

05-24-22 Build Keynote Satya Nadella

## **Build 2022**

### **Satya Nadella**

**SATYA NADELLA:** Good morning and welcome to Build! It's great to be gathered here once again as a community of developers.

Every day, this community builds products and services that reach every corner of the world and touch everyone. And we all now have a profound understanding of how impactful all of the technologies you build are to navigating the challenges of people and planet.

Just think about that for a moment. What you dream, what you build is changing how we live and how we work. And so, we all have to answer two questions: What can we build? And more importantly, what does the world need us to build?

It reminds me of our own story as a company. When Microsoft was founded 47 years ago, it was a time of both great opportunity and crisis. The same month that Popular Electronics featured the Altair 8800, Newsweek's cover featured the three-headed dragon of inflation, recession and the energy crisis. Our company's founders grabbed the magazine on the left and ran with it, creating some of the world's most beloved developer platforms and tools.

Today, in 2022, we all need to pick up both these magazines and relate them to each other. That's our collective opportunity and responsibility.

So, in the spirit of asking what technology can do, I decided to have a little fun with DALL-E, which was trained on our cloud. I asked it to generate images that depict all of you building technology for some of the challenges of the world. And as you can see, it came up with many exciting possibilities across our society and economy. This conference is about helping you as developers realize these opportunities.

And today, I'll talk about 10 technologies that are coming together as a powerful platform to help you build what's next. Let's dive in.

It all starts with your flow as a developer. And when we think about your flow, it's more than the speed at which you work. It's about getting in the zone and staying there. We want to make it easy for you to go from idea to code, and code to cloud, and cloud to the world.

Our approach starts with GitHub, the home for every developer, whether you're helping digitally transform a Fortune 500 or starting a new Web3 project. Flow begins with your dev box. Today, you have to set up a new dev box whenever you start a project, investigate a bug in an older codebase, onboard a colleague, or learn a new technology. It's too complex, it takes too long, and it's too hard to keep workstations secure, compliant and up to date.

Last year, we introduced GitHub Codespaces to address this. Codespaces is your entire developer environment, including the full power of VS Code, hosted in the cloud and accessible via the browser. You can spin up a new, Linux-based dev environment for any sized project with pre-built images in seconds. In fact, GitHub's own dev image starts in just 10.

It's truly your on-the-go dev environment, accessible across the cloud, locally on your machine, with Codespaces and Windows Subsystem for Linux. You can save your changes and bring the entire environment to WSL locally, and when you go back online, you can commit your changes and return to Codespaces.

While Codespaces is a great environment for web development and building cloud-native apps, if you're building for the PC, embedded, mobile or console, you often need a Windows dev box and other client-apps to do your development, and you often need a solution managed by IT.

That's why I'm so excited to announce the Microsoft Dev Box. It's a fully managed service built on Windows 365 to spin the beefiest of VMs with all the tools, connected services and network resources you need pre-provisioned, so you can get to code quickly while ensuring that your dev box always stays secure and compliant.

It'll be compatible with, in fact, the Azure Game Dev VM we announced at GDC, so game developers can access all of the tools they use every day to test and build their games in the cloud. And from an IT perspective, it can be managed alongside all the other devices your company manages through Microsoft Endpoint Manager.

Now, let's take it one step further to help you stay in your flow. What if you had an AI pair-programmer that you could collaborate with and would help you learn new frameworks and best practices? That was our motivation for GitHub Copilot. It draws context from the code you're working on to finish the lines you start or even suggest entire functions. It's incredibly powerful.

We introduced a preview a year ago and the feedback has been fantastic. And we are already seeing more than 35% of newly written code being suggested by Copilot in popular languages like Java and Python. And one-third of the people who signed up for the preview are frequent users.

I'm excited to share that we will make Copilot generally available for developers later this summer, and everyone who's registered for Build today will have free access through GA. We're also making Copilot free for students as well as verified open-source contributors.

The next technology I want to talk about is the cloud, and in particular, the ubiquity of the cloud. We are seeing a major acceleration in the buildout of a highly distributed computing fabric from the cloud to the edge, as we embed computing everywhere. Azure is the world's computer, with 60-plus datacenter regions, connected by over 175,000 miles of fiber, over 190 points of presence. This delivers faster access to cloud services while addressing the critical data residency requirements.

And we're not stopping there. With Azure Arc, we're bringing the power of Azure anywhere. Customers now can run mission-critical data workloads in their own environments, while meeting latency and regulatory requirements. With Azure for Operators, we are partnering across the telecom industry to bring computing to the 5G edge through both private and public MEC.

We're extending our infrastructure beyond our planet, bringing cloud computing to space. Azure Orbital is a fully managed ground station as a service for fast downlinking of data, and we are partnering broadly across the space industry. In fact, NASA is testing these capabilities aboard the International Space Station. For astronauts, the integrity of the space suit is everything. And so, we have built a model with Azure Machine Learning to look for signs of damage on an astronaut's gloves after spacewalks. And they're running it locally on HPE's Spaceborne Computer.

And another example is Ferrovial, which builds and manages some of the world's largest, busiest airports and highways, from Dallas to Madrid. Let's take a look.

(Video segment.)

**SATYA NADELLA:** What a great example of a developer bringing together the power of Azure, Azure for Operators and Azure Space.

Now, let's talk about apps, and in particular, how they will be distributed in the future.

With Windows 11, we also announced the new Microsoft Store, which we rebuilt from the ground up to give you maximum openness and flexibility. We want you to bring your apps to our store, regardless of whether they were built as a native Windows app, a progressive Web app, or any other app framework. Last year, we announced our waitlist for Win32 apps, and today, I'm happy to share that the wait is over.

We also want to help you maximize your reach so your customers can find the apps and experiences that matter the most to them. That's why we are excited to announce Microsoft Store Ads, a new way to help you reach the right customers at the right time. As we continue to innovate with the best app store economics for you, you have the choice to use our commerce engine with industry-leading revenue share, or you can bring your own commerce engine to your apps and keep 100% of the revenue.

And remember, if you're targeting Windows, you're effectively targeting every device. With Windows 365, we are bringing Azure computing to Windows computers, removing the boundaries between the device and the cloud, so your Windows applications can be accessed on any device.

We're also further integrating Windows 365 and Windows 11 with new capabilities announced just last month. You can now switch between your local PC and your cloud PC with just one click.

The real magic happens when you accelerate Windows Cloud PCs with Azure Public MEC. We are testing these scenarios today. Let's take an architect using an iPad at a building site to share designs with site managers and workers.

On a busy worksite, there's no time for slow loading and clunky design rendering, and Wi-Fi isn't always reliable. Azure Public MEC ensures that compute-intensive applications, where sub-millisecond latency needs to exist between the device and the virtual machine, work instantly every time. And because Windows 365 streams your personal data, apps, content and settings to

any device, when the architect moves back to the workstation in her office, the project designs on her cloud PC are live and ready to go, and she can pick up where she left off.

Between Windows 365 and Azure Virtual Desktop, you have the most diverse set of options so everyone can work how and where they want.

What's happening with Windows is also happening with Xbox. We are bringing Xbox everywhere with Xbox Cloud Gaming. More than 10 million Xbox players have streamed games across 26 countries. As a game developer, you can publish to the Xbox Store and your game can be accessible on every device.

In fact, Epic Games is doing exactly that with Fortnite. Fortnite is now available to stream with Xbox Cloud Gaming via an iOS device, Android phone or tablet, or Windows PC through the Web browser for free. We're starting with Fortnite and will look to add more games people love in the future.

And one of the coolest things is that, as a game developer, there's no need to change your code. You can unlock this same opportunity for your customers. You can break free of device restrictions that too often get in your way. You can get started today by building on Windows to take your apps everywhere or Xbox to take your games everywhere.

Next, let's talk about cloud-native app development.

Today, all cloud-native development starts with writing a microservice. Microservices help teams build, deliver, scale applications faster than traditional architectures. Containers have emerged as the industry standard to go from code to cloud-native, and to scale from zero to infinity as your load and business requires.

I've been super excited about the response to our preview of Azure Container Apps, a fully managed serverless container service for building and deploying modern apps at scale. It allows you to create modern apps based on open standards. It's all built on Kubernetes so you can get the benefits of the open ecosystem, but no prior Kubernetes experience or knowledge is required. You can easily deploy containerized apps from Docker Containers or source code repos without having to worry about managing your own infrastructure. And your apps can scale in response to HTTP traffic or event triggers. And today, I'm thrilled to announce the general availability of Azure Container Apps.

And for those of you who want to access the flexibility and the power of full Kubernetes ecosystem, we're also announcing updates to Azure Kubernetes Service. It's the fastest way to spin up managed Kubernetes clusters so you can focus on developing your apps, using the open-source Kubernetes APIs and the community tools like Helm, Kustomize and of course, GitHub Actions. You get the best DevSecOps. You also have managed Kubernetes available to you in more regions than any other cloud provider.

And another game changer is Azure confidential computing, which helps you build not just for scale, but also for confidentiality. It protects data when in use, as well as at rest and in transit,

thanks to the enclaves that encrypt and isolate code and data in a Zero Trust environment, where even Azure, as the cloud provider, doesn't have access.

We are now the only cloud provider supporting the latest confidential-capable CPUs from Intel and AMD, as well as confidential GPUs with NVIDIA. And we are the only ones with a confidential code-to-cloud serverless offering. And we are seeing fantastic applications built with Azure confidential computing across regulated industries, including healthcare.

BeekeeperAI, a UCSF spin off, used our tools to create a unique platform for clinical AI that allows developers to access datasets from multiple institutions, where AI models are validated and tested without moving or sharing the data, or revealing the algorithm. This eliminates the need to de-identify or anonymize protected health information because the data is never visible or exposed in the first place.

The solution they built uses Azure confidential computing VMs running on 3<sup>rd</sup>-Gen Intel Xeon processors with Intel Software Guard Extensions. We're making these VMs generally available at Build, and they bring orders of magnitude more encrypted memory compared to previous generations.

All of these capabilities are available to you today, enabling you to bring even your most sensitive workloads to the cloud.

Now, let's turn to data.

When you step back, it's really your most important architectural consideration. We are moving to a world where every app will be intelligent and adapt in real time. Analytics won't just be a back-end process. It will become a critical part of the product experience.

Today, though, developers struggle with the silos that too often exist between databases, analytics and governance. Data engineers, data scientists and business analysts struggle with the complexity of making data integration, data warehousing, MLOps and BI all come together. And data and privacy officers struggle to meet the regulations when governance is not deeply integrated into the products they're supposed to monitor. That's why, instead of thinking about databases, analytics and governance as separate things, we are architecting these pieces together as part of one intelligent data platform.

And today, I'm very excited to announce the Microsoft Intelligent Data Platform. There is no better example of why this matters than e-commerce. Take, for example, a shopping website that needs to create personalized experiences, including product recommendations. Real-time personalization means the website needs to capture new interests from current sessions and instantly match historical data to online session data in real time and at limitless scale.

It's absolutely mission critical during peak events like Cyber Monday. The site needs to connect and correlate customers with products, inventory, suppliers, logistics and even social sentiment, all while respecting customer data privacy. They can build their application on Cosmos DB or SQL Hyperscale.

With Synapse Link, all committed transactions across these databases are now available in real-time analytics in Synapse. With Azure ML, they can easily retrain models based on this real-time learning and close the loop by deploying those models back on to the application. And, with Power BI, insights can be shared across the company. All of this can happen within just a few seconds, so the website can offer the right product to the right customer at the right time.

And, of course, your service has to ensure it's compliant with GDPR and other privacy regulations. And that's why we have deeply integrated all of our data services with Microsoft Purview, enforcing governance across all these data assets wherever they reside. This is about bringing together all of our data products into one architectural fabric so you, as developers, can shift your focus to your creativity, to what you can uniquely do instead of spending time integrating data services and governance.

I just talked about writing cloud-native apps and how they can be powered by this rich data fabric. The other thing I want to talk about is how apps can be infused with intelligence using AI to both improve existing scenarios and introduce new, more ambitious ones.

We are seeing a paradigm shift of large AI models becoming powerful platforms themselves. In Azure, we have built the most powerful AI supercomputer in the world, and we've used it to train these large, state-of-the-art AI models. We've trained Turing for rich language understanding, Z-Code for translation across hundreds of languages, and Florence for breakthrough visual recognition. And our partner, OpenAI, has trained the GPT family of models for humanlike language generation, DALL-E for realistic image generation and editing, and Codex for code generation in a more than dozen programming languages.

But it's not enough to talk about these models. It's all about applying them because the real power for all of you comes when you can build upon and embed them in your applications. That's why we are building models as platforms in Azure, so you can apply these advances to any use case.

In fact, we ourselves are doing this today across our product line. Nuance's DAX solution is applying the latest advances in AI to document patient encounters at the point of care, reducing the burden on physicians. Dynamics 365 automatically summarizes customer conversation into succinct summaries and allows the customer service agent to pass them off to relevant subject matter experts.

These same capabilities are also coming to Microsoft Teams to help reduce meeting overload. And with our Azure Cognitive Services, you can embed these capabilities in your own applications. We are announcing updates this week at Build, including summarization for documents and conversations.

We're also seeing customers in many industries use our Azure OpenAI service to do transformative things. CarMax, for example, has used the Azure OpenAI service to generate new marketing content based on thousands of customer reviews that would have otherwise taken someone years to summarize.

Similarly, inspired by GitHub Copilot, we envision a world where everyone, no matter their profession, can have a Copilot for everything they do. In fact, we are working with Autodesk to fine-tune the Azure OpenAI Codex model to add such an interface to 3D workflows in Maya. With simple natural language commands, creators will be able to add objects in space and modify them in a scene, change the color and attributes of assets without having to navigate a series of menus, and group, duplicate, position objects in relation to each other.

With these actions alone, they have cut what would have been several layers of deep menu commands down to simple, natural sentences. And we're just beginning to explore what is possible in this new realm of human-machine interface and no-code development.

We just talked about the progress we're making with our AI supercomputer, but if there's one thing that I believe that will surpass the AI supercomputer, it's the supercomputing capacity that will be distributed across all the devices we carry and use throughout our lives.

In the last two years, we have created 50 times more GPU compute on Windows compared to the biggest AI supercomputer. And with the advent of NPUs, this will only grow. We're entering a world where every Windows computer will draw on the combined power of CPUs, GPUs, NPUs and even a new core processor: Azure Compute. In this hybrid, cloud-to-edge world, you will be able to do large-scale training in the cloud and do inferencing at the edge, and have the fabric work as one. We are very excited about how it will enable both us, as well as our partners like Adobe, to create a whole new class of experiences in their applications.

We also want you to build cloud-native applications with these rich AI models where they can offload for local inference wherever NPUs are available. Whether that's Android, iOS or Windows, you should be able to optimize for the device and the local capability you have. It's why we've built a powerful, cross-platform development pattern for building AI experiences that span the cloud in the edge. This pattern allows you to make late-binding runtime decisions on whether to run inferencing on Azure or the local client. It can also dynamically shift the load between client and cloud. We are calling it the Hybrid Loop.

We are bringing this to you as a cross-platform capability through the Onnx Runtime and Azure ML, along with an AI toolchain to make heterogeneous targeting easy. And to bring this to life, we're also announcing Project Volterra. It's a Windows Dev Kit with an ARM CPU and NPU, and it includes native ARM64 Visual Studio and .NET support to provide the same fast, familiar and highly productive experience that you're used to. Let's take a look.

(Video segment.)

**SATYA NADELLA:** Project Volterra will be available later this year, and you can get started today by building on our cloud and taking advantage of all these fantastic AI tooling and services. And it's just the beginning of what will be possible.

Now let's talk about low code, no-code.

In a digital-first company, the developer workflow influences how the entire company works. Fusion teams of pro developers and domain experts integrate across all functions and disciplines, enabling closer collaboration. And we're continuing to innovate to make this even easier for you.

Today, we're announcing Express design in Power Apps. You can upload a PDF or a PowerPoint or even a hand-drawn sketch, as you can see here, and Express design will convert it into an app within seconds. It's pretty fantastic.

This includes a complete UI that's built from your design with a back-end data schema. It's easy to check and modify all of the created components. You can further modify any part of the app to match your own design. And if you want to add more screens, you can simply repeat the process until your app is completed.

For example, the interaction designer who is using Figma every day can build a prototype in Figma, upload the design into Power Apps and have a working app in seconds, leaving developers to focus on data and backend logic. And, by the way, this is all built on large AI models that I just talked about.

And we're not stopping there. Today, I'm very excited to announce Power Pages, which enables anyone to build modern, secure and responsive business websites fast. Power Pages offers a wide set of design, development and security options, along with simple authoring and advanced admin options.

We've already seen people use Power Pages to do meaningful and much needed work. In Portugal, Edgar Simões wanted to encourage the public to donate goods to help an influx of refugees. He used Power Pages to host and build his app, allowing him to launch a site with no HTML or web skills. He went from idea to impact in just days. Let's take a look.

(Video segment.)

**SATYA NADELLA:** Thanks, Edgar, for your commitment. And thank you, Nataliya, for all you're doing as well.

Now, let's talk about the future of work and how we're making apps more contextual and people-centric, so you can build a new class of collaborative applications.

It starts with Microsoft Graph, which underlies Microsoft 365 and makes available to you information about people, their relationships and all of their artifacts. Today, we are seeing developers around the world enriching their apps with Microsoft Graph. In fact, more than half of the Microsoft 365 tenants are using custom-built and third-party apps powered by the Graph. With Graph connectors, ISVs can extend their applications and have them be discovered as part of the user's everyday tasks, whether they are writing an email, meeting on Teams, or doing a search.

For example, data from an app can appear directly in an organization's search results, as you can see in the experience Figma is building here. You can compose email and at-mention files from these apps in-line, and you can access them in a Teams chat, too.



And another way you can create interactive experiences is by building live, actionable Loop components using adaptive cards, like partner Zoho has done. Your users can make decisions and take action, like updating the status of a ticket, right in the flow of work and updates are always live, like this one, across Outlook, Teams and Zoho.

When you combine the Microsoft Graph with Microsoft Teams, you combine the data that describes how people work together with the place they work together. It's incredibly powerful, and developers are extending their apps into Teams and embedding Teams in their apps. In fact, monthly usage of third-party apps and custom-built solutions on Teams has grown 10X over the last two years. And more and more ISVs are generating millions in revenue from customers using apps built on Teams. Adobe, for example, is integrating Sign with Teams, enabling users to share documents and sign them during Teams meetings from anywhere on any device.

At Build, we're introducing new capabilities to help you design rich, collaborative experiences. One of the updates I'm most excited about is Live Share Experience for apps in Teams meetings. We've offered real-time, collaborative software development capabilities for some time now in Visual Studio Live Share. And now, we're bringing real-time collaborative meeting capabilities to Teams with Live Share Experience.

Working on apps in meetings should be multi-player. With Live Share, your apps can go beyond passive sharing so your end users can actively co-watch, co-create, co-edit content together, making live meetings truly interactive.

Our partners are reimagining meetings with Live Share. One example is Hexagon. Hexagon builds manufacturing solutions for engineering teams and is using these capabilities to reinvent how collaborative engineering reviews are done. With Live Share, Hexagon brings engineers as well as simulation and quality experts together in real-time. They can highlight, annotate, edit 3D models and simulations. This connects traditionally siloed phases of product development together right within Teams.

And Skillsoft is using Live Share to deliver interactive training on Teams. With synchronized video annotations and instructor controls, Skillsoft is creating new ways for people to learn together. All of these capabilities will be available to you with Teams SDKs.

Finally, let's talk about the metaverse.

In the last two years, we have seen how immersive experiences in gaming can help create community and connection. And we've also seen the proliferation of HoloLens scenarios across industries, transforming everything from telehealth, to remote maintenance, to collaborative design and training. We are building out a diverse set of platform capabilities to help you build even richer applications that transcend both business and consumer scenarios and extend across all device platforms.

Our approach starts, in fact, with Teams. With Front Row in Teams Rooms, we are bringing people face-to-face, placing remote attendees at eye level like they are in the same room. We are

including hardware like intelligent cameras in meeting rooms for more natural, personal and inclusive connections. And with the latest AI and Teams Rooms experiences, we're dissolving the walls between digital and physical participation so people can present a PowerPoint together as though they were in the same location, even when they're apart.

Presence is the ultimate killer application. When you and I can remotely join the same meeting, but feel like we are in the same room to create and collaborate together, that's finally within reach.

I've been using Mesh for Teams internally, and I'm very excited. With personal expressive avatars, my digital twin looks and interacts like me. It represents me the way I want to be seen, even when I don't have my camera on. I can interact with all of the artifacts in the meeting. Think about what that unlocks with product design as an example! Or, I can drop into a persistent immersive space for more serendipitous meet ups. Think of this as the office watercooler. And we'll have more to share as we open up these experiences in the coming months.

And we're building Mesh as a platform service in Azure to enable you to build your own immersive worlds. You can build your metaverse experiences on Mesh that are accessible from anywhere on any device, whether it's HoloLens, VR headsets, phones, tablets or PCs. And we are taking this approach to ensure that developers are not bound to any one device platform.

Take, for example, the experience WEF is showing at Davos this week. In partnership with Accenture, the experience they've built on Mesh educates people on their 1 Trillion Trees initiative to unite and promote reforestation efforts worldwide.

The same paradigm shifts – the digitization of people, places and things and their interactions – are also happening in the industrial metaverse. We are helping companies optimize their operations and automate, simulate and predict every business function and process, using IoT, Digital Twins, Mesh and the HoloLens platform.

And we see this in action today across every industry, from A-B InBev, which is using our metaverse stack to track every bottle, from the wheat field through manufacturing to distribution, to Novo Nordisk, which is using mixed reality to optimize its pharmaceutical production while meeting regulatory requirements, and Marks & Spencer, which is using the data generated by every surface, screen and scanner in their stores to track store performance and customer experience in real time.

Kawasaki Heavy Industries, as the world's leading robotics company, is using our platform to create an industrial metaverse solution that brings their distributed workforce together with their network of connected equipment. Let's go to Kobe, Japan, and take a look.

(Video segment.)

**SATYA NADELLA:** It's so great to see what Kawasaki's doing and how they're using the metaverse to build the factory of the future.

Everything I've talked about today, and that you'll hear about this week across all the themes of this conference, accrues to creating more opportunities and options for you. And we are building Microsoft Cloud as a platform to help you build your own platforms. Here at this conference, we're announcing more than 50 new products and features spanning every layer of the stack. These are your platforms to build upon.

As you've seen, our vision for you is ambitious. You can set up a complete dev box in seconds and write code with an AI pair programmer, always staying in your flow. You can embed computing across the cloud and the edge where it's never been before. You can distribute your apps to any device. You can scale to infinity with the latest cloud-native technologies. You can build your apps with unified data fabric, as well as infuse them with the output of the world's most powerful AI models.

You can leverage all the computing power that will be created at the edge for the next-generation, AI-driven applications. You can amplify your work with fusion teams of citizen and pro developers. You can be at the forefront of emerging platforms, like collaborative apps as well as the metaverse.

Ultimately, this is about maximizing the opportunity for every developer by providing the richest set of technologies and platforms that empower you to dream and build what is next, breaking free of any arbitrary restrictions and removing friction so you can reach every person and every corner of the planet.

So, I'll end where I started. The question is not just what we can build, but what the world needs us to build.

Our Imagine Cup competition is a great example of the possibilities. For 20 years, students have shown us what's possible when they come together and apply technology to help solve the world's challenges. Over the years, more than 2 million participants have had awesome impact with technology. Let's take a look.

Thank you all so very much, and enjoy the rest of Build.

(Video segment.)

END