

Keynote

Satya Nadella, Chairman & CEO, Microsoft

Sydney, Australia

Thursday, April 23, 2026

SATYA NADELLA: All right. Good morning. It's fantastic to be back in Sydney and in Australia. This morning, I got up and I tried to go down memory lane thanks to Copilot and try and get back to the first time I showed up in Sydney for one of the developer conferences. It was, I think, in 1993.

I remember distinctly getting through the airport with my luggable compact computers. The conference was all about we had introduced, I think, for the first time back then Visual BASIC for applications. That was to replace the macros in Excel 4, Excel 5 had come out and we had Visual BASIC. I distinctly remember, in fact, having a session this time around, there are a few more people for my session.

And then I remember going to one of the banks and having a big discussion on going from macros to Visual BASIC. Is it a good thing or what have you? But little did I know that whatever, 30 odd, 35 years nearly after that, the hottest thing in tech is again, guess what? Agents in Excel or command line interfaces for coding agents. I mean, it's fascinating that as the frontier moves, so does the arrival back to where you started to use that T.S. Eliot metaphor.

To me, I want to spend most of my conversation today and time today to talk a little bit about that new frontier, and how are we all collectively going to push that forward with essentially, what is a new input. The outcomes that matter still remain very much the same.

When we think about any organization, any institution, even individually, what we care about is the excellence in outcomes, which is how do we show up and change that engagement with customers or the employee experience, or the efficiency with which any given business process gets completed. Or even bending the curve in innovation, whether it's software development or drug discovery, how are we going to truly change that? And that's the frontier that we are trying to move.

And if you look back at it, throughout these transitions that we've had in technology, a lot of this has come down to us changing the way we work with technology. But more than the way we work with technology, how technology reshapes the work we do. That's been the story of the movement on the frontier.

And it starts, quite frankly, with that mindset. That curiosity, the posture that we need to have that when there's something new, you go all in. You learn it. That means you really need to think about your skills. You need the new tools. And perhaps one of the most interesting things and important things that we'll talk about in the AI era is codifying the outcome, perhaps even more robustly going forward into these evals.

If you think about AI, anyone who's doing applied AI starts by with this mantra called the product is the eval. And that hill climbing on evals, and that subtle, nuanced way all of us judge what great looks like good looks like inside of your organization for any outcome becomes the real currency through which we can then change the use of technology and have it reshape our organizational frontier.

And so, to that end, we are building this end-to-end tech stack. It's got three parts to it. And for the rest of my talk, I want to unpack each of these in some detail. And the through line across all of this is going to be two things: Intelligence and trust. Because ultimately, you really want the intelligence to show through across each layer of the stack and amplify what you're doing, move that frontier, move that outcome.

But you really need trust because it quickly can become a vicious cycle versus a virtuous cycle if you really don't have intelligence and trust. And so, we'll talk through how those all come together. So let's start with the top of the stack, which is around these high-value scenarios and high-value applications or agentic systems.

And for us, it starts with the way we're architecting Copilot. If you think about Copilot, whatever, three years ago or so, we started with mostly a chat-based response, request response type of pattern. And especially with GPT-4, it became a new way for us to discover knowledge inside the organization or on the web.

And then, subsequently, we had reasoning models come out that helped us not just have essentially new search, but we had reasoning plus search. Even the chat interface started to pretty drastically change with reasoning models. Then we started getting better at long-run agents that you could then tool calling.

And so, that started really the beginning of agents like Researcher and Deep Research and what have you. So the first thing we started doing was to say, okay, let's take even chat and that reasoning and make sure that you have access to all the models. So in fact, when I use Copilot today, I'm mostly doing auto, and auto is picking the best model for my intent.

And then I have a drop down with all the models available. And that is how the chat interface itself has evolved. You used to be in the beginning, you used to have a router, but now the router is becoming smarter and smarter in being able to do the intent classification, pick the right models for the intents.

Next, we've now taken that idea of having multiple models and put them even in the context of agents. So if you go to a Researcher, you now have these new modalities. For example, you have a modality called Critique. And I love that a lot because after all, how do we work with people? We critique each other and our ideas. That's the way you work with even models.

In fact, another mode that I like a lot is Council, which is essentially assigning a particular task or a research topic to a group. And then you get the best of all the models come together, and it's fantastic to see it.

If you have not used it to look at both the models, not only doing critique, but also finding biases in each model, and then in fact, insights from each model. So even when the research report comes back, you gain deeper insights because that you have multiple models at your disposal.

Now, the thing that I'm really excited about is this next big thing that we've done. And in fact, it's the today is the general availability of Agent Mode across Word, Excel, and PowerPoint.

(Cheers, applause.)

Yeah. It's kind of like 1993 again. Actually, in some sense, all of my life's work is right here. And the reason I say that is because if you think about what's happened is Excel became a super popular tool because it helped us intuitively get number sense. It had representational power in the canvas. The rows and columns have information content.

As soon as you see a spreadsheet, you kind of make sense of what that spreadsheet is. That's why the canvas has lasted this many decades, in fact. And so, now, think about what Agent Mode does. What Agent Mode does is it takes that information in the sheet, the positional or spatial information representation, and then mixes it with essentially these powerful reasoning models and the ability to take action, not just reason, but take action.

And so, when you put that together, it's magic, which is you're instructing the model to create a very sophisticated model, an Excel model. And then, when it creates the model, it's represented in this canvas, which you know how to even spatially inspect and what's driving what, what's next to what? And then you can then further reason about it.

And in fact, you can take that entire reasoning trace plus the sheet plus the model and share it. I mean, think about the cognitive work inside the company and how it's getting reshaped. It has familiarity of a spreadsheet, but everything around it has changed. And that I think is when you know have a powerful set of technology.

And so, we're really, really excited about this Agent Mode, and of course, we're not stopping there. The next thing that is also built into Copilot now is Cowork. The ability to now assign asynchronous tasks in a queue to Cowork, which by the way, could be creating a Word document or an Excel spreadsheet or what have you. And that then you can open up in Agent Mode and round trip this entire way of working.

It does feel like when we moved from faxes and interoffice memos to emails and spreadsheets and documents floating around the enterprise. It feels like that time again where the way we work, the artifacts we create, the workflows that we're involved in are fundamentally going through a sea change. And to give you a little bit of a flavor for this Copilot ecosystem at work, I wanted to introduce my colleague, Nirja. Nirja, take it over.

(Applause.)

NIRJA TRIVEDI: Thanks, Satya. Hello, everyone. I am so excited to be here with you all in Sydney, and I'm even more excited to show you the latest Microsoft 365 capabilities that Satya was just talking about, including Cowork, multimodality, and using agentic experiences in the apps you use every single day.

We have a lot to show you, so I'm going to get started, starting with Cowork. Cowork lets you build out assets in parallel using Work IQ data. Work IQ is more than just work content; it is the intelligence layer that powers Copilot and the apps that you use every single day. It makes sure that it understands your reasoning and understands your decisions, the workflows, the people that you talk to, so your work is tailored and personalized to you.

With Cowork, you also have access to a breadth of different prompts like organize my inbox, organize my week, prep for a meeting, and research a company. So really a range of different prompts. If I scroll down even more, you see that I have access to all the list of tasks that are running as well, the ones that have been completed, and that I can reference the assets being made like this Excel doc, as well as ones that need user input from me to continue.

And then finally, the ones that are in progress that I can slam my laptop shut and then come back and Cowork will still keep working. So for this demo, I'm going to use Cowork to help me build out the version two of my product launch. To do that, Cowork is going to look through some of my assets. It's going to analyze a recap email that was sent earlier and then build out four different assets in parallel.

The first one that it's going to build out is a launch budget, using a previous launch that I already had. It's going to take some assumptions from that launch as well. The second thing it's going to do is that it's going to create a new product pitch deck for me, using a technical specification doc that I have, and I also wanted to include that second version launch budget as well.

Lastly, I wanted to mobilize my team by sending a chat message to my B team with highlights of the news and schedule a meeting with two of my stakeholders this afternoon. Okay, that was quite a bit of different things I want Cowork to do. Awesome. It hasn't even started yet, but I wanted to do four different things, requiring many different skill sets.

To get us all on the same page, it's going to build me a launch budget. It's going to create a product pitch deck, send a message to my team, and schedule a meeting. All right, let's get started.

As Cowork gets to work, you'll see on the right side that I can track its progress as it continues, as well as all the tasks that it has as well. So the ones that I went over like the email recap as well as the budget, the pitch deck, and all of those. And then, on the chat pane, you can also see its chain of thought as it goes through each of the tasks. And then finally it tells me exactly the summary of all four tasks again that are running in parallel.

Now, this is sped up for the sake of time, but you could see here that it also includes the first action item. And this is the message, that Teams message that will be sent to my team. Again, I

always have access to review before anything is sent, be it a Teams message, an invite, an email, I can always review.

All right, this looks good to me.

I can also interrupt Cowork. Let's say, like we always do, we think of a last minute task as we're working. I can send that to Cowork. For example, I want to send an email with an exact summary with the budget recommendations. I realize that I've sent many messages, but I haven't really included that net new asset.

Cowork will work on sending that email, but it won't stop the previous list of tasks. Everything will continue moving forward. And then you can see here that that next step is already ready. It's found some time for me to meet with my stakeholders.

And this looks great to me, so I'm going to hit Create there. And I just have a few more assets to continue.

On the right side, you can see that my progress bar is continuing. I have four out of six items done, as well as different output folders as well.

Finally, the last task is ready. It has that email, again, with that launch budget attached to. And that looks good for me to send. That's the last thing. And then, in just a few moments, the task is complete.

You'll see here on the right side, I can verify all the tasks that are complete, all the folders that have the output of all the tasks that I have, finally, the input folders as well.

Let's dive in to the first asset and see how it looks.

All right, not bad. This is the pitch deck that was ingesting that tech spec doc as well as the launch budget. You could see here, there's an exec summary. It's really beautifully designed. It also includes my brand colors as well as the version two opportunity as well. And then finally, it includes a bit of that launch budget as well. This is a net new creation by Cowork.

All right, that looks good. Now that I have the first asset, let me run over to my Teams meeting to see if any of my stakeholders have responded to my chat.

And I see a couple of them have started to respond. Here's that chat message. I see a few of them are responding, and let's check this one out. One of my team members said, "This is great, but I want to know more about the improvements of V1 versus V.2." Makes sense, good question.

I could either go to Copilot, compare the two technical specification docs, copy and paste it back into teams. or I can use Teams for what it's used for, a canvas to work with both my collaborators as well as my agents. I'm going to do just that.

I'm going to go up and I'm going to add agents. And so, I want to add Copilot to this response. I'm going to add Copilot to this chat. I can @mention Copilot, as well as any of the agents that you saw up there, and I can ground it on the specification doc and click Send.

Within a few moments, Copilot will pull in the right responses, a lot richer than if I quickly gave the answer myself, or as if I had to go back and copy and paste it from chat. You can see here, it's a great response, and this looks good to me.

I did notice there was one last stakeholder that had some news, and this is actually really good news for me. She says there's about \$100k and under spend this quarter, and she wants to put that towards the Marketing budget. You don't see that a lot, \$100 extra K, so this is awesome.

I'm going to go back to Cowork, and I'm going to figure out what am I going to do with this extra budget? Here I am, back at Cowork with that first task that Cowork completed for me. That was the Excel doc.

I'm going to open that up, that Excel sheet, and right when I open it up, I want to call out a couple things.

The first thing is right when I open up the sheet, I see that it's automatically saved into my OneDrive, so I don't need to go and find it in the task if I need to.

The second thing that I want to call out is that it's also applying the security label as well. Here, I can see that it has the right security label that is using the underlying assets. Satya mentioned trust and security are part of everything that we do, and you see that here as well, using that underlying confidentiality label.

And now, I'm going to go and activate Copilot so we could use it in the agentic apps that we use. Here, I have access to multi-model functionality, I could choose whichever model I want for the job, including our latest GPT models, as well as our latest Claude models. And then I'm going to have Copilot help me figure out, what am I going to do with that extra \$100k.

As I click Send, Copilot is going to help me build out a dashboard with clear charts and visuals, and it's also going to show me clear recommendations about how to distribute across the marketing tactics as well.

And as it starts building out, I can see on the right side that I can build the dashboard, creating new sheets with recommendations and charts. And I can also see its chain of thought reasoning as it goes through this analysis.

And then finally, I get this beautiful image with pivot tables, Excel native charts, as well as a full in-depth analysis here, too. If I scroll down, I can see here a rich analysis, including key numbers. It tells me exactly what each dashboard contains and then finally, it tells me what the core strategy is as well.

If I go back to the grid, I can click on any single cell. And I have the ability to edit a formula or change any single formula as well, without breaking any of the grid because it's Excel native.

We went through quite a few different things today, four different things. We were able to build a completely new pitch deck with using Excel imagery. We were also able to build a completely new Excel sheet and then finally, mobilize my team as well, through a series of Teams messages as well as Outlook.

Now, I hope this gets you excited because you could see this is just one way that Copilot fits into your workflow, but I hope you can see how it fits into your workflow as well. Now, I hope this gets you excited to use Cowork right in the Copilot app and Copilot in the apps that you use every single day. Thank you.

(Applause.)

SATYA NADELLA: Thank you so much, Nirja. It's wonderful to see even the evolution of Copilot in just the last few years. I mean, in the first instance, it was a bit of a UI on a model, but now you can think of even the Copilot system as essentially, a cognitive loop or a harness that you have at your disposal for you to be able to have more agency, more ambition on what you want to create, how you want to, in some sense, initiate the trajectory of a task. And that, I think, is what you see in all of those features.

And then it's also becoming quickly multiplayer. You're not just single-player mode, but when you're in Teams, you're interacting with not just Copilot, but all the other agents as well. And that, I think, is what is going to probably unfold very rapidly throughout this calendar year.

Now, we are seeing fantastic momentum. The great news is, this is not about this rate of diffusion here happening slowly. It's, in some sense, happening all very fast, everywhere, at the same time.

In Australia, we've now got some fantastic momentum around Copilot usage across all of the top companies that are represented on this slide. I think 17 of the 20 top companies in the exchange here are using Copilot deeply. All the banks are using it, multiple, I think there's 57 government agencies using it. It's fantastic to see the momentum.

And I had a chance yesterday to meet with this team at Australian Post that is looking at Copilot and the adoption of Copilot and it's a complete bottoms-up way. They have all these Copilot champions and team captains who are really making sure that everyone is trained on it and is getting the best out of it.

But the thing that stood out for me is how much care they're taking in making accessibility one of the big beneficiaries of Copilot. In other words, they've used Copilot and its native technology, but it's really driving everyone with all the abilities to be able to fully participate and give them more agency inside of the enterprise, inside of a Teams conversation, in a meeting, inside of being able to use Copilot to really get things done in a much more seamless way. And it was fantastic to see that.

That was the first part of the stack. We go to the next stage, which is the AI platform or the agent platform.

Now, the best way to think about this is in a couple of different layers. There is models. There's context. There is agents. There's harness, and then there's security around it. You want to have access to all the models. You want to be able to have access to all of the rich context. You then want the agents to be built on top of that, using any model with all the context. And then you want the harness that knows how to orchestrate all of it. That's the core of it. And then, of course, you need the security and compliance around it.

I want to unpack a little bit of this platform, and for us, the organizing piece here is Foundry. And Foundry brings all of these things together, whether it's the agent and the multiagent runtimes, whether it's models, whether it's the context layer. And so, I want to walk through each one of these pieces in some detail.

The first is the models. Of course, you have the frontier models, but the reality is, we have 11,000 models in Foundry today. And this is the base material for you to be able to build your own agent and agentic systems.

In fact, all the frontier models, as I said, are there, and they're rapidly evolving. We ourselves, with MAI, have launched, even in the last, I think, month and a half, a fantastic transcription model, an image model in two varieties. One is a very high performance and low cost image model, a new voice model, again.

And so, it's fantastic to see they're really not just the big frontier models, but even the helper models we use all the time for very different modalities inside of our applications are all rapidly evolving. But I wanted to talk about this one model that struck a chord for me, because what is the innovation agenda around models? And how does one go about it and build this as a capability that truly empowers the next level of innovation?

To showcase that, let me just roll the video on GigaTIME.

(Video segment.)

SATYA NADELLA: What was said in the video makes a ton of sense that you can take a technology like this, take what is a complex test that costs a lot of money, takes a lot of time, and truly democratize it. At that point, about every tertiary hospital in every secondary city, now with just a regular slide pathology can get this test done essentially. It's a pretty big breakthrough.

But the question is how that came about. The way it came about is two research institutions that have had this knowledge. They obviously had all this data, and they were able to take that data, train a model, a generative model to essentially do simulation of a test using compute and share it with the world.

That to me is what's happening. This is not about one frontier model company doing things. It's about all of us as an ecosystem at the frontier have the ability to use these new techniques like generative AI, these new algorithms, this new compute infrastructure that's there and turn what is essentially the tacit knowledge we have the information we have the data edge we have.

And turn that not only to move the frontier inside the organization, but you can move the frontier for the society at large. And that, I think, is what's exciting about the true democratic force that's underlying all this.

The next thing that I want to talk about beyond the models is context. In some sense, context around the models is perhaps one of the core assets that's, quite frankly, distributed. Every organization has a unique context, and you want to be able to organize that context such that you can put it together with these models to both reason and take action.

And that's one of the top engineering considerations and architectural considerations. Work IQ. So Work IQ is that rich database underneath Microsoft 365. Everything just shown around Copilot was tapping into Work IQ, and that Work IQ is available headless through an MCP interface for any app developer to be able to build an agentic system around it. And think about this.

Up to now, I've always said that one of the most important databases inside of any organization is the database we don't talk about, which is the database that contains all of the people and their relationship to other people inside the organizations, the communications graph and data, the document, libraries, project files, all that is in Work IQ.

Similarly, another very important database when it comes to thinking about outcomes is what's in Fabric. If you think about all the various dashboards that have gotten created over the decade. And if you think of any dashboard is some outcome, some metric someone's measuring. So there's a rich semantic model that was created. That's again, super important context to be able to give these LLMs and models as you try to drive your workflows.

And by the way, Fabric now has the ability to reach out to any data not just in Azure, but everywhere. You can connect to AWS, you can connect to GCP, bring all of the enterprise data asset as context inside of this Fabric IQ. And then in fact, the outer loop of this can all be Foundry IQ.

What Foundry IQ does is essentially even not only gets you that document database. So if you have a lot of unstructured information, you can then vectorize it, make it a part of your context, but you can also use it as an outer planner to be able to plan across Work IQ, Fabric, and Foundry. That's how these things come together. So context and this IQ layer beyond the models becomes the other second important piece.

Now, the third thing I want to talk about when it comes to agent platform is the agent service itself. With every gen shift in platforms, you kind of need a runtime. And that's what this is. The new runtime for AI agents is this Foundry agent service. And really exciting announcements today up and down the stack.

We have now persistent memory. We have tools used. So, basically, models that then can be attached to knowledge that can then also have attached to tools. You can then, on top of it, build these multi-agent and agentic systems that you can orchestrate using any orchestrator you choose, and you can publish them with one click. And then, as I said, you have memory that's persistent across all of this, and they're long-living agents and they're stateful agents.

And so, the one thing that I'm most excited about is the hosted agent service that now goes into preview mode as well. So it's really exciting to see this because now what you have -- okay. All right.

(Cheers, applause.)

So I think, in fact, this is like that day we started the cloud journey. The first time we talked about containers and said, "Oh, there's a new type of virtual machine effectively for cloud native applications." So think of this as the birth of this new generation of AI-native applications with a hosted service.

Literally, you have a sandbox per session that can bootstrap itself fast. So it's instant on. So in fact, when you look at the service, you will instantly be able to bring up a sandbox. The sandboxes are isolated. You have persistent memory. That session state has persisted. And then, of course, like in any other service in the cloud, you can scale up and down elastically.

So this becomes, I think, a massive boost because one of the things that I've been talking to a lot of developers and they come and say, look, when can we have a hosted service for agents? Because agents are new software. People talk about what's happening to software. Guess what? All software is going agentic and you need to host these agents. You need to have them have their own computer, essentially. And that's what this hosted container sandbox is.

And then it needs to have state, so that's kind of where the persistence comes. You need to be able to scale them up and down. And so, we are very, very excited about this new Foundry service with hosted agents and what you all can do to scale the agentic systems going forward.

Now that you have the runtimes, you have the models, you have the context, you want to start building these applications. And that's where we have the richness of multiple different tools. You can start right inside a Copilot with something like App Builder here, hit a prompt, and here's a career path app that you just built with a single prompt inside of a chat session.

It's a full-stack app that gets generated. You can share it. Or you can go to Copilot Studio, and again, build an agent using a simple set of prompts. In this case, it's an onboarding app that got built or an onboarding agent that got built using Copilot Studio. And then, of course, you have the full richness of the professional developer tools with VS Code and GitHub and now all the agentic tools.

In fact, one of the most exciting things is right here in Australia, we now have in GitHub a 2.3 million strong community.

(Cheers, applause.)

Yeah. And it's growing. It's 26% growth year over year. It's great to see the community continuing to scale. GitHub is not just about the code repo, but it's also now becoming the control plane for all the agents. So we describe it as Agent HQ, so you have every coding agent that you use will be available to you right next to your repo.

So you have your code, you have your coding agents. And in fact, one of my favorite new features is not only do I have my source code, but I have my session traces as a first-class thing inside a GitHub.

Because after all, going forward, it becomes very important for us to not only have the code, but also the traces that generated the code because of the coding agents right with the code. And I think this will become in fact for true provenance of code going forward, that becomes a pretty important attribute.

And then, we're also exposing the tooling through all the form factors. As I said, CLI is back. Everybody loves it. And the reason why is not just because a command line is powerful, but it's command line when coupled with rich agentic models and natural language that have access to everything.

In fact, one of the coolest things that I found myself doing is you can literally go to your repo in GitHub CLI and then say you can connect to Work IQ through something like an MCP and say, "Hey, I was in a design meeting yesterday. We had a discussion on this topic. Can you make sure that my repo is in sync with the discussion?"

I mean, think about that. How crazy is it? With a natural language utterance, essentially, you're able to relate a design meeting that was mostly a transcription to source code all being done in a command line. That's the type of richness in that form factor that you can bring.

And then, of course, IDEs still matter. In fact, IDEs are becoming ADEs. So they're now evolving to basically now the sessions app inside of VS Code is basically an ADE. And sometimes when you go into an Agent Mode or the agentic output, it says, "Please open up in VS Code." And if you want to inspect any code, you want that profiler there, you want the debugger there.

And so, we have all the form factors that developers need and want and want to use as part of a workflow all available to them. And so, to show you all of this developer richness, I wanted to introduce Australia's own Damian to the stage. Damian, take it away.

(Applause.)

DAMIEN BRADY: Thanks, Satya. My name is Damien, as Satya mentioned. And I am one of the team at GitHub, and also one of the 2.3 million developers right here in Australia, from Brisbane.

(Applause.)

Yeah, thank you. (Laughter.) All right.

One of the projects I've been working on lately is Tailspin Transport. And it is a little tool. It's our EV company. And so, I'm going to be working on the website today. I'm going to show you what that looks like.

Let's have a look at the website. This is the Tailspin Transport website. We have our EV vehicles. We've got a vehicles page as well, very comprehensive FAQ, which I'll get to in a second. And I'm going to do some work today to add a search feature to the vehicles page.

And the reason I'm doing that now is because I've got an email from one of our product managers giving me some information that I've been waiting for about exactly how they want this to behave. I use Visual Studio Code, and I use Visual Studio, but lately I've been using the GitHub Copilot CLI, so that's what I'm going to use today.

Let's kick off the CLI. And the first thing I want to do when I get in is have a look at the model that I'm using. I have been playing with GPT-5.4, as you can see right down the bottom, but this is a relatively straightforward task, so I'm actually going to use a different model. And I'm going to use Claude Sonnet 4.6.

Now, you can see just how many models I have available to me, just with the one Copilot subscription, even Opus 4.7, which came out incredibly recently. We use 4.6, and then I'm going to ask it to gather some information together for me, so I know I'm doing exactly the right work.

Now, here's a prompt I prepared earlier. You don't want to see me type that. And what I'm doing is I'm asking Work IQ to find that email from Evie Smith, the PM, and pull that information in together with information from the GitHub issue and build a single set of requirements.

And so, let's have a quick look at that email. We can see that there's a requirement for a search bar with simple search and then Azure AI search, the category sidebar, the price range builder.

And as much as, as a developer, I would love Evie to have put this in our GitHub issue, which is right here, I realized that that's not where communication always happens in our organization. I could copy and paste this, but instead, I'm using MCP to go and reach those two surfaces to pull the information together. Work IQ allows me to reach into the M suite and pull that information out, and the GitHub MCP server also lets me reach into GitHub to pull that information out.

And what's really cool is it's worked out, it can do both of those things in parallel, which is something that I am incapable of doing. I am decidedly single threaded, unfortunately. As that's coming back, it's going to build this one set of requirements for me, after asking Work IQ and getting that information back.

And I could continue with that set of requirements and work synchronously here in the CLI or switch to the Visual Studio Code and keep working synchronously there.

Here we go. We're starting to get the information back. And if I wanted to do that, I could also control it remotely using this new Remote command in the CLI. Even though the work is happening on my machine, I could adjust it and make changes using my mobile or the Web.

But instead, I want to show you a new feature, a new experimental feature called the Rubber Duck agent. This will take a sec, so I'm going to go to one I've already done earlier. And all I'm doing here is asking, can you give me a rubber duck review of these requirements? And what the CLI is doing is it's kicking off the Rubber Duck agent here, and it's reviewing what work has just been done.

Now, the cool thing about this is I was using Claude Sonnet 4.6 but to do the review, it's going to choose a model from a different model family, so we get a real second opinion on this. And it's lucky we did. We found a few issues here. Range is ambiguous, type versus category, the fall back spec is too vague, and a few other things.

I would obviously want to clarify some of these for it before I started work. But this is a demo, so let's just pretend it's perfect.

I'm going to continue with this. I could do it locally, but instead, I'm going to delegate this to GitHub, and that will send this information to the cloud. And the Copilot Cloud agent will pick up that work and start working on that while I can move on to something else.

The thing I'm going to move on to is the other thing that we do as developers, which is a code review.

Now, I mentioned that FAQ page was really comprehensive. The downside of that is sometimes it's hard to scroll through all that information, so it'd be great if we had an agent that we could just ask questions of. And that's what my colleague has put together.

He's used Microsoft Foundry and created a custom agent. It's based on GPT-5.4 mini, which means that it's going to be responsive and relatively cheap. We've got a set of instructions here, and importantly, we also have a file search.

Now, this is one of many ways you can insert real information in, but he just uploaded a markdown document with our FAQ page. That means when we ask a question of the agent here, we get real answers, grounded in truth. They're grounded in that FAQ document rather than potentially hallucinating them.

But we can't just trust that that's going to work. He's also done some evaluation and used automatic evaluation to have a look at how successful that agent is across a number of different areas, so things like tool selection, output, utilization, even things like relevance and at-task adherence. We've really validated whether that agent is going to do the job we want it to do.

Okay, that's the agent. That looks great. Let's look at the actual pull request.

Now, because I've been doing other work, the first thing I did was assign this to GitHub Copilot. And Copilot Code Review has looked through this pull request, knowing the context of our entire repo, and it's made some suggestions. It's made quite a few suggestions, actually. We've got nine comments generated.

And I could scroll through and fix them individually if I wanted, or I could just ask Copilot to fix them all in one batch. And in fact, you can send a message to Copilot from any pull request, whether it was created by a human or by the cloud agent, and ask it to do some additional work.

In this case, that FAQ page involves a UI change, but I didn't see any screenshots in the pull request, which would have been really nice. I asked Copilot to use Playwright to take before and after screenshots so I can see the UI difference. And this is another MCP server we have access to.

And sure enough, there we go. We've got a before and an after with our fields to ask the questions.

Now, that wasn't enough for me. I asked for a couple more screenshots, and I noticed there's actually an issue here. We didn't connect to the Foundry correctly, so there's something wrong in the code here. And this is something I might not have noticed initially, if I was just doing a code review looking at that code.

Okay, I've shown you a few things today, so I just want to quickly recap.

We looked at using the GitHub Copilot CLI to consolidate information, using MCP servers like Work IQ and the GitHub MCP server to start work on that new search functionality. We looked at using Microsoft Foundry to build a custom agent grounded in truth so that we could really answer questions that the customers had. And then we looked at how Copilot Code Review can help us review a pull request, and how Copilot can even improve that pull request as we go.

And trust me, I've only just touched the surface of what we can do with Copilot to help enable our teams to complete more work faster.

Thanks, everyone. Back to you, Satya.

(Applause.)

SATYA NADELLA: I must say, we have come a long way in our naming as well, from Visual Basic to Rubber Duck agents. (Laughter.) I love it. It's pushing the frontier. (Laughter.)

And one of the other super important considerations when you build all these agents, what Damien showed is just pretty stunning, between Copilot studio plus what he showed, you can be assured of one thing, which is we're going to have lots and lots of agents inside of our organization, in our tenant, in our network.

And so, one of the key considerations is going to be security and compliance and observability. And that was the motivation for Agent 365, which is, how do we build a new control plane, just like we did for the previous era, but for the agentic era? Do we need an identity? We extended Entra. We needed security. We extended Defender. We needed data security. We extended Purview. And so, fundamentally, now you have in Agent 365, a platform and a control plane.

And by the way, they don't need to be hosted on Foundry or on Azure. They can be anywhere, again, on AWS, GCP, anywhere, on your private cloud, even, but you have this one control plane that you can use to be able to manage all of these agents and have real visibility into all of these agents, and set governance policies and enforce those governance policies.

Now, one other consideration that perhaps, in the months to come, and definitely in the years to come, will probably become the top consideration is in this age of AI, where you no longer just have a platform, but you have a platform that extends into a learning system, then the question for any software company, any new AI startup, or, for that matter, any bank, insurance, energy, telco, whatever, pick your favorite sector of the economy, is guess what? There's no limits to these learning systems. They can pick up what's novel and new and interesting to know about any business and just make it part of their base weights.

And so, this is a pretty new thing that's never happened before, because it's not like the last time you thought about a digital system, that digital system was a self-improving learning system. It was a system that you could use with your context and your workflows and make sure that that IP that you create on top of it is yours. And that is what you now, and all of us have to do, even in the AI era.

And to do that, that means the first thing is to really think about each of these layers of the agent platform, and make sure that you start with your evals. After all, you're hill climbing with any agent system you build on some eval and making sure that those evals are private. In fact, perhaps one of the most important IP of an organization or a software company that is doing something unique is going to be captured in some private benchmark.

In fact, all the public benchmarks are so saturated, they're more or less becoming noise. And so, therefore, the real game for anything is about hill climbing on benchmarks. And your private benchmarks you create, private evaluation sets you create are going to become the new IP.

The next thing you want to make sure is you decouple. You want to have access to multiple models. Rubber Duck is cool, or Critique is cool, but you want the harness, the context and the models to all be decoupled. It's tempting to couple them to get some gain on an eval, but then you're stuck with that one particular local maxima, but you're not going for the global maxima, where you want to be able to ride every frontier gain or your own model development.

You want that context that's there. You want to externalize as much of the expertise inside of your organization into skills, simple markdown files. In Cowork, one of the cool things is you start writing these skills, and you even compound the skills individually and at an organizational

level. It's externalizing the expertise of an organization. You want to keep that, again, very private.

You want to take all the tacit knowledge of the organization. In fact, the tacit knowledge of the organization is with the people and the way they work and the way they work with AI, those traces, when I talked about in GitHub, now in sessions, you have the traces and the trajectories of all of your agents. That's super important data, not just your source code, but that data is your data.

And you can train your own model. You can take a base model, which is a pretrained model, an open weight model, and train that with these traces, and then put it back in Copilot in Auto, and it will then route to your model.

I think that this idea and this consideration around sovereignty of a firm, because ultimately, what is the future of a firm, the future of a firm is its ability to continuously, monotonically get better at its knowledge creation. And it's able to capture its tacit knowledge that is not only inside of context or in skills, but even in the weights that you control as your own IP.

And so, that's what we're building the agent platform for. And you'll continue to see us even going into our developer conference in early part of June, talk about this in great detail, because we think that this is what, as I said, any start up, any software company, and, for that matter, any part of the economy that now needs to be thinking of all of these pretty strategically.

And speaking of that, in Australia, we have fantastic momentum in terms of people building these agentic systems. I had a chance yesterday to meet with many of them. AMO has taken Fabric and Foundry and built out a fantastic system to be able to improve the signal-to-noise ratio around alerts. Because how do you really have AI help them manage, in fact, some of the most critical infrastructure around the grid in a much more streamlined way and an effective way?

Had a chance to meet the team at the Commonwealth Bank, where they built an extended both their customer-facing bot with all the new power of these models, but also for their agents in their call centers, to be able to take the escalations and have the tools. The productivity gains that they're seeing is tremendous.

I had a chance to meet with even one student as well as a software company called Cogniti that's working with the University of Sydney in enabling, in fact, teachers. We know that a lot of the students are using these tools to learn, but teachers being able to create new course material and change the way they conduct their sessions and so on was fantastic to see.

And of course, I had a chance to meet with Cricket Australia. And look, I mean, it's always awesome. I think Australia has been collecting and Cricket Australia has been collecting stats from, I forget, 1896 or something like that. In fact, the last time I was here, somebody was telling me, "Hey, if you ever wanted to check the scorecards of Steve Smith when he was a 15-year-old, there's no problem. You have all that data."

And it's like, and I'm a big cricket stats buff. In fact, one of the queries I have, which I'm going to do this weekend is, and in fact, I was asking them whether they have basically an MCP server for all of their data. And it's all in Azure blob storage, so I must assume it is there. But one of my things that I want to check out is the conversion ratios.

If truly Steve Smith's under-15 scores are there, I want to see them next to this Vaibhav Sooryavanshi scores and see how good this guy is going to be when he grows up. And so, that type of ability to use data in a rich way and then translate it in their case to the outcome that they desire, which is fan engagement, is fantastic to see. So let's roll the video.

(Video segment.)

SATYA NADELLA: It's so cool. For a stats rich sport, this is just game changer. In fact, they were showing me how you're following the match, and it just keeps coming ball by ball with this unbelievable statistics. And one of the things that they had was these conversion ratios. I guess, I don't know. At least I'm obsessed about it, which is what's the greatest ever test team from different eras, and you can't compare.

One of the hard statistical problems is how do you compare people across generations? And it turns out there is causality and it's called conversion ratio. So if you really like it, at least that's one of the predictable things. And so, I'm going to have a lot of fun this weekend.

(Laughter.)

So the last layer of the tech stack I want to talk about is the cloud and AI token factory. So this is the rich infrastructure that you get to build to power all of these agent systems and applications going forward. And it comes down to getting this one super important formula right, which is performance is really about tokens per dollar per watt.

And you really want to be best in class at it as a hyperscaler and you want to be best in class at it as a country. And that I think is what's going to really create that abundance of tokens that then allows the rest of what we talked about to come alive. And so, we are building out Azure as the world's computer with all of the regions around the world.

And really, we have 70+ data center regions around the world with one fungible fleet. And of course, when it comes to Australia, we have regions in Sydney, in Melbourne, and Canberra, and it's growing. It's 100% renewable energy powered. We're very proud of it. We're very committed to it.

(Applause.)

Yeah. And one of the ways we continue to improve performance and power performance is by innovating on the systems level and even all the way to the silicon level. So we obviously partner with Nvidia and AMD, but we're thrilled about the progress MAIA is making. And MAIA 200 now has 30%+ percent over any leading ASIC out there in the market.

And we'll continue to innovate and co-innovate, in fact, between those agent runtimes as well as the model architecture, plus the AI system itself, because there are very efficient ways for us to be able to really drive that tokens per dollar per watt going forward, and we want to be able to have that ability.

We also are very focused on our sovereignty considerations. And so, this is where we are making sure that you have all of the controls. Not only do we have the regions in country, but we confirm to all the regulatory requirements. You have the ability to manage your own keys. You can encrypt data. You can use confidential computing such that the data in use is also encrypted. So there's tons of sovereignty control that is built in.

And the one other consideration I would also submit, is that you want to think about cyber resilience as well as sovereignty together, because you want the global intelligence of signals when it comes to your overall resilience, in addition to sovereignty. You don't want to sacrifice one for the other. And that's again, part of the design of what we have put together.

And with all of that, one of the things I'm absolutely thrilled to announce is our doubling down, our commitment. We are making our biggest commitment in Australia of Australian dollars of \$25 billion.

(Cheers, applause.)

So this will allow us to expand this ability to generate the tokens and provide the cloud infrastructure. I think one of the things that gets lost is we just think about the tokens. But when I talked about those hosted agents, hosted agents need classic course, classic storage. Same thing, every Cowork session needs storage and compute.

And so, one of the things that we're going to see is the explosion of both cloud and AI accelerators. And that's one of the things that we are really excited about making sure we continue to invest and bring all of that. We're also thrilled to expand our partnership with the Signals Directorate Cyber Shield, because again, going back to--

(Applause.)

Because one of the challenges we all collectively face is how do we protect ourselves, defend ourselves and our organizations against cyber threat actors? And so, that's, again, a signals game. That's an intelligence game. And we want to make sure that we are able to extend that cyber shield arrangement such that the most vulnerable, that is that small business or the consumer is protected when it comes to any cyber attack.

And lastly, we are really thrilled to double down on the skilling investments we're making. We're going to train three plus million Australians around AI skills. And again, there's going to be a real change in the way we work, the tools we use to work. And so, to be able to get ahead of all of the jobs and the economic opportunity that gets created, I think it's going to be very, very important for all of us, including Microsoft, to do the work, to bring the skilling and the upskilling that is required.

So I kind of want to close where I started, which is our mission always remains to empower every person, in this case in Australia, and every organization to be able to achieve more. And I want to leave you with, I think, a video that perhaps best captures what this mission means to us and what it means to us collectively to strive for. Thank you all very, very much. Have a good rest of the day.

(Video segment.)

VOG: Please welcome the President of Microsoft, Australia and New Zealand, Jane Livesey.

(Applause.)

JANE LIVESEY: Welcome to the AI Tour. It's wonderful to see so many creators and innovators, all who are helping shape Australia's next chapter here with us today. Thank you for taking the time.

This is my first AI Tour since taking on the leadership role at the end of last year. After three decades in business, I've been asked, why Microsoft. And for me, the answer has always been the same. It's two words, trust and safety.

much of our world is changing, from health care, education, business and communications. There couldn't be a more important time to be Microsoft. We have the privilege and the obligation to help Australia navigate this change with trust and safety embedded in everything we do.

You just heard Satya outline Microsoft's vision for an AI era, where intelligence moves from experimentation into everyday work embedded in the world's most trusted cloud and productivity platforms.

AI is no longer abstract. It's here, and it's changing how organizations operate. Across every sector, we're seeing the appetite to adopt AI early and responsibly. It's not just about doing things faster. It's about unleashing human ambition, and that's what we call frontier transformation.

It's no longer about pilots or proof of concepts. It's about empowering people, reinventing business process and unlocking the value of data. And this moment isn't just about technology, it's about shared prosperity.

There couldn't be a more important time to unlock the value that AI brings to all of Australia. Through our local ecosystem, Microsoft contributes more than \$36 billion to Australia's GDP, and we support more than 150,000 jobs nationwide. But the opportunity ahead of us is even more enormous.

But none of this will happen on its own. It takes leadership, leadership that is bold and curious and people-centered, and partnership across industry, technology and across all of Australia to

ensure that we adopt AI in a way that is safe, inclusive and benefits all Australians. And that's why this conversation matters.

I first want to thank our partners and our sponsors for bringing these moments to life.

(Applause.)

There are two organizations that are at the heart of Australia's national fabric, Telstra and Westpac, and they're putting this into practice. They're not just talking about what's possible, they're doing it.

Please join me in welcoming Vickey Brady, CEO of Telstra, and Anthony Miller, CEO of Westpac, to talk about how AI is transforming their organizations, their people and Australia's future.

Thank you. Let's get started.

(Applause.)

I couldn't be more excited to have you both here for this conversation today. I've been looking forward to it as part of the AI Tour.

Both Telstra and Westpac are operating in an environment that's rapidly changing. That's something that we're all experiencing. How are you navigating this? Maybe, Anthony, I'll start with you.

ANTHONY MILLER: Certainly, it's one of the most challenging and invigorating environments for a long time. The way we think about particularly the AI opportunity, is that it's just this fantastic way in which people can do their job better. They can do their job faster. They can do it more safely. They can do it more consistently.

And so, the way we think about AI is, how do we get everybody set up to use it, and how are we investing in making sure they can use it and know how to use it and continue to improve how they use it?

Funnily enough, from my perspective, when we think about AI, it's a technology, yes, but really, it's a people issue. It's a people challenge. And so, how well are we investing in our people? How well are we training them? How well are we leading them? How well are we inspiring them and motivating them to use this tool to improve how they do things, day in, day out? That's the way we're thinking about it at Westpac.

JANE LIVESEY: Vickey?

VICKEY BRADY: Yes, so many parallels. I guess I've been trying to think about it in three ways. First off, as a telco, lots of increasing demands, expectation, just demand on our network. To continue to be a leader in connectivity, we've got to be a leader in how we're embedding and

applying AI inside the business. And that's a whole of business transformation, the way we engage with customers, the way we empower our teams, even the way we run and operate our network.

The thing on my mind in that part is always, are we going fast enough? I mean, the world is just changing so fast. As Anthony just said, it's sort of unprecedented. I know that gets used a lot, and I just think to be competitive, we've got to be moving at a scale and a pace.

The second way is very much the skills. Are we giving sufficient opportunity to our teams to grow those skills? We use a data and AI academy. We've just put in specific learning pathways for jobs. We're working with Microsoft at the moment on a pilot for an AI learning agent, which would be great.

And then the final way we think about it is really the national conversation around how are the AI benefits going to be shared across Australia, because AI is really a whole of society transition. It's not just affecting business, it affecting the way we work. It affects the economy. We're thinking about things that are important around trust, around inclusion, around skills. And we do think about some of those groups that are going to need more support, small businesses, regional communities, people facing barriers to digital inclusion.

And again, Jane will announce something today around small business, trying to help small businesses get greater access. But I think businesses need to be in that conversation alongside government and industry and be helping contribute to those solutions.

JANE LIVESEY: If we come back into your own businesses, you've both been quite ambitious in terms of your AI transformations. Can you share maybe a little bit about what you've learned and some of the outcomes you've achieved?

VICKEY BRADY: Yeah, absolutely. I might start with challenges and learnings, lots of those. Top of the list would be just how important getting the foundational tech and data is.

Just as an example, a few years ago, we had 80 platforms where our data was held. Our target is to get to three. We're at 20 right now, so making progress, but you've got to be relentless, I think, relentless on removing complexity and getting much simpler. We had more than 400 software partners. We consolidated down to two to move faster, and then it is skills and access.

And we've got a goal this year to get out that 85% of our teams are using AI tools at least weekly, and we're at 75% now. That's been important foundationally.

If I think about outcomes, I mean, everyone will talk about getting better customer outcomes alongside being more efficient. But if I drill down and just use one example, last November on our website, we launched a generative AI-powered assistant. First time, we had Gen AI in that tool, and we, in March, took that into our app. On our website, it does pretty simple stuff, be able to check plans, activate a SIM. In the app, it can do more sophisticated things, because we know who it is, which customer. It can make changes to plans, accounts, trouble shoot.

Already, of all of that volume coming in on the website and our app, it is dealing with 30% of that without our teams needing to get involved. And we continue to build that capability. Of course, it hands off where it's particularly sensitive or there might be indicators of vulnerable circumstances, but it's allowing our teams to spend more time with customers on those more complex issues. That's one example.

ANTHONY MILLER: I mean, Vickey, I agree with almost everything you said there. That's the learnings we've had.

I think the challenge or the learning that I've really found interesting with this tool, this technology, is the challenge that people need to be open minded. I run, and it is a very regulated industry, banking as appropriately, banking license is a real privilege, and we need to meet a certain standard. And so, therefore, there's a real culture of, that's the way we do it, we must do it this way.

And so, the tool asks you and gives you an opportunity to think about doing things differently. And so, therefore, it's a very big challenge for us to sort of create an environment where people feel they can step back and go, is there a different way that I can do my job? Is there a different way that we can deliver this service or product to the customer, and do it in a way that meets the standards rightly expected of us?

And that is not a skill that people naturally challenge to acquire as they progress in an organization. Equally, if someone has a better idea than you, or a better way of doing things is not a way people have classically been rewarded, to acknowledge someone else has a better idea. And so, therefore, this challenge that sits in front of us is that it asks us to think about things differently and be more open minded.

And so, that's the learning we've taken. And hence my observation around it being, perhaps not withstanding against the technology, it's one of the great people issues or opportunities of all time. We need to ask people to think about things differently.

And thankfully, Jane, to your organization, with the Copilot, we've given every single member of the organization the tool. Every single member of the company, 35,000 employees, is challenged, encouraged and motivated to try and experiment and learn. And then layering on top of that, we have more ambition, more goals and more opportunities as we go forward.

JANE LIVESEY: Amazing. I think as leaders, we're all operating in the heart of Australia at the moment. What are some of the things that you think we should be thinking about and your hopes as leaders for Australia?

VICKEY BRADY: Oh, it's a big question, isn't it? I mean, the first thing I'd say is I think, of course, we've got to consider the risks and challenges, but I really hope we grab the opportunity, because I just see it as to remain competitive, to remain relevant as companies, but as a country, grabbing that opportunity right now and moving fast enough.

And I think for me, that means choices we make today as businesses, as a country, around some of those foundational pieces about investing in skills in Australia, building trust through the risk and governance frameworks, inclusion so people can share in those benefits and getting the digital infrastructure right, so we unlock those productivity benefits. I just think it's a huge opportunity, and those things to position us well are necessary and they're urgent.

JANE LIVESEY: Anthony?

ANTHONY MILLER: I think I agree. I mean, the real challenge, or what really want to see is let's have some ambition. The current discourse publicly is so negative. It's so informed by what's not going well, or what is a challenge, as opposed to there are opportunities this country has today, which are extraordinary. And I think this technology is a contributor to what we should be aiming for, which is to really take Australia forward.

And so, without wanting to be too dramatic about it, I mean, one of the challenges that faces the country is lifting productivity. And classically, to grow productivity, it's about investment in things such as technology and innovation. And normally, that technology and innovation, it only drives productivity when it's diffused right across the community and right across society. And so, I think electricity, it's only when everyone's using it that we truly realized its benefits.

And funnily enough, this technology has so quickly diffused and so quickly being adopted and used by everyone that it sets us up quicker than we've ever had the opportunity to. How are we embracing it? And so, we need an ambition for the country. We need a public discourse that's focused on, how do we improve the country using this?

And it's not the panacea for everything we need to do, but gee, it's got a big role to play. And I think things like regulators, governments, businesses have a real obligation to encourage and motivate and invest in and challenge people to use it to improve how they do their job, deliver better outcomes for their customers.

And I think that's that catalyst will then drive some of what I think we can realize from it. And as I say, I just wish we were much more positive and ambitious around it.

JANE LIVESEY: Amazing. Vickey, Anthony, I can't thank you enough for your partnership and for sharing your insights today. We're really excited about the opportunity that AI brings to all Australians.

Thank you all for joining us today, and we hope you enjoy the rest of the AI Tour.

Thank you.

(Applause.)

END