

05212024 Build Satya Nadella

Microsoft Build 2024
Satya Nadella, Chairman and CEO, Microsoft
Tuesday, May 21, 2024

SATYA NADELLA: Good morning. It's fantastic to be back here at Microsoft Build. Welcome to everyone here and joining us on the web. Developer conferences are always the most exciting and the most fun when there are these fundamental changes that you can sense in the air.

I've marked all my adult life by coming to PDCs and Builds for the last three decades. I still remember distinctly the first time Win32, which was I guess '91, .NET, Azure, right? These are moments that I've marked my life with, and it just feels like, yet again, we are at a moment like that. It's just that the scale and the scope is so much more deeper and broader this time around. Every layer of the tech stack is changing. Everything from the power draw and the cooling layer of data centers to the NPUs at the edge, it is all being shaped by these new workloads, right, this distributed, synchronous data power and workloads, reshaping every layer of the tech stack.

But if you think about, even going all the way back to the beginning of modern computing, let's say 70 years ago, there have been two real dreams we've had. The first is can computers understand us instead of us having to understand computers? And second, in a world where we have these this ever increasing of people, places and things, right, and so as you digitize more artifacts on peoples, places and things, and you have more information, can computers help us reason, plan and act more effectively on all that information?

Those are the two dreams that we've had for the last 70-plus years, and here we are. I think that we've had real breakthroughs on both fronts. The core underlying force is one of the questions I always ask myself, which is like, "OK, this is great, and this is maybe like the golden age of systems, but what's really driving it?"

I always come back to these scaling laws, and so just like Moore's Law helped to drive the information revolution, the scaling laws of DNNs are really – along with the model architecture and the interesting ways to use data and generate data, it's really driving this intelligence revolution. You could say Moore's Law was probably more stable in the sense that it was scaling at maybe 15 months, 18 months. We now have these things that are scaling every six months or doubling every six months.

Effectively, what we have, though, with these scaling laws is a new natural user interface that's multi-modal. That means supports tech, speech, images and videos, and it has input and output. We have memory that retains important context, recalls both our personal knowledge and data across our apps and devices. We have new reasoning and planning capabilities that help us understand very complex context and complete complex tasks while reducing the cognitive load on us.

But what stands out for me, as I look back at this past year, is how you all, as developers, have taken all of these capabilities and applied them, quite frankly, to change the world around us.

I'll always remember this moment in January 2023 when I met a rural Indian farmer who was able to reason over some government farm subsidies that he had heard about on television using GPT-3.5 and his voice. It was remarkable, right?

For me, it just brought home the power of all of this because the frontier model developed in the West Coast of the United States just a few months earlier was used by a developer in India to directly improve the life of a rural Indian farmer. The rate of diffusion is unlike anything I've seen in my professional career, and it's just increasing.

In fact, earlier this month I was in Southeast Asia. I was in Thailand, where I met a developer and I was having a great roundtable, and he was talking to me about how he's using Phi-3 and GPT-4, and he was using Phi-3 to just optimize all of the things that he was doing with RAG. I mean, this is crazy. I mean, this is unbelievable. It had just launched a few weeks earlier, and I was there in Thailand, in Bangkok, listening to a developer talk about this technology as a real expert on it.

So it's just great to see the democratization force which we love to talk about, but to witness it has just been something. And this is, quite frankly, is the impact and why we are in this industry. I would say it's what gives us the deep meaning in our work.

I want to start, though, with a very big thank you to every one of you how has really gone about bringing about this impact to the world. Thank you all so very much.

When I think about the progress we have made since the last time we were here at Build, we've built really three platforms. The first is Microsoft Copilot, which is your everyday companion. It puts knowledge and expertise at your fingertips and helps you act on it. And we built the Copilot stack so that you can build your AI applications and solutions and experiences. And just yesterday we introduced a new category of Copilot+ PCs, the fastest AI-first PCs ever built.

All three of these things are exciting platforms, but I want to start with Copilot+ PCs. We are exposing AI as a first class namespace for Windows. This week we are introducing the Windows Copilot Runtime to make Windows the best platform for you to be able to build your AI applications.

What Win32 was to graphical user interface we believe the Windows Copilot Runtime will be for AI. It starts with our Windows Copilot Library. A collection of these ready-to-use local APIs that help you integrate into your new experiences all of the AI capabilities that we shared yesterday.

Now, this includes no-code integration for Studio Effects, with things like creative filters, teleprompter, voice focus and much more, but of course, if you want to access these models themselves, you can directly call them through APIs. We have 40-plus models available out of the box, including Phi Silica, the newest member of our small language family models, which we specifically designed to run locally on your inputs on Copilot+ PCs, bringing that lightning fast local inference to the device.

The other thing is that the Copilot Library also makes it easy for you to incorporate RAG inside of your applications on device data. It gives you the right tools to build a vector store within your app. It enables you to do that semantic search that you saw with Recall, but now, in your own application, you can construct these prompts using local data for RAG applications.

Today, I'm so thrilled to announce, as well, that we will be natively supporting PyTorch and new WebNN framework through Windows through XML. Native PyTorch support means thousands of OSS models will just work out of the box on Windows, making it easy for you to get started.

In fact, with WebNN, web developers finally have a web native machine learning framework that gives them direct access to both GPUs and NPUs. In fact, last night I was playing with it, turning it on in Edge and seeing the WebNN sample code running. It's just so cool to see it now using the NPUs. Both PyTorch and WebNN are available in developer preview today.

Let's take a look.

(Video segment.)

These are just some of the many announcements today. We are introducing more than 50-plus new products and partnerships to create new opportunities for you. We've always been a platform company, and our goal is to build the most complete end-to-end stack, from infrastructure, to data, to the application extensibility so that you can apply the power of this technology to build your own applications.

And so today I want to highlight our top news for this event across every layer of this Copilot stack. So let's dive right in with infrastructure. We have the most complete, scalable AI infrastructure that meets your needs in this AI era. We're building Azure as the world's computer. We have the most comprehensive global infrastructure with more than 60-plus data center regions, more than any other cloud provider.

Over the past year, we've expanded our data center regions and AI capacity from Japan to Mexico, from Spain to Wisconsin. We're making a best-in-class AI infrastructure available everywhere, and we're doing this with a focus on delivering cloud services sustainability. In fact, we're on track to meet our goal to have our data centers powered by 100% renewable energy by next year.

We are optimizing power and efficiency across every layer of the stack, from the data center to the network. Our latest data center designs are purpose built for these AI workloads so that we can effectively and responsibly use every megawatt of power to drive down the cost of AI and the power draw. We are incorporating advanced data center cooling techniques to fit the thermal profile of the workloads and match it to the environment in the location where it operates.

At the silicon layer, we are dynamically able to map workloads to the best accelerated AI hardware so that we have the best performance, and our custom IO hardware and server designs allow us to provide dramatically faster networking, remote storage and local storage throughput. This end-to-end approach is really helping us get to the unprecedented scale.

In fact, last November, we announced the most powerful AI simple computer in the cloud for training, using just actually a very small fraction of our cloud infrastructure. And over the past six months, we've added 30 times that supercomputing power to Azure. It's crazy to see the scale.

And of course, we're not just scaling training our fleets. We're scaling our inference fleet around the world, quadrupling the number of countries where Azure AI services are available today, and it's great to see that.

At the heart of our AI infrastructure are the world's most advanced AI accelerators. We offer the most complete selection of AI accelerators, including from Nvidia and AMD, as well as our own Azure Maia, all dynamically optimized for the workloads. This means whether you're using Microsoft Copilot or building your own Copilot apps, we ensure that you get the best accelerator performance at the best cost.

For example, you will see this in what has happened with GPT-4.x It's 12x cheaper and 6x faster since its launch, and that's the type of progress you will continue to see as we evolve the system architecture. It all starts with this very deep, deep partnership with Nvidia, which spans the entirety of the Copilot stack, across both all of their hardware innovation as well as the system software innovation. Together, we offer Azure Confidential Computing on GPUs to really help you protect sensitive data around the AI models, end to end.

In fact, we are bringing the latest H200s to Azure later this year, and we'll be among the first cloud providers to offer Nvidia's Blackwell GPU V100s as well as GB 200 configurations. And we are continuing to work with them to train and optimize both large language models like GPT-4o, as well as small language models like the Phi-3 family.

Now, beyond the hardware, we are bringing Nvidia's key enterprise platform offerings to our cloud like the Omniverse Cloud and DGX Cloud to Azure, with deep integration with the broader Microsoft Cloud.

For example, Nvidia recently announced that their DGX Cloud integrates natively with Microsoft Fabric. That means you can train those models using DGX Cloud with the full access to Fabric data, and Omniverse APIs will be available first on Azure for developers to build their industrial AI solutions. We are also working with Nvidia, NIM industry-specific developer services and making them fantastic on Azure. And so a lot of exciting work with Nvidia.

Now, coming to AMD, I am really excited to share that we are the first cloud to deliver the general availability of VMs AMD ND MI300X accelerator. It's a big milestone for both AMD and Microsoft. We've been working at it for a while, and it's great to see that today, as we speak, it offers the best price performance on GPT-4 inference. And we will continue to move forward with Azure Maia.

In fact, our first cluster are live, and soon, if you are using Copilot or one of the Azure OpenAI services, some of your prompts will be served using Maia hardware.

Now, beyond AI, our end-to-end systems optimization also makes cloud native apps and the development of cloud native apps better. Six months ago is when we announced our first general purpose ARM-based compute processor, Microsoft Cobalt. And today, I am really excited to announce the public preview of Cobalt-based VMs.

Cobalt is being used for video processing and permissions management in Microsoft 365, helping power billions of conversations on services like Microsoft Teams already. We are delivering that same ARM-based performance and efficiencies to many customers, in fact, including Elastic and MongoDB. Siemens, Snowflake and Teradata.

In our most recent benchmark data and tests, our Cobalt 100 VMs delivered up to 40% better performance than any other generally available ARM based VM. We are very, very excited about Cobalt getting into the market.

Now, let's move up the stack to the foundation models. With Azure AI, we offer the broadest selection of frontier and open-source models, including LLMs and SLMs so you can choose the model that makes the most sense for your unique needs and your application needs.

In fact, more than 50,000 organizations use Azure AI today, and that is great momentum. It all starts, though, with our most strategic and most important partnership with OpenAI. Just last week, OpenAI announced GPT-4.0, their latest multimodal model, which was trained on Azure. It's an absolute breakthrough.

It has text, audio, image and video as input and output. It can respond and just have a human-like conversation that's fast and fluid. It can even be interrupted mid-sentence. PGPT-4o is also the top performing model on benchmarks across a variety of modalities, and it's always going to get you to the best answer.

It has state-of-the-art performance and understanding the information you actually provide in your prompt, because that's really what matters. What OpenAI I demoed last week, I mean, that was just pretty magical stuff. And yesterday we showed how Copilot will leverage PGPT-4o.

In the video I'm going to play, notice that you can share your screen or session as a prompt in Copilot and have Copilot help you with whatever it is that you are doing.

Let's role the video.

(Video segment begins.)

RYAN: Hey, Copilot how's it going?

COPILOT: Hey, I'm doing great. And it looks like you're about to dive into some Minecraft. Ready to build, explorer and maybe dive a few mobs?

RYAN: Yeah, I'm trying to figure out how to make a sword. My son plays this game all the time and I have no idea what I'm doing.

COPILOT: No worries. I'm here to help you impress your son. I see you've got a crafting table in front of you, which is a great start. To craft a sword, you'll need some materials. Can you open your inventory? Just press E on your keyboard.

RYAN: Yeah, do I have what I need?

COPILOT: Let's see. You've got some sticks. Perfect for the swords handle. But it looks like you're missing the material for the blade. You can use wood, stone, iron, gold or diamond. Let's start by gathering some wood or stone. Shall we head out and find some resources?

RYAN: Yeah, that sounds great. Oh my God, what's that?

COPILOT: Oh no, it's a zombie. Run! You need to get away from it fast. Either sprint away or quickly build a tower block straight up to get out of its reach. If you can find a shelter or dig into the side of a hill, do it now. Just make sure it can't reach you.

RYAN: All right, I think I might be safe now.

COPILOT: Whew, that was a close one. Great job finding shelter.

(End of demo experience.)

SATYA NADELLA: It's really cool to see that. The magic here is that Copilot is your personal assistant that's helping you with any task, helping you stay, quite frankly, in the flow in the moment, whether it's a game, a document or a line of code. As OpenAI innovates, our promise is that we will bring all that innovation to Azure too.

In fact, the same day that OpenAI announced PGPT-4o, we made the model available for testing on Azure OpenAI service. And today, I'm excited to say that it's generally available on Azure AI.

What this means, though, is that we can now have these groundbreaking apps that all of you can build using this capability. One of the coolest things is that now, any app or any website, can essentially be turned into a full, multi-modal, full duplex conversational canvas.

Let's take a look.

(Start video segment.)

SETH: All right, let me ask the agent here.

AGENT: Hey, Seth. How's it going? What can I help you with today?

SETH: I am fantastic. I'm sorry I am in such a hurry. I'm going to go on an overnight camping hike. Where we're camping, it's going to be dark and cold. Should I use these shoes? It's the mountains, and it's probably going to be cold.

AGENT: Hey. No worries. Let me take a look. Ah. Those are the Treksta hiking sandals. While they're fantastic for summer hikes and keeping your feet cool, they might not be the best choice for cold mountains –

SETH: OK, I'm sorry to interrupt. Can you pick the best thing that I should use? And can you just add it to my cart? Just one product.

AGENT: No worries. Let's go with the Trek-ready hiking boots. They're durable, provide great ankle support, and most importantly, they'll keep your feet warm in the cold mountains. I'll add them to your cart now. Ready for checkout?

SETH: Oh, I see that amazing, muchísimas gracias (continuing in Spanish.)

AGENT: De Nada, Seth (continuing in Spanish).

(End of video segment.)

SATYA NADELLA: Seth was in a hurry, but there you go, PGPT-4o kept up with it. It's incredible, and so I really want to thank the OpenAI team for their partnership and really the responsible approach to innovation, helping our industry move forward.

In fact, Sam will be here, joining Kevin in a little bit to talk a lot more about what's coming, because that's the exciting stuff, how do you all sample what comes next? We are also bringing lots and lots of other models in as well, from Cohere, Databricks, Deci, Meta, Mistral and Snowflake, all to Azure AI.

We want to support the broadest set of models from every country, every language. I'm excited to announce, in fact, we're bringing models from Cohere, G42, NTT Data, Nixla, as well as many more, as models as services, because that's the way you can easily get to manage AI models. And we all love open source too.

In fact, two years ago at Build, we were the first to partner with Hugging Face, making it simple for you to access the leading open-source library with state-of-the art language models via Azure AI. And today, I'm really excited to announce that we're expanding our partnership, bringing more models from Hugging Face with text generation inference and with text embedding inference directly into Azure AI Studio.

And we are not stopping there. We are adding not just large language models, but we are also leading the small language model revolution. Our Phi-3 family of SLMs are the most capable and most cost effective. They outperform models of the same size or the next size up, even across a variety of language, reasoning, coding, as well as math benchmarks.

If you think about it, by performance to parameter count ratio, it's truly best in class. And today we are adding new models to the Phi-3 family to add even more flexibility across that quality cost curve. We're introducing Phi-3 Vision, a 4.2 billion parameter multi-modal model with language and vision capabilities. It can be used to reason over real-world images or generate insights and answer questions about images, as you can see right here.

And we're also making a 7 billion parameter Phi-3 small and a 14 billion parameter Phi-3 medium model available. With Phi, you can build apps that span the web, Android, iOS, Windows and the edge. They can take advantage of local hardware when available and fall back on the cloud when not, really simplifying all of what our VS developers have to do to support multiple platforms using one AI model.

Now, it's just awesome to see how many developers are already using Phi-3 to do incredible things. From Amity Solutions, the Thai company that I mentioned earlier, the ICC, which has built a Copilot for Indian farmers to ask questions about their crops, Epic in healthcare, which is now using Phi to summarize complex patient histories more quickly and efficiently. And another very, very cool use cases in education.

Today, I'm very thrilled to announce a new partnership with Khan Academy. We'll be working together to use Phi-3 to make math tutoring more accessible. And I'm also excited to share that they'll be making Khanmigo, their AI assistant, free to all U.S. teachers. Let's roll the video here already.

(Applause.)

(Begin video segment.)

TEACHER: I felt like I was in a place in my teaching career, where I felt like I was kind of losing my sparkle. And I would just feel really defeated when I looked out on the classroom and I would see students that just didn't look engaged.

SPEAKER: Teachers have an incredibly hard job, and what we think we can do is leverage technology to take some of the stuff off of their plate to really, actually humanize the classroom.

TEACHER: By some miracle, we became a Khanmigo pilot school.

SAL KHAN: With new advances in generative AI, we launched Khanmigo. The point is to be that personalized tutor for every student, and to be a teaching assistant for every teacher.

TEACHER: I started to build these more robust lessons, and I started to see my students engage.

SPEAKER: We're working with Microsoft on these Phi models that are specifically tuned for math tutoring. If we can make a small language model like Phi work really well in that use case, then we would like to kind of shift the traffic to Phi in those particular scenarios. Using a small language model, the cost is a lot lower.

SAL KHAN: We're really excited that Khanmigo, and especially in the partnership with Microsoft, being able to give these teacher tools for free to U.S. teachers is going to make a dramatic impact in U.S. education.

TEACHER: I think we're going to make them the innovators, the questioners. Isn't that really just why you wake up every morning, because that's our future, our next generation? And to me, that's everything.

(End video segment.)

(Applause.)

SATYA NADELLA: I'm super excited to see the impact this all will have and what Khan Academy will do. And Sal is going to, in fact, join Kevin soon to share more. And I'm really thankful for teachers like Melissa and everything that they do. Thank you very much.

Of course, it's about more than just models. It's about the tools you need to build these experiences. With Azure AI Studio, we provide an end-to-end tooling solution to develop and safeguard the Copilot apps you build. We also provide tooling and guidance to evaluate your AI models and applications for performance and quality, which is one of the most important tasks, as you can imagine, with all these models. And I'm excited to announce that Azure AI Studio now is generally available.

(Applause.)

It's an end-to-end development environment to build, train and finetune AI models and do so responsibly. It includes built-in support of what is perhaps the most important feature, which is in this age of AI, which is AI safety. Azure AI Studio includes the state of the art safety tooling to everything from detecting hallucinations in model outputs, risk and safety monitoring. It helps understand which inputs and outputs are triggering content filters, prompts shields, by the way, to detect and block these prompt injection attacks.

And so, today we are adding new capabilities, including custom categories, so that you can create these unique filters for prompts and completions with rapid deployment options, which I think is super important as you deploy these models into the real world, if an emerging threat appears.

Beyond Azure AI Studio, we recognize that there are advanced applications, where you need much more customization of these models for very specific use cases. And today, I'm really excited to announce that Azure AI custom models will come, giving you the ability to train a custom model that's unique to your domain, to your data, that's perhaps proprietary.

The same builders and data scientists, who have been working with OpenAI brought all the Phi advances to you, will work with all of you to be able to build out these custom models. The output will be domain specific. It'll be multitask and multi-modal, best in class as defined by benchmarks, including perhaps even specific language proficiency that may be required.

Now, let's just go up the stack to data. Ultimately, in order to train, finetune, ground your models, you need your data to be in its best shape. And to do so, we are building out the full data estate, from operational stores to analytics in Azure. We've also added AI capabilities to all of our operational stores, whether it's Cosmos DB or SQL or PostgreSQL.

At the core, though, of the Intelligent Data platform is Microsoft Fabric. We now have over 11,000 customers, including leaders in every industry, who are using Fabric. It's fantastic to see the progress.

(Applause.)

With Fabric, you get everything you need in a single, integrated SaaS platform. It's deeply integrated at its most fundamental level with compute and storage being unified. Your experience is unified, governance is unified, and more importantly, the business model is unified.

And what's also great about Fabric is that it works with data anywhere, not just on Azure, but it can be on AWS or on GCP or even on your on-premise datacenter. And today, we are taking the next step. We're introducing real-time intelligence in Fabric.

(Applause.)

Customers today have more and more of this real-time data coming from your IoT systems, your telemetry systems. In fact, cloud applications themselves are generating lots of data, but with Fabric, anyone can unlock actionable insights across all of your data estate. Let's take a look.

(Begin video segment.)

VOICEOVER: Introducing real-time intelligence in Microsoft Fabric, an end-to-end solution empowering you to get instant, actionable insights on streaming data. At its heart lies a central place to discover, manage and consume event data across your entire organization with a rich, governed experience.

Get started quickly by bringing in data from Microsoft sources and across clouds with a variety of out-of-the-box connectors. Route the relevant data to the right destination in Fabric using a simple drag-and-drop experience. Explore insights on petabytes of streaming data with just a few clicks.

Elevate your analysis by harnessing the intelligence of Copilot in Microsoft Fabric, using simple, natural language. Make efficient business decisions in the moment with real time, actionable insights, and respond to changing landscapes proactively. Allow users to monitor the data they care about, detect changing patterns, and set alerts or actions that drive business value.

All your data, all your teams, all in one place, this is Microsoft Fabric.

(End video segment.)

(Applause.)

SATYA NADELLA: And we're making it even easier to design, build and interoperate with Fabric with your own applications. In fact, we're building out a new app platform with Fabric Workload Development kit so that people like Esri, for example, who have integrated their spatial analytics with Fabric, so that customers can generate insights from their own location data using Azure's rich tools and libraries, right on Fabric. This is just exciting to see. It's the first time you know where the analytics stack is really a first class app platform as well.

(Applause.)

And beyond Fabric, we're integrating the power of AI across the entirety of the data stack. There's no question that RAG is core to any AI-powered application, especially in the enterprise today. And Azure AI Search makes it possible to run RAG at any scale, delivering very highly accurate responses using the state of the art retrieval systems. In fact, ChatGPT supports, for GPTs, data assistants API, are all powered by Azure AI Search today.

And with built-in OneLake integration, Azure AI Search will automatically index your unstructured data, too. And it's also integrated into Azure AI Studio to support bringing your own embedding model, for example. And so, it's pretty incredible to see Azure Search grow over the last year into that very core developer service.

Now, let's go up to developer tools. Nearly 50 years after our founding as a developer tools company, here we are, once again redefining software development. GitHub Copilot was the first, I would say, hit product of this generative AI age. And it's the most widely adopted AI developer tools, 1.8 million subs across 50,000 organizations are using it.

(Applause.)

And with GitHub Copilot, we are empowering every developer on the planet to be able to access programming languages and programming knowledge in their own native language. Think about that. Any person can start programming, whether it's in Hindi or Brazilian Portuguese, and bring back the joy of coding to their native language.

And with Copilot Workspace, staying in your flow has never been easier. We are an order of magnitude closer to a world where any person can go from idea to code in an instant. You start with an issue. It creates a spec based on its deep understanding of your code base. It then creates a plan, which you can execute to generate the code across the full repo, that is, multiple files.

At every point in this process, from the issue to spec to plan to code, you are in control. You can edit it. And that's really what is fundamentally a new way of building software. And we are looking forward to making it much more broadly available in the coming months.

And today, we are taking one more big leap forward. We are bridging the broader developer tools and services ecosystem with Copilot for the first time. We're really thrilled to be announcing GitHub Copilot Extensions.

(Applause.)

Now, you can customize GitHub Copilot with capabilities from third-party services, whether it's Docker, Sentry and many, many more. And, of course, we have a new extension for Azure, too, GitHub Copilot for Azure. You can instantly deploy to Azure to get information about your Azure resources, just using natural language. And what Copilot did for coding we are now doing for infra and ops.

To show you all this in action, here is Neha from our GitHub team. Neha, take it away.

(Applause.)

NEHA BATRA: Thanks, Satya.

GitHub Copilot gives you suggestions in your favorite editor, like here, where I'm writing unit tests. Copilot is great at meeting you where you're at, regardless of the language you're most comfortable with.

Let's ask for something simple like how to write a prime number test in Java, but let's converse in Spanish using my voice.

(Global language.)

Look at that. Gracias, Copilot.

Copilot is great at turning natural language into code and back again, but what about beyond the code? With the new GitHub Copilot Extensions, you can now bring the context from your connected systems to you.

Now, I can ask Azure where my app is deployed. I could ask what my available Azure resources are, or I could diagnose issues with my environment.

And this isn't just for Azure. As Satya announced, any developer can now create extensions for GitHub Copilot, and that includes any tool in your stack, including your in-house tools, keeping you in the flow across your entire day.

Actually, 75% of a developer's day is spent outside of coding, gathering requirements, writing specifications and creating plans. Let's show how GitHub Copilot can help with that, live on stage for the first time.

Typically, my day starts by looking at GitHub issues. Looks like we want to support a rich text input for our product description. Let's open Workspace and get some help with that.

Copilot interprets the intent of the issue to see what's required, and it then looks across the entire codebase, and it proposes what changes should be made. This specification is fully editable and the whole process is iterative.

But actually, this looks pretty good. Copilot can now help us build a plan on how to implement this change.

All right, that's a great start, but we must not forget about our documentation. Let's edit the plan and have Copilot update our readme.

And then we can even get Copilot's help and starting to implement the code for us.

Now, this was just a simple example, but in a large enterprise code base, there are tens of thousands of files and dozens of stakeholders involved. And that means meetings, so many meetings. Workspace helps you focus on what you need to change. And, by the way, as a developer, I'm always in control. I can see exactly what changes Copilot is proposing, and I can even get a live preview.

All right, let's test out the input.

All right, this looks great. I can go back, and I can edit my code in VS Code, or I can submit these changes as a pull request to share with my team.

GitHub Copilot, Copilot Extensions and Copilot Workspace help you stay focused on solving problems and keeping you in the flow.

Back to you, Satya.

(Applause.)

SATYA NADELLA: Thank you so much, Neha. I mean, I'll tell you, GitHub Copilot and everything that that ecosystem is doing is just bringing back a lot of fun and a lot of joy back to coding. And really, the thing about staying in that flow is, I think, what we all have dreamt for and dreamt about, and it's coming back.

That brings us to the very top of the stack, Microsoft Copilot. We built Copilot so that you have the ability to tap into the world's knowledge, as well as the knowledge inside of your organization and act on it.

Now, Copilot has had a remarkable impact. It's democratizing expertise across organizations. It's having a real cascading effect. In fact, it reminds me of the very beginning of the PC era, where work, the work artifact and the workflow were all changing. And it's just really having broad enterprise business process impact. I always say that it's lowering both the floor and raising the ceiling at the same time for anything any one of us can do.

Since no two business processes are the same, with Copilot Studio, you now can extend Copilot to be able to customize it for your business processes and workflows. Today, we are introducing Copilot Connectors in Copilot Studio, so you can ground Copilot with data from across the Graph, from Power Platform, Fabric, Dataverse, as well as you now have all the third-party connectors for SaaS applications, from Adobe, Atlassian, ServiceNow, Snowflake and many, many more.

This makes the process of grounding Copilot in first and third-party line of business data just a wizard-like experience, enabling you to quickly incorporate your own organizational knowledge and data.

We're also extending Copilot beyond a personal assistant to become a team assistant. I'm thrilled today to announce Team Copilot.

(Applause.)

You'll be able to invoke a Team Copilot wherever you collaborate in Teams. It can be in Teams, it can be in Loop, it can be in Planner and many, many other places. I mean, think about it. It can be your meeting facilitator when you're in Teams, creating agendas, tracking time, taking notes for you, or a collaborator, writing chats, surfacing the most important information, tracking action items, addressing unresolved issues. And it can even be your project manager, ensuring that every project that you're working on as a team is running smoothly.

These capabilities will all come to you all and be available in preview later this year. And we're not stopping there.

With Copilot Studio, anyone can build Copilots that have agent capabilities, and work on your behalf, and independently and proactively orchestrate tasks for you. Now, simply provide your Copilot a job description, or choose from one of our pre-made templates and equip it with the necessary knowledge and actions, and Copilot will work in the background and act asynchronously for you. That's, I think, one of the key things that's going to really change in the next year, where you're going to have Copilot plus agents with this async behavior.

You can delegate authority to Copilots to automate long-running business processes. Copilot can even ask for help when it encounters situations that it does not know much about and it can't handle. And to show you all of this, let's roll the video.

(Begin video segment.)

VOICEOVER: Redefine business processes with Copilot Studio. Create Copilots that act as agents, working independently for you.

Simply describe what you want your Copilot to do. Easily configure your Copilot with the details and needs, like instructions, triggers, knowledge and actions.

Quickly test your Copilot before you deploy, and seamlessly publish across multiple channels.

Watch it use memory for context, reason over user input and manage long-running tasks.

Copilot can learn from feedback to improve.

And you're always in control.

Put Copilot to work for you with Copilot Studio.

(End video segment.)

(Applause.)

SATYA NADELLA: All around this stack is perhaps one of the most important things that we, at Microsoft, are doing, which is wrapping it with robust security. Security underlies our approach with Copilot, Copilot+ PCs, Copilot Stack. We're committed to our Secure Future Initiative. You can see, you'll see us make rapid progress across each of the six pillars of SFI, and the core design principles, which is secure by design, secure by default and secure operations. You'll hear, throughout this conference, in fact, a lot more in Scott's keynote tomorrow, how it underlies everything that we build and everything that we do.

Coming to the close, there are many announcements that you will hear about at Build, but I want to go back to, I think, the core of what I think why we chose to be in this industry and why we come to work every day as developers, which is the mission, ultimately, of empowering every person and every organization. At the end of the day, it's not about innovation that is only useful for a few. It's about really being able to empower that everyone. And it comes down to you all as developers and builders of this new world.

For us, it's never, never about celebrating tech for tech's sake. It's about celebrating what we can do with technology to create magical experiences that make a real difference in our countries, in our companies, in our communities. Already, this new generation of AI is having an incredible impact, thanks to all of you, the passion you bring and the hard work you put in. And I want to leave you with this one unbelievable example of how you're all building a more accessible world, which means a lot to me, using our platform and tools.

Thank you all so very much. Enjoy the rest of Build.

(Applause.)

END