It's kind of like a website. It's kind of like an app. That's the power.

Then Copilot Studio is another way to think of it as build out of your agents. Again, you're using a natural language prompt. You can go build an agent. You can ground it in some data, and there you have it. You have an agent.

And so, we are building out Foundry, in fact, without any cliffs, to be able to then say, okay, we now have agents. Can I even orchestrate multiple agents and build full agentic systems? Microsoft Foundry allows you to do that as well, because the multiagent frameworks are there. All of the tools that you need in order to be able to build multi-agentic systems are there.

We also now, of course, have all these models. We have 11,000 models. Why do I need 11,000 models? In fact, even for me this week, it is a great reminder of why having this diversity of models is going to really make this world a better place.

In fact, we launched and we announced in Foundry as an open weight model, a model called GigaTIME. There was an article published in *Cell*, and quite frankly, it gave me goose bumps when I read about it and learned about it, because it's work Microsoft did with both the University of Washington, as well as Providence Hospitals.

And what it does is, if you think about the challenge of someone with cancer, today, when you try to understand whether immunotherapy will work, you have to model the immune system, essentially. And so, there's a test for that. The problem with that test is it's a complex test. It takes a long time and it costs a lot of money.

And so, what this particular model does is it does a simulation of the test. You can literally take a simple pathology of the tumor, literally get a microscope image of a tumor, and then use, essentially Gen AI to fill in. It will do the simulation of what is the basically, spatial proteomics of the immune system, and say, you now can have a test.

I was thinking about it. With that model, that means some entrepreneur in Mumbai is going to pick it up and say, you know what, I'm going to build an entire simulation lab in silico, and make it available to every city in Maharashtra and in this country, so every hospital can start doing what now today is done in the few tier one hospitals and costs a lot of money.

In some sense, when you're playing the game, sometimes you forget why you're playing it. This is why we are playing it. We are playing it so that we can actually use this innovation to make a real difference in the world. It's exciting to see how this is all coming together.

Now, we're bringing, as I said, all this together in this toolchain. This toolchain is going to become what you will use. In fact, GitHub is where you would start interfacing with all of this app stack. If you think of the new app server as Foundry, the tools are Visual Studio and VS Code and GitHub.

And talking about GitHub, it's really exciting to see the growth of GitHub in India. I mean, it's tremendous. By 2030, India is expected to be the number one country in terms of GitHub participation.

(Applause.)

It's tremendous to see, not only just the participation, but the ambition level of the projects coming out of India just continuing to improve and get better. And so, it's fantastic to see that.

Now, we are building out GitHub, by the way, as Agent HQ. What I mean by that is, you go to GitHub and you will have access to all the coding models that you could use in the context of your repo. And so, that's the idea. It's not about having one model, but you have your code repo. Just like Work IQ, think of this as code IQ. That's really it. Your code is your most important resource. You want that, and then you want many, many models over the years to keep working and creating more of that repo, and the repo features.

And so, that's what – and by the way, the form factors, if people say, oh, wow, will I only use autonomous agents? Will I use only an IDE with Agent Mode? Will I use a command line? Guess what? I use all three, all day long. I mean, that's it. The form factors of how I interface with AI are going to remain essentially, always some IDE, some CLI interface, and a new chat interface, if you will. And that, I think, is the way you will interact with it.

But to show you all of this in its full glory, I want to invite up on stage my colleague, Karan.

Karan, take it away.

(Applause.)

KARAN MV: Thanks, Satya. A few weeks ago, at GitHub Universe, we announced Agent HQ, which is an agentic platform that unites every agent, making them native to the GitHub flow you're already familiar with. Let me show you a demo of what that looks like.

All right, so I have a demo app here that's a storefront for some of the nostalgic toys which might have in common in many of our households. However, as you can see, the app doesn't really have any search or filter options, and I want to add that functionality.

I have my code repo here on GitHub. With the new Agents Panel, I can now kick off an agent tech task from anywhere on GitHub by just describing the task and delegating it to the Copilot Coding Agent. Copilot starts working on it in the background. I can assign a task to Copilot from anywhere I work, on github.com, in VS Code, in the CLI or even on my mobile.

All right, while working across multiple repositories and tasks, I have delegated many of these tasks to Copilot. And I want to know what's going on with it. Now, I can see all of my agent sessions in a single mission control view, both ongoing as well as completed once over here.

Not just that, I can also start off a task from right here while using custom agents that I have created or have been built by partners. These custom agents are domain experts, which extend the Copilot Coding Agent across all of your tools and workflows.

All right, so let's dive into the session that we just kicked off and see what Copilot is working on. I can see the progress of Copilot here on the left, and also the draft PR on the right.

I'm a developer, which means I am still in control, so I can even steer an ongoing session midway to have Copilot work in a slightly different direction. This will take a few minutes, so I let Copilot continue working on that in the background, but I want to show you one of the agent sessions I kicked off earlier for the same feature.

I can see it's completed. And then there's a summary as well, and, oh, nice screenshots as well. Thanks to the playwright MCP, I can also take a look at the files it has changed over here, and also ask it for any change, if I need, right over here.

All right, amidst all of this, if I want to have a chat with Copilot to ask it something or so, I can jump into Copilot Chat right here as well. Copilot provides me a choice of various different models from different providers, including GPT-5.2., which just launched last night and already available on GitHub Copilot. How fast it is.

All right, not just that, I can also bring in my own models from Microsoft Foundry. I have a few models that I've deployed on Microsoft Foundry, including a couple of them that I had finetuned myself. And I have brought in all of these models to Copilot by connecting them to my deployments in Foundry. Now, I can actually use all of these models directly within Copilot, be it on GitHub or VS Code.

Now, I was just doing something, and I finetuned one of the model, just for fun to see if it can converse, almost like a developer in Mumbai. I just asked it something simple, saying, what is Kubernetes, your pods, notes, (foreign language)? And this is coming, and it says, simple language with (foreign language). This is what containers are. This is what Kubernetes is. And it's so awesome, because this –

(Applause.)

This is coming from my own model that's deployed on Foundry, which I can use in Copilot. That's pretty awesome.

All right. Speaking of deployments and everything, I had also created a PR to implement a new product detail page sometime earlier. And I can see that Copilot has already reviewed my PR and also left a few comments on it. Let's take a look at some of the suggestions it has left, like this one here, which speaks about DB connection pool exhaustion, or this one right here, which speaks about pipe safety.

This is the new agentic Copilot code review, which combines LLM detection, tool calling and also deterministic scans that help you with smarter reviews. Not just that, I can actually ask Copilot to help me implement all of these suggestions in a single PR, which will, again, kick off an agentic task for me.

All right. We all know how important security is. Thankfully, I'm also a security admin for this account on GitHub. With the new integration between Microsoft Defender for Cloud and GitHub Advanced Security, I can now specifically filter for runtime risk issues in production that have been identified by Defender, all right, right here.

That's really nice, and a campaign to address all of these alerts was created. And oops, I see some of them in my repo as well, but not to worry, because Copilot can also help me fix these security issues. I can just select all of these, and then say, assign to Copilot. And then Copilot will go off and then create a PR to address all of these security issues.

All right, now we have multiple different agents working on different tasks in the background, but I want to write some code myself as well.

Here is the agent session for the search and filter function that we were looking at earlier, and I want to make some changes to it. I can continue iterating with this by opening this directly in VS Code, which will open up the session for me. I can either check out the branch or even apply the changes directly in my current branch, and I can continue working.

I will open up the familiar chat interface here in VS Code, and I'll switch to my Agent Mode, probably use one of the models I deployed for Foundry. And I will ask it to make one small change, which is, I also want an autocomplete feature for the search bar as well.

This can continue working in the background. And you would have seen that just like on github.com, I also have a mission control view in VS Code as well, where I can see all of my local agents, my cloud agents, my background agents, which is running as well.

All right, so I want to implement a new feature. The toys are great, but I want to make a bundle of it and almost gift them as a play box. I know many of you would love that as well. But I need to plan how to go about implementing it.

With the new Plan Agent in VS Code, I can ask Copilot to create an implementation plan and also keep iterating with it until I am confident and ready with it. I had created one of this plan earlier as well, which I'll just open up right over here.

You can see, it gives me a detailed plan for all the steps that needs to be accomplished. Whenever I'm ready, I can continue coding this, or I can hand it off to an agent as well to work on it in the background. I can hand it off to the CLI agent, and I will say, include the changes. That's fine. And this will start working on this agent in the background.

All right. Some of this will take a bit of time to complete. Meanwhile, when all of this task would have been completed, let's see how our app would look like. It will look like something like that, with all of the filters that are there, and also my search and auto complete feature, and I can go start playing my (break?) game.

(Applause.)

All right. In the past few minutes, we went from so many different ideas, planned them and also coded them using Copilot Coding Agent in the cloud, on github.com, in VS Code, in background, using CLI agents, had that code reviewed with the Copilot Code Review Agent, and also implemented the security fixes using Copilot Autofix Agent, all the while having ability to look at all of my agent sessions, steer them in the single mission control view on GitHub or even in VS Code. That is Agent HQ.

Thank you. Back to you, Satya.

(Applause.)

SATYA NADELLA: Thank you so much, Karan. And so, when you look at all of these tools that Karan showed, what does it do? It inspires you.

Recently, we had Thanksgiving in the United States, so we had some time off. And so, what does one do when you have time off? You build. (Laughter.) And so, I built an app.

And so, by the way, this is my regular laptop you are seeing. This is, in fact, my Azure subscription. And this is, in fact, my app that's deployed right now, I think, in Canada Central.

This is my repo. This is my GitHub repo. This is my code. In fact, my standard way I work is I have a Windows 365 setup so it travels with me everywhere. And then, of course, I use GitHub, and then I use Codespaces. I think it'll show up with the Codespace, so it's like turtles all the way down. I have my Codespace running on GitHub, running on Windows 365. That's my virtual environment that I travel with.

And of course, as Karan was showing, every morning, I get up and my favorite place to go is here, which is I go to my repo, and I go to my Copilot, and start assigning it tasks. Every day, there's some multiple agents I kick off like.

This morning, I said, please upgrade this entire app to GPT-5.2, And it's actually done. And when I'm done with my keynote, and I'm on a flight back, I'll get there on my flight Wi-Fi and be able to actually look at the PR and push it, if you will. In fact, I'll test it once with my Codespaces, and then I'll commit. But in fact, it shows you the power of how you can literally start really taking simple tasks to big updates and start powering them, powering through them.

What did I build? What I built is essentially a new deep research tool. One of my life's ambition is to figure out how to get a job in the Copilot Team. And so, I'm preparing for it. I'm desperately trying how to get competent enough to get hired. And so, therefore, I said, Okay, let me build a new, more sophisticated, deep research application.

And so, the thing that I did is, I said, look, there are all these models available. Why don't I create decision frameworks that work on top of the model? And one of those decision frameworks is Andrej Karpathy, who's the famous AI researcher, came up with this LLM Council, he called it.

The idea is you will have, like a selection committee, you will have multiple members of the committee. In this case, I can have GPT-5.1 or 5.2 today, Claude, Gemini, LLaMa, Grok, whatever, pick your favorite selection committee. Then you elect even a chairman. You can say, who is the chairman? Is it Claude or Gemini or GPT, or what have? You pick your chair, and then go ahead and issue a query and ask it to do stuff.

And then another one I did was DxO. This is a pretty cool decision framework. This comes out of health care. In fact, we first implemented it – in fact, we have implemented it in health care, where you have just different roles.

For example, one role is a lead researcher who does, let's think of it as, exhaustive breadth-first research. There's a critical reviewer, and by the way, you can assign different models to these different roles. I have a lead researcher that happens to be Claude, a critical reviewer that happens to be a GPT.

And what does the critical reviewer do? That particular role finds method problems, say, recency bias or period bias, whatever are the issues, which are cognitive biases. Then you can have domain expert, data analyst, so different roles.

And in fact, when we applied DxO in health care, it turned out that having multiple roles work together in a multiagent system outperformed any single model. It's intuitive. Why not have all the smart people working with you on a particular decision versus one? And so, that was DxO.

Then I implemented even another one called Ensemble. And the idea with Ensemble is you fire off parallel queries, anonymize the responses so there is no bias, even in selection, and then synthesize.

In fact, last night, I said, well, let's even implement a regular old debate. This is a new feature I added, and thank god it's short up here. And basically, think of it as chain of debate. In fact, I think of all of these as, instead of thinking of them as chain of thought, thinking of all of this is chain of debate. And I said, let's actually have a full on debate.

The bottom line here is, you can have different formats of even debates. You can have pros and cons, swat. You can have risk versus impact, counter factual versus devil's advocate. You can even pick your form framework. You can have critiques. Each person makes an argument, everybody else critiques. You can have multiple rounds of it. I was just – I had a lot of time, I guess, in my hotel room yesterday, but so I went on and added features like this.

I said, okay, let's have – oh, by the way, I even added, because I was spending so many tokens, I said, man, I might as well actually add cost control to my thing. And so, I said, do I need fast, lower cost or deep, higher cost? I'm picking balance here. I can pick the different research analyst, system architect, risk officer, three types of roles, then have multiple rounds. And so, you get the point.

Of course, you build all this. What does one good old South Asian do? Use it to select the best ever Indian test cricket team. (Laughter.) I think we need one.

And so, if I go to my history, you'll see it. I've been at it. I've been working it. In fact, I have a fantastic MLB lineup as well. But since we are here in Mumbai, I'll show you some of the stuff. I'll show you the council one. I'll show you what it came back with. In fact, let me increase the font here.

This is the chairman synthesis. It shows me obviously, Sunil opens with Sehwag, Rahul, Sachin, Virat. Man, it picked – this is the synthesis, and it shows VVS at six and seven, and then it goes on.

Here's the interesting thing. The chairman's report says that there was unanimous, essentially consensus. Gavaskar, Sehwag, Dravid, Tendulkar, Kohli, Kapil Dev, Ashwin, Bumrah, but the big debate, of course, was do you need an additional batsman or not? And thank god, they selected, at least being a Hyderabadi, I'm really proud that they made VVS as Lakshman in there. (Laughter.)

And here's the interesting thing. It weighted Phi-1 and Claude's inclusion. I mean, think about this. These two models bid for my man, Lakshman, because why, of his crisis management and made the point about why.

And then the other debate was, who is the captain? In fact, I think – oh, by the way, the Kumble versus Zaheer was also a very good one, which is it chose Anil versus Zaheer. Zaheer is going to be 12th man. It depends. I guess that's where I think there's more work to be done. I would pick Zaheer some days and sometimes, depending on where you're playing. Captaincy debate, it picked Kohli.

The key thing is, I see the annotations of the chain of debate, and it gives me insight. In fact, if I go back to show you one other case, let me go back to my history again. Let's see, let me go to the DxO view. I think this is interesting, because the DxO view is slightly different. It says, here is what the lead researcher did. Then it says, here is what the critical reviewer did.

But here's the cool thing. Look at the way it's found the bias, the error bias. The critical reviewer's job was to say, okay, when I'm selecting something, are you mixing up your stats. I mean, I always say, sometimes you could say, wow, the fitness levels of the modern teams are so great and wonderful, and what have you. But man, to watch Gavaskar go and bat without a helmet, or even, I don't know, Hazari go uncovered wicket in play in England, how do you really equate for that. That's the era bias.

And so, to me, to be able to have all these models debate that issue is what this is all about. Anyway, so you get the point, which is, you can use these decision frameworks as the new form of metacognition.

If you have all this abundance of models, the ability to literally do what I did, which is, just build your own multiagent system, that's the new commodity. I mean, the reality is literally, all of us are going to be doing this work, just like how we do spreadsheets or documents.

And if we have that power, then the question is, what are we going to do with it? And one of the things, I think, is exercise better judgment, better decision and in high stake situations, in health care, in financial services, in insurance, in supply chain. And that's the beauty of it, to be able to think about these as metacognition frameworks for us to apply better judgment in a world where we have all of this.

One of the other things I should mention is I built all this. I even made it into, in fact, Copilot. This is all in my Azure private tenant, and I then did a rep basically, an open API to it. And so, then I made it into a Copilot agent. And then I said, great, let me go deploy. And guess what? I was caught. (Laughter.)

In other words, the Microsoft DLP, data loss prevention, came and said, you can't deploy it because this has to be deployed inside the Microsoft tenant. You just can't point it to an external agent, because that's a massive x-fill problem.

And that is thanks to this other part, which I think is the runtime of agents we have built called Agent 365. And Agent 365 takes our identity management, our Defender, our things like Purview and data loss prevention, all of that, and brings it to the agent world. Just like what we have done for end user computing, if you will, we're now doing it for agents, because if people are going to build powerful agents like the one I built, you want full compliance. You want full visibility. You want full governance of it.

And that's what the Agent 365 does, because you can't just run around and say, I have thousands and hundreds of thousands of agents without that security and governance and visibility. And that's really what Agent 365 delivers.

The beauty of all this is there is significant momentum, significant ambition level, to translate this ultimately, into people doing unbelievable things.

This morning, I had a chance to meet with, for example, the folks from Adani Cement, who are really doing this invoice to payment and reduce something that took 20+ days to four hours.

I had a chance to meet the team from Aditya Birla, where they took their mobile app that's doing complex products, and said, okay, here's a conversational interface that just makes it simpler for them to be able to interact.

I had a chance to meet with the folks at LTIMindtree. They're doing a complete wall to wall transformation, using things from Copilot to Copilot Studio to a Foundry, to do every business process and re-look at it from an AI-first perspective.

Mankind, it's a beautiful use case of really putting that power of AI and the expertise in the hands of the people who are used – the medical reps who are visiting doctors in rural areas to be able to make sure that the medical rep is confident in answering the questions and then improving care outcomes.

Yes Bank has done a fantastic job of really, the invoice turnaround times. They're also doing a wall to wall transformation with 40+ plus use cases, but really seeing the ROI in each one of these examples.

Now, one of the chances I had, again, was to meet with the team that built this new agentic system called the Maha Crime OS. This was sponsored by the Chief Minister of Maha Rashtra, who is here with us.

(Applause.)

And it was really great to meet both the SP, as well as the people in the team, and in fact, the investigative team. And the pride they had in being able to use this system to deliver justice to a citizen in Nagpur, who was unfortunately a victim of a crime, but the ability for them to use this agentic system to speed up the time to justice was just fantastic to see. Let's go ahead and play the video.

(Video segment.) (Applause.)

SATYA NADELLA: It's a great example, quite frankly, of using the power of all this technology to help citizens of this country. And more importantly, it also is a great example of how quickly this technology is diffusing. I mean, if you thought about any previous era of tech and the rate at which it would come to impact in a context like this, versus where we are, and that, I think, is the opportunity in front of us. It's fantastic to see.

The last piece I want to talk about is powering all of this, of course, ultimately, is the infrastructure. You've got to run essentially, a token factory. And when it comes to token factory, one of the most important things is this, one simple equation, yet very, very important, which is performance per tokens is all about tokens per dollar or per rupee, per watt. You have to optimize.

And this, by the way, is going to be the most important thing. Whether any country, any community, any state that is able to really have the infrastructure that delivers this efficient frontier of the token factory is just going to have more abundance of tokens that can then be deployed to drive the productivity, which directly is going to be correlated to the GDP growth in the economy.

And so, to that end, we are all in. We are building out Azure as the world's computer. We are investing in building out Azure in India. We now have multiple regions. In fact, in Maharashtra, we have a central India region in Pune. We have the West India in Mumbai. We have South India in Chennai. We have a DC in partnership with Geo.

We are also very, very excited about our new South Central region that is going to come out in Hyderabad next year, which is going to be on the next big data center region, and we'll keep expanding. And to that end, I'm really – by the way, the region in Hyderabad is going to be 100% renewable. That's the other cool thing, is we are really focusing on sustainability as well.

(Applause.)

And to us, we are also building all of this with making sure that there's real sovereignty, because after all, at the end of the day, this is not just about bringing the capabilities, but it's also ensuring that there is real sovereignty controls.

And so, for that, we have a full set of options. You can use public cloud. You can use public cloud with sovereignty controls. That means you're the key management locally. You can use confidential computing in concert with it. That means when the data is being processed, even in use, at rest, in flight, it's encrypted, and the keys are with you. We also have the private cloud option. We also, obviously have a region that's operated by local partners like Geo or NIC and others.

And so, we have a full set of options. Think of it as a portfolio that you can then use that portfolio for the different workloads you have, and have the assurance of your own sovereignty and control over it.

We are also doing, in fact, sovereign data control. Copilot and all of the things we talked about are all being locally processed, so none of the data is going outside of the data boundary of India. And so, we're very excited about all of these options being available.

Now, one other point I would also make, since we talked about Crime OS, and what have you, is one of the biggest challenges for any online infrastructure is cyber and cyber risk and cybersecurity. And it's an intelligence game. And ultimately, whoever has the best intelligence is going to be able to defend you the best.

And so, Microsoft sees trillions of signals every day. We see every nation state and every cybercriminal everywhere in the world. And we are able to harness that to protect. In fact, in the demo Karan showed, it was very cool, which is he basically was able to take your production instance, have Defender check in real time, what are all the vulnerabilities, and then have you even go back and fix those vulnerabilities. And that's not a static list. That is a dynamic list that is being monitoring, essentially, what is the cyber activity and the cyberattack graph.

And so, to me, really thinking about sovereignty and cyber resilience are two sides of risk. When you think about even you managing your risk, you want to think about it, because if you get disconnected, you have a sovereign everything except it doesn't have the best intelligence. You're going to be very, very vulnerable.

And so, therefore, it's a much more nuanced discussion that needs real thought. And then, at the end of the day, it's not a one-time decision, it's a continuous monitoring of that portfolio of risk and opportunity.

To that end, we are very, very committed to this country. And therefore, one of the things that I was excited to have a chance is meet the Prime Minister earlier this week and even announce the biggest announcement.

(Applause.)

It's our biggest investment in Asia, \$17.5 billion, which builds on top of the \$3 billion we announced earlier in the year, in fact. And so, we are obviously very excited to bring all of this infrastructure to empower, ultimately, all the citizens of this country, all the developers of this country.

And so, that's where I want to close out with, which is to us, all this is about empowering every person and every organization in India to achieve more. To us, whether it is things like GigaTIME, which allow us to think about making what is personalized care or even precision oncology possible in every tier two city in this country, because an entrepreneur here will take an advance like that and parlay it into health outcomes that matter; or the Maha Crime OS that's going to deliver the justice to someone who was a victim of some crime in Nagpur.

To us, ultimately, that's why we do what we do. All of us working together to realize the opportunity ahead of us in this age of AI is An exciting time. I hope you all have high ambition for what you can do. We sure do, for you all.

Thank you all very, very much.

(Applause.)

END