Build 2023
Satya Nadella
Tuesday, May 23, 2023

SATYA NADELLA: Good morning, it's fantastic to see you all. Welcome to Build, all of us. You made it here. It's fantastic to be back together, physically, and everyone is going online. These developer conferences are special time, special places to be, especially when platform shifts are in the air.

I distinctly remember my first PDC in 1991, driving up 101 into Moscone Center, and my life changed after that developer conference. So it's exciting to be able to come back to Build 2023 with that same sense of anticipation of something big that is shifting around us as developers.

To just sort of put this in perspective, in fact last summer I was reading Mitchell Waldrop's *Dream Machine*, while I was playing with DV3, as GPT-4 was called then, DaVinci 3, and it just brought into perspective what this is all about.

I think that concept of *Dream Machine* perhaps best communicates what we have really been doing over the last 70 years, right, all the way starting with what Vannevar Bush wrote in his most seminal paper, "As We May Think," where he had all these concepts like associative memory or Licklider, who was the first one who even conceptualized the human computer symbiosis, the mother of all demos that came in '68 to the Xerox Alto, and then, of course, the PDC that I attended, which was the PC Server 1 in '91.

In '93 is when we had the Mosaic moment, and there was the iPhone and the cloud, and all of these will be one continuous journey. And then, in fact, the other thing I've always loved is Jobs' description of computers as bicycles for the mind. It's sort of a beautiful metaphor, and I think it captures the essence of what computing is.

And then, last November, we got an upgrade, right? We went from the bicycle to the steam engine with the launch of ChatGPT. It was like the Mosaic moment for this generation of the AI platform. And now we look forward to, as developers, what we can do going forward.

And so it's an exciting time. And in fact, every layer of the software stack is going to be changed forever and there is no better place to start than the actual developer stack, right? We as developers, how we build is fundamentally changing.

In fact, when I think about how we build, I think about first Codespaces right? Being able to set up that environment in seconds versus minutes. A dev box, you know, instead of waiting for a day for your managed dev box to be set up, you have it in less than an hour. If you think about Copilot, what it does to you and Copilot X, in terms of driving the overall flow and productivity which is what, 54% or so up?

And then of course GitHub Actions and Azure development of dev environments really making that possible for you to stay – in fact, one of the things that I keep on bugging Scott Guthrie about for years is, "Hey, I want to stay in VS Code, I want to stay in Command Line, and let me do everything there."

We are close, close to that dream, and to me, that is bringing back both the joy of programing and the flow of programing. That ability to be able to stay on task, it's just so wonderful to see.

So how we build software is radically different, but what we are going to build as developers is really the story of this developer conference, and what we've built, we've been – you know, it's not like I came in on January 1st and said, "Let's start doing press releases," but it does feel like that. It does feel like we're – every week there is something new and, you know, infusing this new AI stack across all layers of it, right?

So we started with tooling in GitHub – or rather, Copilot in GitHub. We did Copilot in Power Platform, and when it comes to productivity, Copilot in Microsoft 365, Copilot in Viva, with business processes it's the Copilot in Dynamics 365. When it comes to industry, it's sort of workflows, and what Nuance has done with DAX or the Security Copilot, or the Copilot for the web in Bing and Edge, features in LinkedIn, which are driven by AI, and of course the AI infrastructure with Azure OpenAI APIs and everything else around it, right?

So every layer of the stack is profoundly changing. And today, as part of this developer conference, we are going to have 50-plus more announcements, but I want to highlight five of them.

The first is we're bringing search grounding in Bing to ChatGPT. We are very excited about this. Yeah, you can clap.

(Applause.)

Look, ChatGPT is the most sort of fast-growing consumer app we've ever seen, and search grounding is a very key feature, right, so that all the information is current and grounded by what you have from the crawl and the index. And so it's fantastic to see that. We are excited to be able – it's going to launch in ChatGPT Plus, immediately and quickly coming even to the free tier, and this is just the start of what we plan to do with our partners in OpenAI to bring the best of Bing to the ChatGPT experience.

Next, we are bringing Copilot to the biggest canvas of all, Windows.

(Applause.)

You are going to hear a lot from Panos tomorrow about it, but I think that this is going to make every user a power user of Windows.

Let's roll the video.

(Video segment.)

SATYA NADELLA: So cool.

(Applause.)

So we're going to talk a lot more about Windows tomorrow when Panos is up here. The other thing that we're also very excited to launch is the Copilot stack, right? After all, we built all these copilots with one common architectural stack. We want to make that available so that everyone here can build their own copilot for their applications. We will have everything from the AI infrastructure to the foundation models to the AI orchestration, all the way up to your copilot and its extensibility.

In fact, the other thing that we're going to do is have common extensibility across all of these services, right? Whether it is ChatGPT, Bing Chat, Microsoft 365 Copilot or all of the Microsoft Copilots, and of course, with your own copilots, we can share the same extensibility model. This is one of the most powerful things for your developers, for every developer to be able to write a plugin and have it reach billions of users across all of these surface areas.

So to be able to show you everything in action from both the plugin extensibility to all of the copilots we announced, let me invite up on stage Yusuf Mehdi to come show you all of this.

Yusuf, let me throw it over to you.

YUSEF MEHDI: Thanks Satya.

(Applause.)

Hi, everybody. We're making fast progress on delivering our vision of your copilot for the Web and for business. And today, as Satya said, we're excited to announce that we're going to bring ChatGPT and Bing together with the default search experience to give you higher quality answers and more timely answers.

Let's take a look. Here I am in ChatGPT. And as you can see now, Bing is the default, and when I come in and select it, I can now ask sort of real time queries. For example, let's ask what I should expect to hear from about Build and .NET. And what you can see is the results now are more up to date. They include fresh content, and they include citations. In fact, if you can see the links on that page there, you can click those and those will take you straight to a webpage that's sourced by Bing.

We're also excited. Yeah, absolutely. You can clap.

(Applause.)

We're also excited to announce that we're going to bring interoperability between ChatGPT and Bing for plugins. So you write them once and they're going to run everywhere.

So as you can see here in ChatGPT, I've got Zillow and Instacart enabled, but I want to show them to you here in Bing Chat. So we'll flip over, and you can see again, I've got the same plugins now in both Bing Chat and in ChatGPT.

And what we're going to show you now is I'll do a search here for houses in Chicago, and I can ask for a set of criteria. I'll learn a little bit about the neighborhoods, and now I can automatically call Zillow by saying, "Hey, give me three houses in a certain price range that meet my criteria." And what you can see is now I get these great options and I'm also going to get all of the other great things you get with Bing like helpful city guides and maps and prompts.

I'm going to show you now how we're going to further add value to the plugins that you write. They're going to work not just in Chat and ChatGPT. They're going to work across the entire web courtesy of the Edge browser.

So here's an example. I'm on a webpage here checking out a recipe for a cake, and now I can call Bing Chat and ask it to tell me, "Hey, give me the ingredients from this webpage." And notice Bing can read the context of the webpage, understand those ingredients, put them into Chat, and then I can say, "Hey, give me a shopping list for this," and it'll automatically call the Instacart plugin. It takes those ingredients right off the page and puts them into Instacart shopping, and with one click, I can get those now delivered to my house. This is an incredible productivity benefit for people.

Let me show you now now – yeah –

(Applause.)

Let's show you how you'll be more productive at work. Here I'm going to use Microsoft 365 Copilot. Now I'm in Microsoft Word and I'm going to need some help for drafting a legal contract. I've got a legal contract here and I need some help with California law. So I'm going to call three plugins from Thomson Reuters to edit this document.

The first thing is I'll go into Copilot, and I'll pull it up, and I'll say, "Hey, help me understand how to edit the limitation of liability using the Practical Law plugin." It'll read the document, find the paragraph and make that change.

Next, I want to know if this is enforceable under California law, so I'll call in the Westlaw plugin that will do that analysis and it'll come back and give me an analysis about it from a legal perspective.

And finally, since we're making lots of changes, I'd like to know the summary of all of these changes. And with Document Intelligence, I get a simple table that shows you all of those changes in an easy to read format.

By joining the power of Microsoft 365 Copilot in Word, with the support of these real powerful plugins like Thomson Reuters, now you can draft a legal contract in so much more powerful way.

Let me show you one more. Here I am in Teams Chat, and I'm engaging with Microsoft 365 Copilot to track website changes. Copilot will just call the Atlassian plugin to help. Atlassian Jira specializes in project and issue tracking, so it'll pull the Jira ticket automatically with the plugin. And now, all I have to do is assign an owner, using the Azure Active Directory, and that's it. It's all done. So with plugins and Microsoft 365 Copilot, you can intelligent reason across all of your business apps and the data stored in the Microsoft Graph to keep you in your flow.

Finally, as Satya shared, we're excited to announce the Windows Copilot. I think it's going to change how you use your PC forever. Let me show it to you.

Here I am in the coding project on my PC, but I want to configure my PC to help me be more creative and more productive. All I have to do now is invoke the Windows Copilot. I now just come down here to the taskbar. I click on that, and now will pop up the Windows Copilot on the right. This side pane here will be consistent across every app that you use on your PC.

And just like with Bing Chat, I can now ask it questions like, how can I adjust my system to get work done? And not only will I get a bunch of great suggestions but watch this. I can now with one click take action on those suggestions.

For example, I can put into Focus Mode. I also know that developers, we like dark mode. There's a suggestion here for dark mode. With one click, I'm now here on the dark side. And to really get going, I want to get that coding playlist going, so I'll pull the plug-in from Spotify and say, "Give me a great coding playlist." In this case, "Chill Vibes" will come up, and now I'll have it ready to go.

And finally, there's a suggestion here that says, hey, to organize your PC, let's take advantage of Snap. With one click, it snaps all the windows right in the place I need them so I can be super productive.

What do you think? (Applause.)

So as you can see, we have an incredible array of powerful AI-powered copilots. We've got over 50 plugins already available for customers and thousands more coming. I can't wait to see what you're all going to build.

Thank you very much. (Applause.)

SATYA NADELLA: Thank you, Yusuf. We have, as Yusuf said, fantastic momentum already building. And this is about really creating that opportunity for developers to reach all users across all of these surface areas. And we are so excited to see how you go about exploiting that opportunity in the weeks and months to come.

Of course, when we talk about the AI platform and the copilot stack, the next thing for us, which is really exciting, is AI Studio. This is the full lifecycle toolchain for you to be able to build your intelligent apps and your copilots, everything from being able to train your own models to be able to then ground whether it's OpenAI or any open source model with data that you bring, built-in vector indexing in Azure Search, built-in support for RAG, or retrieval augmented generation support, built-in support for prompt engineering with Prompt Flow and Orchestration, and of course, built-in support for perhaps the most important feature, which is AI Safety.

One of the things that we have been hard at work is to build into the toolchain AI Safety. We've been at work on AI Safety for the last five years. We have principles which we have translated into a core set of processes that we implement across our engineering stack. And then, of course, we have all of the compliance and oversight. But the real challenge is not just to have these things outside the engineering process, but to build it into the everyday tool chain.

And that's what we're doing with AI Studio, and it starts with testing. There is the Responsible AI dashboard that helps you during the testing phase to ensure that what you're developing is safe. We have grounding and, in fact, the Prompt Flow is perhaps one of the best features for you to be able to ground your models. You have provenance, provenance for media, provenance support for images and videos, and watermarking for your neural voice that's going to be available to all of you as you build your applications, and deployments in time.

That's perhaps one of the most critical things, is we have taken all of the safety work we did, for example, for the launch of Bing Chat, and really made it available as just a set of features for any developer to use, right? You can take an OSS model and use the AI Safety service to really make it, at the deployment time, safe. And of course, then you can even monitor the model for model drift. And that way, then you can make sure that it's not just a one time, but you're continuously looking to make sure that you have safe deployment.

We're very, very excited about AI Studio helping every developer out here to be able to build AI applications but build them with safety first. Let's roll the video.

(Video segment.)

VOICEOVER: Introducing Azure AI Studio, a full lifecycle tool to build, customize, train, evaluate and deploy the latest next generation models responsibly. With just a few clicks, developers can ground AI models with their structured and unstructured data to quickly and easily build customized, cutting edge conversational experiences for their customers.

Developers can take advantage of a new model catalog that works with the popular models organizations use, including those from Azure OpenAI Service, Hugging Face and many other open source models.

With Prompt Flow, developers can combine relevant data from your organization and create a detailed prompts to get better results. Prompt Flow works with foundations, internally developed for open source models, and uses popular open source tools, LangChain and Semantic Kernel. And because the AI systems we build are designed to support our AI principles, with Azure AI

Content Safety, we are making it easier for you to test and evaluate your AI deployments for safety.

This is Azure AI Studio, the trusted tools you need to build the next generation of AI applications.

(End video segment.)

SATYA NADELLA: Cool. (Applause.) And of course, all applications start with data, and we are really thrilled to be announcing Microsoft Fabric. This is a product that we've been working very, very hard on over multiple years, and it's finally coming together. It's perhaps the biggest launch of a data product from Microsoft since the launch of SQL Server.

It really brings together compute and storage. It unifies compute and storage, it unifies all of the full analytic stack product experiences. It brings together governance, so it unifies governance with analytics. And most interestingly, it unifies the business model, right, across all the different types of analytics workloads, whether they're SQL, machine learning. Whatever job you want, you can use the same compute infrastructure.

And this unification, at the end of the day, is what I think will fuel the next generation of AI applications. Let's roll the video.

(Video segment.)

VOICEOVER: Introducing Microsoft Fabric, a unified data analytics platform, one product, one experience, one architecture, one business model. Unified data is stored in OneLake, a SaaS data lake for the entire organization. Data is integrated and stored in an open format, allowing one copy to be used to train machine learning models, visualize data and run queries on the lake and data warehouse.

A unified experience brings together all the tools data professionals need, pipelines for orchestrating data movement, experiments for training machine learning models, Semantic models for defining key metrics and much more.

And for business users, Fabric brings together data for collaborating and doing ad hoc analysis in Microsoft 365. Unified governance, security and compliance is built in for all your data. And with Copilot for Microsoft Fabric, AI helps everyone be more productive, whether it's writing SQL statements, building reports, or setting up automations based on triggers.

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

(End video segment.) (Applause.)

SATYA NADELLA: What the AI supercomputer did for the infrastructure layer, Microsoft Fabric will do to the data layer for this next generation of AI applications; very, very exciting.

These are just the five of the 50. We have 45 more to discover throughout the conference. We're, again, really on a fast pace to build things that help us build this next generation of applications but build them with safety first.

Now, one of the things that I think we should ask ourselves as developers is why do we build? Why do we build technology?

This relationship between economic progress, and economic growth and technology has been there for a long time. In fact, this graph, when you see it, it's pretty stunning that for most of human history, we didn't have much economic growth, nor did we have much technology. And then something happened 250 years ago, right, which was long in building, by the way. From perhaps the enlightenment to scientific revolution to the Industrial Revolution was close to 400 years. But then there was real progress. You see that slope going upwards.

And then, of course, over the last 70 years, the information technology has played a role across all of those sort of seminal moments on the march towards that dream machine. And of course, we now enter the age of AI, and we get to define what the slope looks like, going forward, for economic growth.

But it's not even just economic growth on its own, right? We don't build just because we want economic growth. We want economic growth so that we can have human development index growth. We want the lifespans to go up. We want education and prosperity and standard of living to go up everywhere. That's why we build. That's why we innovate. That's why technology exists. It's not for technology's sake, but it is for that broad impact.

And to me, this all came together in January, when I was visiting India, I had a chance to see this demo. And it sort of had a really profound impact on me, because at some level, it sort of motivated me to go into this next wave with that much more rigorous ensure that this time around, this technology reaches everybody in the world.

There are two things that stood out for me, right, that things that we build can, in fact, make a difference to 8 billion people, not some small sort of group of people, and to be able to do that by diffusing or diffusion that takes days and weeks, not years and centuries, because we want that equitable growth. We want trust in technology. We want to ensure that we protect those fundamental rights that we care about, and that we do this in such a way that we manage our energy transition, given the finite resource we have in our planet. That's, at the end of the day, what grounds us in our mission to empower every person and every organization on the planet to achieve more.

And so, I want to leave you with this video of what you, as developers, are going to do in the days and weeks and months and years to come, is going to have the most profound impact of any technology to 8 billion people all around the planet.

Thank you very much and have a fantastic Build. (Applause.)

(Video segment.)

END