

Ignite KEY01

Satya Nadella – CEO, Microsoft

Callie August, Seth Juarez, Judson Althoff, Rajesh Jha, Charles Lamanna, Nicole Herskovitz, Pavan Davuluri, Scott Guthrie, Charlie Bell, Vasu Jakkal, Sherrod DeGrip

Tuesday, November 19, 2024

(Applause.)

SATYA NADELLA: Hello, and welcome to Ignite! It's really fantastic to be back here in Chicago with all of you and everyone joining from around the world. It's always fun to be in conferences like this, at times like this, when major platform shifts are in the air. They're exciting. The buzz, in some sense, I live for these.

In fact, I was — this morning, I was reflecting, what is it? Like 32 years ago, right here in this very conference center, at the spring COMDEX is when we launched Windows 3.1, and that was like a big deal. It was just actually a few months before I joined Microsoft, and it made Windows what it finally became. It was just the most major release of it.

And in fact, in 2015, Ignite itself started in Chicago, and that was right in the middle innings of the cloud. And so here, to be back again during the middle innings of AI is just fantastic. And so, I'm really excited about it. And you know, given I see this room, it's breadth first. So I'm going to give a keynote that's breadth first. So I'm going to try and cover everything that we have for you throughout the show.

With every platform shift, it's always perhaps good to build a deep context and understanding of the underlying forces. And today we talk about them as scaling laws. Just like Moore's Law, we saw the doubling in performance every 18 months. With AI, we've now started to see that doubling every six months or so.

Now, in fact, there's a lot of debate. In fact, just in the last multiple weeks, there's a lot of debate, or have we hit the wall with scaling laws? Is it going to continue? I mean, the thing to remember, at the end of the day, these are not physical laws. These are just empirical observations that hold true, just like Moore's Law did for a long period of time. And so, therefore, it's actually good to have some skepticism, some debate, because that I think will motivate, quite frankly, more innovation on whether it's model architectures or whether it's data regimes or even systems architecture. So it's a good thing to have.

In that context, though, if anything, we are seeing the emergence of a new scaling law which test time or inference time compute. In fact, OpenAI's o1 is a good example of it. And features like the Copilot Think Harder is built on o1, is all about using test time to solve even harder problems.

So ultimately, though, all these breakthroughs manifest in three capabilities that are

exponentially getting better. The first is this new, universal interface that's multimodal. It supports speech, images, videos, both as input and output. Second, we have these new reasoning and planning capabilities, which essentially we have new neural algebra to help solve complex problems. We can detect patterns involving people, places and things. You can even find relationships between people, places and things using this new algebra.

And third, we now have this capability to support long-term memory, rich contexts, and in fact teach these models to use tools. So if you put all those things together, you can build a very rich agentic world defined by this tapestry of AI agents, which can act on our behalf across our work and life, across teams, business processes as well as organization.

Now, you know I love this stuff. We are going to celebrate a lot of technology over the next multiple days, but it's worth reflecting. In fact, given we are in Chicago, there was a professor of philosophy right here in University of Chicago some 50 years ago, John Haugeland, who said the following: "The trouble with artificial intelligence is that computers don't give a damn, but we do." And that's what really grounds us. Amidst all this rapid change, we remain grounded in our mission to empower every person and every organization on the planet to achieve more using this technology to make a difference for themselves, for their teams and for the world.

It's not about tech for tech's sake, but it's about translating it into real outcomes. And today, I want to focus on AI and this transformational power as it drives growth in business. It improves efficiency; it improves operating leverage. And to do that, we are building out three platforms. Copilot, Copilot Devices, and Copilot and AI Stack. That's it. Those are the three platforms. But before I dive into all of the news across the platforms, I want to talk about that something that we are prioritizing above all else. That is security.

At this conference, you'll hear us talk a lot about our Secure Future Initiative, the progress we are making, the principles of secure by design, secure by default, secure by operations, but most importantly, our commitment to continuous improvement. Because this is not just a destination or stand down, one milestone thing. We'll never be done. In fact, we're only as good as our ability to defend against the next novel attack. And we are seeing real momentum, even with our customers and partners around implementation of zero trust using our tools and practices. In fact, a great example of this is what the U.S. Navy has been able to do. They've been able to respond to all the executive orders and mandates and meet, in fact, their zero-trust goals years ahead of schedule. And it's fantastic to see that.

We continue to invest in security. You can see the updates. We're making tons and tons of announcements that you'll hear about across the security stack. The one thing, though, I want to point out is Purview. It's probably the product for this conference because in the age of AI, data governance takes on an even more critical, central, important role. And in Purview, we're introducing updates to prevent everything, from oversharing risky use of AI such as malicious intent detection, prompt injections, misuse of protected materials. So there's a lot in Purview.

At the end of the day, though, we recognize that when it comes to security, it's fundamentally a team sport. And that's why we want to partner. And we are partnering broadly with the security community. And today, I'm really excited to announce our Zero Day Quest. This is the first big announcement that I'm excited about because it's a new hacking event. There will be \$4 million in rewards focused on securing cloud and AI. It's the highest rewards of any public hacking event in the industry, and the quest starts today and will culminate in an in-person hacking event next year. So we're very, very excited about it.

So with that, let's just dive into each of the platforms, and starting with Copilot. Copilot is the UI for AI. It's rapidly becoming an organizing layer for work and how work gets done. Every employee will have a Copilot that knows them, their work, helping them unlock productivity, enhancing creativity and saving time. And Copilot Studio will allow you to create agents that automate business processes. And every IT department will have a control system to manage, secure and measure impact. That's it. Those are the three basic concepts of the Copilot ecosystem.

Over the past year, we've seen incredible momentum of what lean did for manufacturing, AI will do for knowledge work. It's all about increasing value and reducing waste. Just take risk analysis at Bank of Queensland Group. In the past, when an incident occurred, they would have to go through thousands of documents and write reports, and now they use Copilot to be able to synthesize everything that happened and create that first draft. That means analysis that took weeks, it's just taking a day.

Vodafone is another great example. Their legal team used to manually analyze and draft and renegotiate thousands of contracts that govern their massive base station network. And now, they use Copilot to determine which contracts need to be renewed, which need to be scrapped and to keep track of all the expiry dates. Vodafone is also personalizing all of their customer service by leveraging Copilot, as well as Azure AI to manage customer inquiries. Their virtual assistant engages in more than 45 million monthly customer conversations and is reducing the average hold time by more than one minute.

And of course, we're just getting started. We're continuously innovating, shipping hundreds of updates since we made Copilot generally available a year ago. And it starts with actually work on the fundamentals. In fact, you could think about the Copilot responses today. They are more than two times faster than on average, and the response satisfaction is improved nearly threefold. And going forward, we think about Copilot adoption in three fundamental ways. The more employees use Copilot and Copilot agents, the faster they'll be able to realize value. And second aspect is: how do you ground Copilot or extend Copilot in your world, in your ecosystem? So Copilot is a platform you can extend with agents to scale what you can do, right? So if you have one employee, one Copilot, and one Copilot can have thousands, thousands of agents. And finally, it's about measuring ROI.

And let's start with Copilot Pages. Now, Pages is this first artifact of the AI age. And it's

pretty magical, right? You see, in this video, we're bringing rich artifacts to Pages. You can add interactive charts, tables, code blocks, math equations, complex diagrams. You can use Copilot directly on the page to iterate on the content, and also control what happens on the page directly from a chat. It's truly this multiplayer canvas that enables you to ideate with AI and collaborate with other people.

In fact, I use Pages all the time. In fact, if I'm prepping for a meeting, I just give it a prompt. The Copilot combs through everything about the customer, let's say, from the web, LinkedIn, all the business applications like our CRM. In fact, all the work artifacts are documents, emails, Teams messages, presentations. And then I can put all of that just as a first draft into my Pages and share it, let's say, with my account team, everyone working with me in my office. It becomes much more of a real-time way for every one of us who's working on that topic to be on the same page. And then it becomes, in fact, the first draft for what comes next, which is my meeting notes, right? Then again, it's real-time shared with everybody involved.

Copilot is also not just in Pages, but it's deeply integrated, right? It's not just about charts and pages, but it's deeply integrated into the entirety of Microsoft 365 system. Let's say it's starting with Teams. Copilot can reason over all the past meetings and chats and transcripts and get you quickly up to speed. With Teams, in fact, screen understanding, which we are announcing today, it's pretty exciting. Copilot can even answer questions about presentations and documents shared during the meeting as well, right? You are doing a presentation. Here is a Copilot that understands exactly what you're presenting and answer questions about that presentation.

In Word, Copilot creates a draft based on other Word documents, PowerPoint, PDFs, emails, meetings. You never have to start from a blank page. In PowerPoint, you can enter a prompt of what you want to do or what your presentation is about. And the narrative builder generates essentially a spec for your presentation, right? An outline with topics, and you can edit it, refine it, its suggestions, and it will create your first complete presentation.

And in Outlook, this feature is something that I've sort of gotten so used to. It's the quickest daily habit I've built, which is prioritize my inbox, is a total game changer because it helps you quickly get to the messages that matter. Analyzing your inbox based on both the content of the mail, but also the role of the sender and the context. And that's, I think, it's just like someone went and labeled every piece of email with a blue arrow and said: why is this important? Right? I mean, think about it.

And I think what we're doing in Excel, though, is perhaps one of my favorite things. Just like what GitHub Copilot did for software developers, Copilot in Excel will do for data analysts. I can start with a very high-level prompt, that in the past, for example, I would have given a data analyst — let's say I'm in a manufacturing plant, and I want to figure out how to improve production rate, right? So I give a high-level prompt and say, hey, come back with an analysis of what can we do here.

And Copilot uses this advanced reasoning capabilities to build out a full plan for strategic analysis. I can easily change that plan, adjust it as I want to, and then it just goes and executes the plan. And throughout I can see it actually doing the work. It first does visualization. It generates heat maps, scatter plots, histograms. It figures out, in fact, the key drivers of production, does a comparison of all those key drivers, figures, and it calculates and then summarizes the insights and the actions in minutes. Like, how amazing is that, right? It takes data analysis and makes it available to everyone who has an Excel spreadsheet.

If you sort of look back, and you say Excel was one of those products that at population scale improved number sense. I think Copilot in Excel with Python will improve analysis sense across the world. So I'm really, really excited about this.

Now, if you have all of this richness of Copilot that you're using across the length and breadth of Microsoft 365, let's talk a little bit about how you can extend Copilot. Today, I'm really excited to announce Copilot Actions. With these Actions, you can use Copilot to reduce the amount of time you spend on repetitive, everyday tasks that you do. In fact, the best way to conceptualize Actions is for those of you who use Outlook rules, this is Outlook rules for the age of AI, and it works across the entire M365 system, not just in Outlook. So it automates everything from asking for, let's say, a status update from your team, compiling weekly reports, to scheduling emails, requesting feedback on a document. Actions are a very simple, but yet powerful way for you to scale what you do, right?

So whatever was the thing that you had to do multi-step, you just create one of these actions, and it just does it for you. You can discover templates for actions which you can reuse in your everyday work. It's just a simple interface, and we're not stopping there. And today, we're introducing new agents you can use within the context of your team. Again, the best way to think about these are as just your teammates. They're scoped to specific roles with very specific permissions. Just like we have permissions and roles. For example, a facilitator agent is someone you can add to your Teams meeting, and the facilitator will help keep the meeting focused, moderate the meeting, chat, as well as the follow-up and action item.

Our project manager agent in Planner will help automate, in fact, all the key steps in a project management workflow. It'll create a new plan from scratch. It'll help oversee what's happening across the project, task assignments, content creation. And next is even self-service agents, right? So for these agents, provide really useful information, answer questions and policies. But not just that. When it comes to HR and IT, these agents will help you complete the task. Think of these as just augmenting your HR and IT departments.

And we're also announcing SharePoint agents. Every SharePoint site will now have a built-in agent. These agents provide instant access to real-time information and insights from your knowledge base in the flow of your work.

We're also giving you the ability to easily create your own agents using Copilot Studio. Sometimes we sort of mystify these agents as things that somehow require a lot of effort to build, but it's really pretty straight forward. In fact, our vision is that it should be as simple as creating a Word doc or a PowerPoint slide or an Excel spreadsheet, that's it. When you say "agent," think creating a doc.

This example is a good one, right? I can create field service sources, in this case, a SharePoint site in my Dynamics CRM. I can easily configure it to meet my specific needs, and you have an agent in seconds that's just now integrated into Copilot.

You can also make these agents autonomous using Copilot Studio, and they can always raise an exception in Copilot for input. Remember, even an autonomous agent from time to time will need attention, and it will need UI, and that UI for interacting with us is Copilot.

Now, just last month, we introduced 10-plus autonomous agents in Dynamics 365, that do everything from optimizing supply chain, to helping customer service teams resolve issues. For example, take a look at a sales qualification agent, autonomously researches all the leads that are there in your system, and flags the best prospects for you, and then grabs a personalized email that you can, of course, edit and send. And we are already seeing customers using capabilities like this.

In fact, McKinsey has built an autonomous agent that reduces client onboarding time by as much as 90%. DAO has built agents to optimize their shipping process or freight shipping process and projecting millions of dollars of savings even in just the first year.

And of course, when we talk about extensibility that includes changing how you interact with your bespoke business applications too, that's fundamental. The idea that I have to go to one business application at a time just goes away in this world of agents, and so we are very excited to share. Many of our partners have built their own agents and connectors in Microsoft 365 Copilot. That includes Adobe, and obviously LinkedIn, SAP, ServiceNow, Workday. Even companies like Cohere, they are building AI-first agents they are also integrating right into Copilot.

And now, to show you all of this in action, I wanted to introduce my colleague, Callie, on stage to take you through the entire Copilot ecosystem. Callie, over to you.

(Applause.)

CALLIE AUGUST: Thanks, Satya. Let's imagine for a moment that I work in sales. My day is comprised of a lot of important but time-consuming tasks, which can take away from the time I spend building customer relationships and closing deals. That's where agents come in, created in Copilot Studio. Agents range from simple prompt and response to fully autonomous. They help with everything from prioritizing leads, scheduling meetings, to fulfilling orders. Let me show you.

I have an autonomous agent that monitors and fulfills customer orders on my behalf. It alerts me when there's a cross-sell or up sell opportunity. I can see one of my long time customers place a big order for more products. The agent is able to quickly fulfill this order and send me an alert that this is a great opportunity to offer the customer additional products. In seconds, this deal just got the potential to be a lot more lucrative.

Let's take the next step and get ready for the pitch meeting. I need to quickly get up to speed on the latest product updates and roadmaps. For that, I go to SharePoint. Over the years, SharePoint has become the most used solution in the enterprise for knowledge and business process. This site is where we keep all of the latest product specs, roadmaps and training materials.

Now, every SharePoint site has an agent, so all this rich information is immediately accessible and useful. I can find just what I need in seconds by simply asking the agent. Now, let's jump back to this chat. From here, I can interact with the same SharePoint agent by simply @mentioning it. I can ask it to give me an overview of the relevant products, availability, and pitch deck.

Because the agent is grounded in this SharePoint site, it pulls from the right source of truth and gives me just what I'm looking for.

Now, it's been a while since I've met with this customer and, of course, I want to put my best foot forward and understand where they're at. Copilot points me to the sales agent, which surfaces all of the details that I need, like account highlights and opportunities, our most recent sales engagements and interactions with the customer.

I can see that I have a new counterpart on the customer side, and I want to learn more about them. For that, I'm going to turn to LinkedIn. I can simply @mention the LinkedIn Sales Navigator agent to see their background and experience, and who else I might know in their network. It looks like we have several connections in common.

Within seconds, I'm up to speed on my new partner, right in the flow of work. Now, I'm ready for the meeting, and I never even opened my CRM.

I'm going to fast forward to after the meeting. It was super successful, and now I'm going to work up a quote. Let's step over to SAP's Copilot Joule. I asked it to reference the latest pricing sheet, and I want it to reflect the appropriate discount given the size of the order. SAP Joule generates a detailed quote I can quickly review and send to the customer.

Now, I'm in great shape to land the business and it's not just SAP. Many of our largest partners are creating agents to bring their unique knowledge right into the flow of work, from HR to finance to sales.

Take ServiceNow, which covers customer service, human resource management and workplace services. Or, Adobe, which can help your team build world-class marketing, design and visual content.

Copilot will empower every employee, and Copilot Studio Agents will transform every business process. Back to you, Satya.

SATYA NADELLA: Thank you so much, Callie, thank you.

(Applause.)

SATYA NADELLA: Now, let's take a look at one more very important consideration, which is measurement. After users start using Copilot and all these agents, one of the fundamental things that all business leaders want to do is to figure out, how do we measure ROI?

And so today, we are very excited to announce Copilot Analytics. Here, if you take, let's say, a sales territory manager, they can now correlate the specific Copilot usage to a business metric, like their win rate over time. And it's not just Copilot, it's Copilot and all the agents that you have built. You can look at their usage and start tuning even the usage to the business KPIs.

Our goal is to show how Copilot usage is ultimately directly translating into business outcomes across sales, marketing, finance and more, and there are many, many more examples. This is fundamentally the process, because it's really a question of change management. Think of Copilot Analytics as a tool for all of us to change how work, work flow and work artifacts are all getting done. That, I think, is ultimately how we get ROI.

So, that's a look at Microsoft 365 Copilot, Copilot Studio and Agents and Autonomous Agents, that end-to-end system for AI-driven business transformation.

Now, let's move to the next platform, which is devices. In the age of AI, even the devices are fundamentally getting transformed, with both AI and cloud. What is happening here is that we are fundamentally taking what's happening in the cloud with AI to the edge, and think of all of this as one, continuous, distributed computing fabric.

Over the past year, we have introduced an entirely new class of Windows PCs designed to unleash the power of that distributed computing fabric across the cloud and the edge. We call these Copilot+ PCs.

We are working across the entire ecosystem. It's fantastic. Just like in the cloud, there's all the silicon innovation. Silicon innovation is back in a big way in the client, whether it's Qualcomm, AMD, Intel, all building fantastic systems for PCs going forward. And of course, we're working with all of the OEMs, and now you can see the real manifestation of all of that in these 40-plus FLOPS, all on the client with the NPUs.

Now, when it comes to the fundamentals of these PCs, also, whether it's battery life or performance, they're best-in-class. That was the other thing. This is just not adding an NPU, but it's about making Windows and Windows PCs just fantastic on fundamentals. And, of course, ultimately, it's about developers, right? What happened 32 years ago with 3.1 was about applications, and here we are with Copilot+ PCs.

We are back at it with Adobe, and WhatsApp, all who are seeing the capabilities of these new PCs. They're bringing their best applications, leveraging these NPUs to really deliver breakthrough AI experiences.

We're also delivering entirely new end points. Three years ago, we introduced the Cloud PC category with Windows 365, which securely streams your personal Windows desktop from the cloud to any device, whether it's iOS, Android or even a mixed reality, let's say a headset from like something like Meta Quest. We've seen unbelievable momentum and adoption for these Cloud PCs to remote workers, temporary workers, IT developers, all of the frontline scenarios in particular, even around disaster recovery.

In fact, it's my go-to developer desktop. It gives full access to GitHub, Code Spaces, VS Code Azure SDKs, is all the CLIs set up in one place, and I can access it everywhere. Our Windows App, now, gives us one-click access to all of these Microsoft virtualizations on any device.

Today, we are announcing, in fact, Windows App is coming to Android. We're excited about that. We are announcing mobile application management. I know this is something that IT has wanted for a long time, and it's both to iOS and Android. This means any employee can work on Windows 365, even on their personal devices like this iPad, because your corporate apps and files stay secure, all managed on the Cloud PC.

Windows 365, itself, is growing rapidly. It's grown by triple digits year-over-year, and it's now used by some of the world's largest companies, including Wells Fargo, Johnson & Johnson, Siemens.

And today, though, I'm really excited about the next big step here. I'm really thrilled to announce Windows 365 Link. You can see it right here. It's a simple, secure, purpose-built device for Windows 365. It's admin-less, password-less, and security configurations are enabled by default and cannot be turned off.

Windows 365 Link expands the PC category, or the Cloud PC category, by connecting you directly to your productivity in the cloud with no data or information left on any device.

Let's take a look. Let's play the video.

(Video segment.)

Now, I'm really excited about this device, and I'm also pleased to say that it's going to be

available in April of next year, so really looking forward to it. The form factors on Windows 365 Link give you another choice for Windows endpoints. But look, we fundamentally recognize how mission critical Windows is, and we are committed to both its security and resilience as a first-class priority.

The latest release of Windows 11 has over a dozen new security features, and most importantly, they're turned on by default, including device encryption enabled across all devices. In addition, we are excited to announce this new Windows Resiliency Initiative. Super important. We are doubling down on our commitment to make Windows secure and reliable for customers for all their mission-critical workloads. As part of this work, we are making changes to low-level operating system access. We're introducing new features in partnership with the entire ecosystem, establishing new guidelines for safe deployment practices.

One example of this, which is, I think, something that is really exciting is Windows Hotpatch, which works across your entire Windows estate to apply critical security updates without requiring a restart. We also continue to push the envelope on Windows security and resilience across both the cloud and the client. And thanks to point-in-time restore, customers who use Windows 365 can be up and running in minutes and roll back a cloud PC to its exact earlier state, so it's pretty awesome to see all of this support for all the mission-critical workloads that run on Windows.

That's what we're doing with Copilot devices. And now I want to get to the final third platform, which is Copilot and AI stack. Now, the way we are approaching this is pretty simple. What we're doing is we're taking all of these apps we're building with Copilot and agents in Copilot Studio and exposing every layer of that tech stack so that you can use that to build your own Copilots and agents. That's it. That's as simple as it is. Every app is becoming an AI app, and over the past year we've seen unbelievable momentum in what people have been able to build.

Take a great example like NASA. Data scientists who built the Earth Copilot have all these enormous amounts of geospatial data, contains tremendous insights. Everything from climate and air quality that can be super helpful for urban planning or disaster response, but its scale and complexity are quite difficult to analyze. And so, the Earth Copilot makes it possible for anyone to navigate all this data using just natural language for the first time. For example, you can see how air quality right here in Chicago has changed over the years. That's an analysis you can now just do in natural language.

On the other side of the world, Toyota has built Obeya, meaning big room in Japanese. Their ambition is to create this big room of AI agents that are accessible 24/7 for all their engineers. It's all grounded in their engineering designs, regulatory information, even handwritten docs. You can OCR them, put them into this Copilot. Engineers can ask it anything from how to make a car run faster or something super specific about some emissions output. And so, these are all the examples.

In fact, right after my keynote, you're going to hear from Lance from Blackrock in conversation with Judson on what they're doing with Aladdin and Azure. It's really an exciting work that they're doing, so it'll be great to listen to Lance. These are some examples on how the stack is being used today, and we're just getting started. Adding, in fact, capabilities across every layer of the tech stack here, starting right at the infrastructure layer. We continue to build out Azure as the world's computer. Over the past year, we have added and made so many data centers in 15 countries, data center investments in 15 countries across six continents. We now have 60+ data center regions, more than any other provider.

We're innovating to build these data centers sustainably. I'm really proud of this. In fact, we just announced two data centers in Northern Virginia built completely with low-carbon, cross-laminated timber to reduce embodied carbon footprint. This new construction model will really reduce the carbon footprint of our data centers by 35% compared to any conventional steel construction, so it's exciting to see this. — (Cheers, applause)— And when people say data center is the computer, increasingly we really think of the system starting right from the construction all the way.

At the network level, we're delivering innovation in hollow-core fiber. This technology delivers absolute breakthroughs, whether it's speed or bandwidth or power efficiency, in fact, compared to traditional fiber. Obviously, photons travel faster in air compared to glass, and earlier this year, we demonstrated fiber loss at the lowest level ever achieved in optical fiber. This low fiber loss is absolutely critical for data center to data center connectivity, and we now have production routes of hollow fiber running. In fact, we're going to add 15,000 additional kilometers planned over the next 24 months.

We're not stopping there. We're extending our cloud to the edge, and today we're going further and announcing Azure Local. — (Cheers, applause)— This is, again, something that many of you have asked us to do, which is bring Azure Arc all the way to all of the edge with Azure Local extends Azure services across hybrid, multi-cloud and edge locations with one central control plane. It brings Azure services to customers' distributed locations, whether they're in retail, hospitality or manufacturing, so that they can run their mission-critical workloads, some of these new AI workloads across cloud and edge.

A great example of this is how Armada has taken Azure Local and in fact, with Starlink connectivity, is helping Marriott with full resilience even in the midst of real extreme weather events. Now, when it comes to silicon, our Cobalt 100 VMs became generally available last month, and customers and partners, whether it's Databricks or Siemens, Snowflake, are seeing up to 50% improvements in price performance with these virtual machines. In fact, our own media processing capabilities in Teams are now all 100% on Cobalt 100 as well, and we're not stopping there with silicon innovation.

Our approach to silicon includes deep commitment to security, and I'm excited to announce our first in-house security chip, Azure integrated HSM. This is it. — (Cheers, applause)— This is a dedicated hardware security module that hardens key management,

managing encryption and key signing, that can remain within the bounds of the device without compromising performance or security. And starting next year, it'll be part of every new server deployed on Azure, enhancing security for both confidential computing as well as general-purpose virtual machines and containers. We are very, very excited about this new silicon innovation.

Beyond hardening security, ultimately, when we think about systems innovation at the silicon layer, it is about removing bottlenecks that stand in the way, whether it's performance, latency or resource constraints. That's the opportunity we have as a hyperscaler. When you look at the workload, you look at the system, you say, what can we offload to silicon just to kind of get rid of some of these constraints. And so, that's why today we're very excited about expanding Azure boost with our first in-house DPU. DPUs are processors specifically architected to accelerate data-centric workloads, absorbing multiple components of a traditional server into a single piece of silicon.

It runs, in fact, cloud storage workloads at three times less power and four times the performance. I mean, what this will do for storage is what SmartNIC did for hosts in the network, these are going to do for storage. And of course, when it comes to AI, we are continuing to build out these new data center intelligence factories. We're extending all of Azure as the world's computer to basically be these intelligence factories. Tokens for Watt plus dollar is the best way to think about the new currency of performance. It's all about maximizing that value and doing it in the most efficient way.

The pace of innovation across the industry is simply phenomenal, and we are working with our partners to bring you more choice across cost and performance. That includes our deep partnership with Nvidia, which of course spans bringing their own workloads, whether it's Omniverse or DGX cloud onto Azure, but of course, working with them at the core system level on AI infrastructure.

In fact, last month we brought new clusters with H200 that became available. We're very excited about it. And our systems stack optimization, we have done between H100 to H200 continues to push on the total performance that we can deliver to anyone doing inference or training. And today we are announcing the preview of Nvidia Blackwell AI infrastructure on Azure.

(Applause)

Now, Blackwell is pretty amazing. It's got this 72 GPUs on a single NV Link domain, and then you combine it with InfiniBand on the back end. These racks are optimized for the most cutting-edge training workloads and inference workloads, so we are very excited about having Blackwell. And we're also continuing to work very closely with AMD. We're the first cloud that offered VMs powered by AMD's MI300X GPU, and we're using that infrastructure to power Azure OpenAI, so that's in our fleet today. And today we're introducing, in fact, Azure HB v5, which we co-engineered with AMD. It's up to eight times faster than any other cloud virtual machine, setting a new standard for high-performance computing, and it will be generally available next year.

(Applause)

And now, of course, when it comes to our own AI accelerator, MAIA 100, I'm really excited to say that it's right now live in the US East region supporting Azure OpenAI inferencing one. In fact, one of the most impactful applications we built is our customer service today, and it's all running today on MAIA 100 during all of the customer support workloads. It's exciting to see MAIA make it into the fleet, and we'll continue to improve it and scale it. We are contributing, in fact, our own systems innovation to the industry.

Last year at Ignite, we showed our first-generation MAIA RAG, which was that liquid cooling sidekick. These can support cooling of GPUs and AI accelerators, not just our own MAIA systems, but we're going to bring that innovation to Nvidia GB200. And you'll see this in the show floor because this is the type of systems innovation to get that performance. Ultimately, this is all about really being able to bring compute storage network edge silicon to deliver that TCO performance for all of your workloads.

Now, with that infrastructure layer, let's move to the next one, which is data. There's no AI without data. In order for you to build your AI applications, you need to be able to rendezvous your data with your AI compute effectively, and we are building out this full data estate to just do that. At the core of our platform is Microsoft Fabric, which we introduced last year at Ignite. It brings together all of your data, as well as all of your analytical workloads, into one unified experience with OneLake.

In fact, in Fabric, you can easily unify your data no matter where it lives in Azure or whether it's on premise, whether it's on Amazon or GCP, you can create this data layer for AI workloads in one place. The momentum actually has been incredible. We have over 16,000 Fabric customers, including 70% of the Fortune 500. Today, we're taking the next step with it. Until now, if you think about a typical data architecture, you required separate services for your operational stores and your analytical stores, and a lot of data had to be shuffled between these two.

We are very excited to announce that we are bringing our flagship operational database, SQL Server, natively to Fabric with Microsoft Fabric Databases. Just like Fabric simplified every aspect of an organization's analytical needs, we want to do the same for operational databases. Now, with Microsoft Fabric, customers have an enterprise data platform that serves all of their use cases. Whether it's batch data, real time or even massive transactional performance, all in one unified product. And all of that data is in open-source formats in Fabric's OneLake. This new database experience enables you to autonomously provision a database that's secured by default in seconds for OLTP applications, while simultaneously creating that connection to analytical workloads right there inside of Fabric. You have both your operational and analytic workloads, essentially like a SaaS service.

It integrates with developer tools like VS Code and GitHub, so you can now utilize your unified data estate and build applications. This creates a unified data platform with the ability to apply AI across all of your data, operational and analytical data. And for those that need additional customizations and control, our Azure databases offer just that and are optimized for AI.

In fact, one of the things that we're very, very focused on is vector search at scale because that's one of the most important operations when it comes to AI and data. The DiskANN was developed by Microsoft Research to power that low-latency, high-scale, cost effective vector search. We ourselves have been using it for a while now for our 400+ billion vector indexes in Bing, and with 10,000 or so real time updates and query latencies of less than even ten milliseconds.

And today at Ignite, we are bringing this very powerful DiskANN technology to our Azure databases, including both PostgreSQL as well as Cosmos DB. Excited about this really cutting-edge stuff coming to databases.

Now, let's talk about so you have your infrastructure, you have your data, and now it's time to build some apps. When it comes to applications, every application is an AI application. And every new generation of apps has brought a changing set of needs. Right? When it is the web or mobile or cloud you need to build a new app platform. The same thing is happening with AI. AI is transforming how we design and customize and manage apps today, and that's why we are building out a first-class app server for the AI age, announcing Azure AI Foundry.

With Foundry, we are unifying all of our models, tooling, safety and monitoring solutions into a single experience integrated with the most popular developer tools available as a standalone SDK and a portal. It all starts, in fact, with our approach to models. We know models and model choice obviously sit at the core of every use case. You want to be able to optimize for COGs, latency and performance, and we want to help you choose the right model for the right job. In fact, we now have 1,800 models in our catalog.

OpenAI continues to make unbelievable innovation. They're setting the pace around model innovation in everything they're doing right. Even the latest open models are all available on Azure. In the last six months, Azure OpenAI consumption has more than doubled. But we also support all these other models, like open-source models from Meta and Mistral, as well as providers like Cohere, so that you all can choose the right model.

In fact, no application uses one model. If you look at even Microsoft apps, they're all thousands of models that are being optimized, fine-tuned, distilled, that all come together. And so new categories of models are also emerging. In fact, we've added over 20 industry models designed for very specialized use cases from partners like Bayer, Page, Rockwell, Siemens, Site Machines and others. With all of this model diversity, it's never been more important to have the right tools, though, to choose the models for the job at hand.

With Foundry, we are adding model experimentation capabilities. Meaning for the first time, you'll be able to experiment with several of these models, compare the outcomes, and choose the best model that works for you. In addition to our own offerings, we're also announcing new collaboration to help developers accelerate this model customization. From data prep and generation to training eval experimentation with fine-tuned models, these are all the considerations of an app server.

Our work with Gretel Labs and Scale AI helps developers remove data bottlenecks, make data AI ready for training. We are working with Statsig to help customers configure run these fast A/B tests using different models. New integrations with weights and biases brings a comprehensive suite of tools for tracking, evaluating, optimizing models using Azure OpenAI service.

This is all available on Foundry. We are excited also to introduce a new service to help you simplify the creation of all of these AI-powered agents. Our new agent service helps developers build, deploy and scale AI apps that automate business processes. I showed you before how you can use Copilot Studio to build agents just with a few clicks, but as developers you want a code first approach to building an agent, and that's what the agent service really enables. You can build agents that are grounded in data wherever it lives. Public data from the web, enterprise data in Microsoft 365, SharePoint. Or you can leverage, even Fabric's OneLake to unify your data across all of your clouds.

These agents then can take action, right? So, you want to be able to give action space to these agents. And you can take those 1,200-plus connectors we have in Logic Apps that we have been using in our app services, and you can connect it to the agent runtime. We also know that the multi-agent stuff is becoming pretty exciting. It's evolving very, very quickly. And we are making sure that even at this very early stage, our agent service supports all the multi-agent frameworks, right? Which is Magentic-One, Autogen, Semantic, Kernel. You can use all of these within Foundry and agent services to build out your applications.

AI apps have specific operational considerations as well. Right? Because when you build an application at your organizational level, you want to be able to manage AI costs, performance, safety and security. And that's why we are very excited to share that we are bringing really new management capabilities to Foundry. For example, we are introducing AI reports that will help developers document and share their applications, use cases, and most importantly, eval results. Because I think one of the ways we're going to be able to think about, and talk about, and reason about what your application does, is through the eval of the models on which it's built.

Safety is the most important feature of any AI app, and we'll continue to ensure we have the best tools to build these secure AI apps, including things like prompt, shield, and so that you can detect and block manipulation of outputs for your business content. And today also we are announcing risk and safety evaluations for image content. That's become a very important consideration right in Foundry. Now to show all of this AI

Foundry and how you want to build applications, I wanted to invite up on stage Seth, to show you all of this in action. Seth, take it away.

SETH JUAREZ: Thank you so much, Satya. Let me show you how Azure AI Foundry offers everything developers, AI engineers, and AI and IT pros need to build and manage transformative AI apps. AI Foundry makes all of this easier with three groundbreaking capabilities. Our model experimentation tooling, our new agent service, and observability features like evaluations that help manage these apps even after they're in production. And by the way, everything I'm going to show you is actually live. OK.

In this demo I will use GPT-4o, our new real time audio model, GPT-4o audio, the real time API, and image processing models all seamlessly integrated into the experience I'm going to show you in a little bit. Foundry offers over 1800 models for you to choose from. It's pretty amazing. Once I've chosen my model, I can use Azure AI agent service, call APIs when needed, and perform actions with function calls. In this example, I have a main customer facing agent that is then delegating tasks to other agents to understand what the customer wants, what their purchase history is and to recommend products based on multimodal conversation with the customer.

These agents can run micro-intelligent tasks to provide a cohesive, personalized experience. These agents can also be grounded to a variety of data sources, including Fabric; it's pretty awesome. Fabric's OneLake makes grounding your agents really easy because it brings together your data across not only Azure, but also data in AWS, GCP, or from SaaS platforms such as SAP and many others.

And for CIOs, IT pros, business decision makers, we know compliance, measurement and observability are paramount. With Foundry, you get evaluations with your own data. All evaluation results are securely stored, ensuring you have the data you need for audits, compliance checks and continuous improvement. You also get experimentation and AI tracing that captures precise information, allowing you to correct issues proactively. Let me show you how all of these capabilities I just showed can create a hyper personalized customer service experience that solves real world problems in ways that we couldn't even have imagined a year ago.

OK, so I'm going camping, and my friend sent me this picture, and I want to make sure I have the right gear for a winter camping adventure, so I started a chat. Now notice this is something that's now routine, AI interacting with me with chat and images. The chat is great and we all love chatting, but what if we can break out of the actual chat box and do a more personalized experience? Because it feels like AI today it's all just like inside of your chat box. It feels like there's a lot of Copilot and stuff, but what if we could do a little bit more and maybe make it a little more personal, hopefully with a phone call? Hello? Is somebody there?

COPILOT AGENT: Hi, Seth. Yes, I'm here. I see you're getting ready for a winter camping trip. Can I send the concierge page to your browser to provide more details?

SETH JUAREZ: You certainly can. Send me to the concierge page, please.

COPILOT AGENT: Sending you to the concierge page now. You'll find more detailed guidance there for your winter camping trip.

SETH JUAREZ: This is awesome. What do you recommend for this camping trip?

COPILOT AGENT: Hey, Seth: Super exciting that you're gearing up for a camping trip. Since you've already got the Skyview two-person tent, let's build on that. Here are a few recommendations. First, Mountain Dream sleeping bag.

SETH JUAREZ: Is there a way that you can write it up for me so that I can see it, maybe with pictures?

COPILOT AGENT: Absolutely.

SETH JUAREZ: Right now, what it's doing it's going to take a couple of seconds to write an entire article for me, just for me. And you're probably wondering, well, hey, it's taking a couple of seconds. When's the last time you wrote a personalized article for someone that calls you in 10 seconds?

COPILOT AGENT: I'll put together a personalized recommendation article with all these products and their details, complete with pictures. Give me a moment and it should pop up on your screen shortly.

SETH JUAREZ: You already did it. This is amazing. I love everything here.

COPILOT AGENT: Fantastic to hear, Seth. I'm glad you love the recommendations. Enjoy your camping trip! If you need any more help, feel free to reach out. Have an awesome adventure out there!

SETH JUAREZ: How about them apples? Remember, like I'm serious, this was running right here because I wanted to show you something real. This isn't a peek into the future. It's happening right now with Azure AI Foundry. Back to you, Satya.

SATYA NADELLA: Thank you so much, Seth. Yeah, that's pretty awesome. You know, taking one of those multimodal models and then the app server Foundry and to be able to build these sophisticated applications. That's the app platform today.

Now, let's go to dev tools. We have the best developer tools for this AI era. GitHub Copilot is by far the most widely adopted, used AI developer tool, and we are focused on making it even more of a game changer.

Until recently, if you think about my own usage of Copilot, you can use the editor with completion, and you use chat no more. Now, with Copilot Edits, you get both of them to come together. We are bringing chat and inline editing together, so that you can easily

make inline changes across a set of files. You can have a working set of files and then just use natural language to be able to change across all files. Copilot has gone full multifile.

And with Workspaces, Copilot is the most advanced in the first agentic AI-native IDE. Copilot Workspaces leverages agents from start to finish. You can go from sort of basically, an issue to a spec to plan to code, all in natural language. And just last month at GitHub Universe, we added agents on both ends of the workflow, one for ideation and one for automatically building and repairing code.

And we're not just stopping with just writing code. We are building software agents across the entire lifecycle, from testing to deployment. And we can tackle these complex code maintenance tasks, like upgrading an app framework. Right here, you have this example of a Java framework where you're applying updates, iterating until, in fact, that all of your code builds and all the tests pass. And so, that's just an agent that's doing that for you.

Agents can also improve, in fact, performance. This is pretty cool, which is performance engineering agent effectively, by creating performance benchmarks, running them, iterating on the code, until, in fact, you find a solution, and you pass the performance evals.

And finally, agents can also help you go seamlessly from idea to implementation all the way to production by creating all the resources on Azure and deploying it. Think of all of the DevOps functions and having an agent for it.

And all of this innovation is what we're working on. It will just ship in the months to come. We're very excited about it.

Now so far, we've talked a lot about how AI can drive productivity by understanding, fundamentally the language of business, but AI can also drive fundamental business transformation by understanding the language of nature and science.

Science itself is becoming computed science, and that's one of the reasons why we are focusing on delivering the systems and AI innovation to power breakthroughs in material science, chemistry, physics and more. And our new frontier in AI-powered science is moving from static prediction to dynamic prediction, meaning not just predicting the shape of the molecule, but understanding the dynamics their motion and how they interact, which is a critical step in developing new materials and new medicine.

Earlier this month, in fact, Microsoft Research published in *Nature* this AI-driven simulation system that can accurately model protein behavior down to individual atom orders magnitude faster than ever before. This is, I think, a real breakthrough that will then help biomedical research and really, advances in areas such as drug discovery and protein design and enzyme engineering, all of these functions, because you need those dynamic systems to drive them.

And this is not theoretical. We're bringing together these advanced AI models and agents to help scientists reason over and orchestrate across the entire scientific method effectively. And we are already delivering on this vision of a platform for scientists with customers around the world.

Novartis, for example, is using generative AI to design hundreds of new molecules for drug discovery projects, helping accelerate the process. Nissan has partnered with us to create a model to predict EV battery performance over time, improving it by something like 80%. Unilever is running lots of simulations to really accelerate their R&D, using AI. In fact, the Institute for Protein Design at the University of Washington, scientists are using our cloud to engineer new proteins from scratch that promise to be absolute game-changers in medicine and sustainability and other fields.

Let's just take a quick look.

(Video segment.) (Applause.)

SATYA NADELLA: It's amazing to see it. Congratulations again to David, who was just awarded the Nobel Prize in Chemistry for his work. And it's great to see RoseTTAFold is open source, and it's available, in fact, right in that Azure AI Foundry catalog.

(Applause.)

Yeah. It's the most accessible tool for scientists in the field. And so, it's great to see all of the progress.

Talking about AI for science, though, this new era of discovery can go into hyperdrive with quantum computing. To simulate the world, you have to fundamentally break free from the limitations of the von Neumann architecture and classical computing. And for that to happen, we need reliable qubits, not just noisy qubits that are currently available. And our Azure Quantum provides that unique virtualization technology that can extend any type of qubit, detect errors in physical qubit, correct them and reliably help compute all over them.

In fact, earlier this year, our partner, Quantinuum, and us, we achieved a first-time ever record of reliable logical qubits; a huge, huge milestone for the industry. And then just a couple of months ago, we went further. We announced 12 reliable logical qubits. That was the record. And today, I'm really thrilled to announce yet another milestone, this time with Atom Computing. We just doubled the previous record, creating a machine of 24 logical qubits.

(Applause.)

I mean, to just kind of put this in perspective, these logical qubits are all entangled, making it the foundation for the world's most powerful quantum machines. And to give you a sense for why this matters, if you add 100 of these reliable qubits, you will have scientific quantum advantage, could be achieved. And so, and that will unlock, as you can imagine, the computing power to go solve some of the most pressing challenges we have.

Discovering these solutions, or discovering these types of workloads, or running these type of workloads on classical computers is obviously impossible. And so, therefore, we are excited about these next generation of quantum computers that go beyond today's noisy qubits.

We are building, in fact, to that end, a first-of-a-kind, commercial offering with our partner, Atom Computing, which will combine our science solution in Azure and our quantum computing platform, that virtualization layer I talked about, with Atom's quantum hardware into a full discovery suite for companies, as well as labs. This is going to be available to really accelerate scientific discovery.

I want to close out by talking about sort of our core mission. This was a preview of what you'll see throughout this week, across all of the three platforms, but at the same time, our mission is to empower every person and every organization on the planet to achieve more, one community and one country at a time.

As we enter this middle innings of AI, it's up to us to empower human achievement. That's why, over the past year, we have helped train over 23 million people in AI and digital skills. And we are very committed to working to help millions more learn how to use AI, because learning these skills will change people's lives. And in fact, it already has. Let's roll the video.

Thank you all very, very much. And have a fantastic Ignite.

(Applause.)

ANNOUNCER: Please welcome to the stage Judson Althoff, EVP and Chief Commercial Officer, Microsoft, and Lance Braunstein, Head of Aladdin Engineering, BlackRock.

JUDSON ALTHOFF: Good morning, everyone. Welcome to Ignite. Satya used this frame earlier in his keynote when he talked about computers not having a lot of empathy in AI, and he talked about enriching employee experiences, better connection with customers, building out AI-first business processes and bending the curb on innovation.

I'm delighted to be up here with my friend Lance to talk a little bit about BlackRock's journey and the context of what you've been working on. We've been on this journey together, Lance, now, since 2020, as partners from on-prem to cloud to digital and now AI, so I thought maybe it'd be good for you to share a bit of what that journey has been like inside of BlackRock.

LANCE BRAUNSTEIN: Amazing, and thank you for having us, and welcome to Ignite. So look, our vision, five years ago, when we first started on this journey was quite simple, to get out of the data center business. Like, we are an investment firm, and we build the pre-eminent investment platform, Aladdin. If you're not aware of Aladdin, it is an investment platform that supports portfolio construction, order and execution, all things post trade. It is a pivotal tool for investors. And what we don't do for a living is data center operation.

So five years ago, we went on a journey to explore the opportunities to outsource, and we found Azure to be the best of breed. And so in the intervening years, we have migrated Aladdin, which is a SaaS, and so we've migrated all of our clients onto Azure, and it's been a profound unlock. We completed that journey just this year, and it has made us better, stronger, faster.

So what does that mean? Better, stronger, faster is we believe that we get to market quicker. What used to take quarters now takes weeks. We believe we're more secure. We encrypt all data at rest, in transit. We certify all services. And importantly, when you think about the business that we're in, where we're really driven by market, the idea that we could survive and manage through stress tests like index rebalances, there's one that's super important to us, and we have actually found that on Azure, we've been able to do that much more effectively, and we are seeing profound improvements in performance.

All of this has been in close partnership, and not always the easiest client. I suspect Judson, Satya and Scott would say that, but I think we've made each other stronger. When I think back on the journey that we've been on, I believe that Azure and Aladdin are stronger for it.

JUDSON ALTHOFF: Yeah, super. We appreciate that, and agreed, the partnership has gone both ways. You've pushed us to make Azure better for all of our customers around the world, but it's not just been about transforming Aladdin and taking that to market together. It's also been about your internal transformation. You've used products like our Dynamics CRM offering and our Copilots to help enrich the employee experience at BlackRock. Maybe share a bit about what that's been like.

LANCE BRAUNSTEIN: Yeah, it's been a journey, and I was going to make a bad joke about transformer being transformative. I won't. Let's take it as read that I did that, and it was really bad. But you know, our employee experience has really been informed by two areas of investment, one in productivity, and the other in driving investment signal, what we call alpha.

I'll start with productivity. As Satya said earlier, we believe in a Copilot for every job family. Every job family should be augmented by an assistant that is driven by a large language model that gets you to a first draft of an email, a pitch or code, much faster. And in that spirit, we've implemented MS365, and we've implemented GitHub Copilot, and to real profound impact.

What I would say is the exciting thing for us is this is just the beginning. This is as bad as those tools are ever going to be. They're only going to get better. And so we are so excited to be on the journey to continue to see that productivity gain. On the investment signal it's all about driving more data into the investment life cycle with richer models.

And driving alpha to us means finding signal in the data, so if you're a relationship manager, ensuring that you could find the right product at the right market moment to put in front of the right client. That's alpha, so driving that signal is really profoundly supported by Azure ML and all of the foundational models.

And again, we're already seeing great benefit from working with our partner.

JUDSON ALTHOFF: So we both have a passion for bending the curve on innovation, and you're not just using AI inside of BlackRock, but also then infusing it into the product capability to deliver the market. Maybe share just a bit about that as well.

LANCE BRAUNSTEIN: Yeah, look, I'll zoom out for a quick second, and I realize we're coming to time, but the three pillars of our AI investment are productivity, alpha, and we talked about those, and then client engagement. We think we could profoundly bend the curve as you describe it, in our client experience.

And how are we doing that? We have always delivered our investment platforms through a set of user interfaces and application programming interfaces, now delivering a set of natural language interfaces, and what that means is you can have a conversation with Aladdin. You can have a conversation whereas you used to navigate user interfaces.

So imagine, last year, if you had the question, what is my exposure in my portfolios to Silicon Valley Bank, now you can just ask that question. You can ask the platform that question, and that massively reduces the barrier to entry. It makes your investment decisions much more efficient, and we're already seeing clients use that. We roll that out earlier this year, and again, built on the foundational models, Azure ML, and we're super excited by what the future brings.

JUDSON ALTHOFF: Awesome. Thanks so much for the partnership.

LANCE BRAUNSTEIN: Amazing.

JUDSON ALTHOFF: Really appreciate it. Thank you.

LANCE BRAUNSTEIN: Cheers.

JUDSON ALTHOFF: So we talk about these four pillars pretty consistently with our customers. We have thousands of referenceable customers in each of these categories. We released a blog earlier this week that has the real world evidence that talks about the ROI that our customers are seeing across industry.

We're really excited to be on this journey with you as well. We're not just explaining how we do this with our customers, we're doing it ourselves. Inside of Microsoft. Employing all of these capabilities has allowed us to generate more pipeline, have faster close rates, take hundreds of millions of dollars of expense out of our call center, and we're excited to be sharing all that in more detail with you here today.

Our next sessions are going to be awesome in terms of getting to the next level of detail on everything that Satya talked about. And we're going to start with Rajesh Jha in talking about our Copilot, the Copilot devices, and then flow into Scott Guthrie and Charlie Bell for our Azure platform and our security offering.

Please welcome to the stage Rajesh Jha.

ANNOUNCER: Please welcome EVP, Experiences + Devices, Rajesh Jha.

RAJESH JHA: Hello everyone. I'm Rajesh, and I lead our product team focused on Microsoft 365, including Office, Exchange, SharePoint, Teams, Windows and Surface. Thank you for your continued confidence in us. My team and I are committed to driving security and quality in all of these and we deeply, deeply appreciate the trust that you place in us in running these products.

Today, you heard Satya talk about three major platforms that we are building. I'm going to spend more time on Copilot, specifically Microsoft 365 Copilot.

Now, it has been one year since we made Microsoft 365 Copilot available. Usage is growing. Customers are rapidly deploying and adopting it. Thank you to the IT pros, support folks, engineers, our partners, CIOs and to this technical community here at Ignite. We could not do this without you.

So when we launched Copilot, we integrated it into the applications, Word, Excel, PowerPoint, Outlook and Teams that hundreds and millions of people use every day, giving people AI right in the flow of their work, and we continue to integrate Copilot across apps and services, from One Note to Stream, Forms to OneDrive, and more.

Now, another important call out is Copilot's deep integration into the Microsoft 365 application. Over the next couple of months, we will change the Microsoft 365 application icon to the new Microsoft 365 Copilot icon. You will continue to see us bring more and more Copilot capabilities and native AI to this application. This truly makes the application the front door to Copilot.

Now, what makes the Copilot unique is the grounding that it has. First, our Copilot understands the web. Second, our Copilot is crucially grounded in a user's work context, who do they work with? What do they work on, their meetings, their email, their conversations, their documents, and that is your organization's Microsoft Graph.

But more than an individual's data, IT can now connect the organization's business application data to the graph so the Copilot can reason over line-of-business data as well. We have hundreds of connectors already, and we are going to add many more.

Now, this comes up often, and so – I mean, it's important to note that Microsoft handles the Copilot's prompts, responses and data flows no differently than the way we handle your emails or documents or meetings. Microsoft is the data processor and your data is not used in training models.

As Satya spoke, based on your feedback, we continue to rapidly evolve Copilot, adding features both for end users, but also for IT for better governance and management. And now, with GPT-4o and enhanced orchestration, we continue to dramatically improve performance. Copilot responses are now more than two times faster on the average, and response satisfaction has improved by nearly three times.

Satya showed you the Copilot in the Office applications and some of the great features that are available and coming soon, but let's talk about what else is new.

We use a variety of models in Copilot depending on the specific scenarios, and today, we are expanding Copilot's capability in Teams meetings by leveraging one of these multi-modal models. We know online meetings can be much, much richer than the old school content calls, because insights and actions happen, not just in the spoken meeting, but in the chat and in what's being shown on the screen.

Copilot can now understand, recap answer questions based on visual content shared on screen, and that content can be a Word document, an Excel document, a web page or web application in addition to these spoken words in chats.

In September we introduced a new way of working with Copilot: Web, Work and Pages.

Pages are what we consider the first new digital artifact of the AI age. You can think of them alongside traditional documents, traditional documents and Pages. Pages are dynamic, persistent canvases that are designed for collaboration between humans and AI. They turn your Copilot interactions into durable work assets that are built to be shared and co-authored.

Pages are in depth with the rest of your Microsoft 365 data making them searchable. Is this a Word document or Excel document. And because these are all on SharePoint, the same policies for content management and compliance that you have in place apply to Pages as well.

We will continue to innovate here, and today, we are announcing that, in BizChat, you will be able to create rich artifacts from board blocks, Mermaid diagrams, tables and charts. And so, with a simple prompt, you can find data buried in the graph and transform that into a beautiful, interactive, and also collaborative piece of content, and then the

formatting and the interactivity is retained when you transform this Copilot response into a Page.

Now, look, these demos give you a sense of just how powerful Copilot is, but people have told us they want the Copilot to do more on their behalf, even delegate entire tasks to Copilot. So I'm excited to announce and introduce Copilot Actions. These are simple but powerful ways for you to scale and automate what you and your Copilot can do. Let's take a look.

Copilot Actions can be accessed directly in the Microsoft 365 app. You can create a new action from scratch, or if you're looking for inspiration, discover from a series of pre-built templates designed to help you. You can catch up on things that you've missed. You can gather information from your colleagues or simply to create new content.

Now, let's say that you need to create an aggregated weekly status email. Working with various stakeholders to provide updates, it can be time consuming. Now Copilot Actions does this for you. the templates are intuitive. You just select the parameters you need, add the context and choose how often it should run. You preview it, you make any final tweaks and your actions are ready to go.

Copilot will then reach out to your colleagues on your behalf, asking for those status updates. It even follows up with anyone who didn't respond. Copilot then generates a consolidated email for you, for you to edit and send it, and so with just a few clicks, you offload a pretty time consuming weekly process to Copilot and think about all the processes that you can use Copilot action for. You'll be able to manage all your AI automations in a single pane of glass where you can run, schedule or turn off the actions.

So as you can see, Copilot is not just your personal AI assistant, but it's becoming your portal. It is becoming your user interface to the breadth of AI capabilities that are available. Copilot Actions is one aspect of that.

Now, since the very start, we've been clear that the Copilot is individual. It is your Copilot. It advances your personal agency. It is grounded in your unique work context, your permissions. We think of agents as working alongside your personal Copilot to transform business processes. Copilot helps you. Agents can help everyone. Agents will tackle the process and workflows. That is how work happens. Whether these are mundane like expense reports, or the more strategic like financial reporting.

These agents will often be built on top of existing content. There were many enterprises that content resides in SharePoint. SharePoint. Over 2 million sites are created, and over 2 billion files are uploaded every day to this SharePoint service. And so, think about the knowledge that this represents for your organization and the value of having Copilot reason over your SharePoint.

So today, we are announcing that every SharePoint will have a ready-to-use agent. You can easily ground agents in SharePoint and relevant files, sites and folders. And the

SharePoint agent inherits all these permissions. It is your expert on all the content on the site, not just to find documents, but also to answer questions.

Now, you can also create a custom agent scoped to a specific content with a single click, and then share that with others. And then, when you add that SharePoint agent to a Teams chat, anyone can @ mention the agent and ask a question. And so, the agent is contributing, just as a teammate would, fully respecting the organization's security and privacy and permission controls. So if you really want to unlock insights, and you want to unlock knowledge and the IT that resides in your organization's SharePoint, this is the easiest way to get organizations, getting the value from Copilot on that knowledge.

So now I'm going to hand it over to Charles. But let me summarize my thoughts quickly. First, thank you again for the trust you placed in us with the rapid adoption of Microsoft 365 Copilot.

Second, we believe that every employee needs, deserves and wants a Copilot. And let's be real for a minute. If you haven't already deployed Copilot, do you really think your employees are not already using consumer GenAI tools at work? I think we know the answer to that question: with Microsoft 365 Copilot is how you can empower them safely, securely and compliantly.

And thirdly, every business process will have an agent, and many of these agents will be customized and built in Copilot Studio. Charles, tell us more.

CHARLES LAMANNA: Thank you, Rajesh.

As you've heard, Copilot is transforming personal productivity and the way that people work with AI. With Copilot and agents, we're also going to change how companies work with AI. An AI-first company starts with people, with every employee augmented by Copilot as their personal assistant, and it also consists of agents that transform every business process and even complete tasks autonomously.

And today, we are excited to introduce the Copilot Control System, designed to help every IT team confidently adopt Copilot and agents. This layer helps you stay secure and compliant by enforcing permissions, creating data governance policies, managing risk and even measuring impact.

So an AI-first company has all these components. Copilot for everyone. Agents for every business process and all operating within the control system. And one of the more magical parts of how this all comes together is the seamless collaboration between people and AI. And today, we're announcing Interpreter, a new pre-built agent in Microsoft 365 designed for deeper collaboration, available directly in all of your Teams meetings. Interpreter provides real-time speech translation during meetings, helping everyone easily overcome language barriers. You can even have Interpreter simulate your speaking voice for a more personal experience. Let's see how it works.

SPEAKER: (First in Spanish, then quickly the English below.)

We need to review the progress of the project. Can you please tell us how the tasks we should take are going?

CHARLES LAMANNA: Isn't that pretty cool, to see it work end to end?

And we are extending agents across core business applications like Dynamics 365. Just last month, we introduced 10 pre-built agents in Dynamics 365 to help you automate business processes. These agents can do everything from automatically qualifying sales leads, engaging suppliers and collaborating with service reps to resolve customer issues.

And today, we are also introducing the employee self-service agent in Copilot. With this agent, we are reimagining the core employee experience, starting with HR and IT processes. This agent comes prepared in Biz Chat, providing a streamlined experience that helps employees get real-time answers to workplace policy questions and even acts on key tasks like applying for leave or getting IT support.

And this agent, it connects to your organization's knowledge in SharePoint and Teams, plus data from Workday, SAP and even ServiceNow. And all of this is configured in Copilot Studio using Microsoft-built connectors.

And across Microsoft 365 and Dynamics 365, these pre-built agents help you quickly boost collaboration and automation. But we know that you also want to use solutions from other companies right in the flow of work. That's why we're thrilled to partner with several companies to ensure their agents are natively available inside Copilot.

The Adobe Marketing Agent empowers marketers with campaign planning and measurement tools, workflow optimization and scaled content generation directly within Microsoft 365 apps and Copilot. This helps marketers where they work, driving impact and efficiency for all of their brands. SAP Joule integrates with Copilot and enables employees to streamline their workflow. It provides unified, two-way access to both Microsoft 365 and SAP business applications, enhancing productivity for everyone.

And ServiceNow's Now Assist integrates with Copilot to streamline workflows and enable self-service capabilities within Microsoft apps. This combines two AI assistants to better understand your user intent, drive automation and boost productivity where employees work.

And Workday's Copilot integration, it leverages your Workday and Microsoft data to help managers with performance reviews or streamline HR tasks, or even enhance employee productivity through intelligent suggestions and insights.

And we also are seeing AI-native companies like Cohere adding their agents too. The Cohere agent for Copilot integrates powerful language models, multilingual RAG and grounding technology to provide secure, customizable access to unified data insights

across your entire organization. And all of this is just the beginning. Copilot is your one-stop shop for accessing all these incredible agents from all these incredible companies.

But the reality is, it's not one size fits all. Every company has unique processes, unique data and unique workflows, and we are going to support agents that you create for your own processes. And this is where Microsoft Copilot Studio comes in.

Copilot Studio is a low-code platform that makes it easy to create your own agents. You can chat with these agents or even have them running autonomously in the background. And with Copilot as your UI for AI, you'll be able to work with all of these agents in one place.

Let's see how easy it is to use agents with a real example from Thomson Reuters. They provide expert solutions for legal, finance and other regulated industries. And today, we'll focus on their services for acquisitions, where the assessed value and risk through a lengthy and highly manual due diligence process. And with Copilot and agents, they've been able to improve their productivity by over 50%. Let me show you.

So here is a standard Copilot experience. I can ask questions about my calendar or my email, but something that Thomson Reuters has done, which is really incredible, is they've leaned into all these different agents to help with common business processes and workflows, and each of these can be quickly built and deployed to users securely using the Copilot control system.

So I'm going to click through into one of these agents called MergeMate. And what this does is it helps me identify any risks associated with an acquisition I'm working on. And as part of this, it doesn't just do a quick response, say, looking over a few pieces of information; there's an in-depth analysis of documents and business data. It even goes out to Azure AI search indexes to pull down the information that matters to me. And it can give me really insightful guidelines and suggestions for deeper investigation, like, for example, identifying compliance gaps due to inconsistencies in regulatory filings. And this is something that you can't get with just a standard model or internet data. You need your corporate enterprise data there too.

But what Thomson Reuters did is they didn't want to just settle for an agent that you chat with. They also adopted an autonomous agent, and they call this one Co-Counsel because it's the co-counsel to the lawyers, and it will run automatically in the background whenever a new data pack is uploaded to a SharePoint site. And what it's done is it's provided great findings, and it's done the initial review and the initial scan of all the information to save more time for those lawyers. That's where that 50% productivity improvement comes from.

And whether it's that chat based or that prompt based or the fully autonomous agent, it's all easily accessible right inside of Copilot. And in the case of both of these agents, it was built inside of Copilot Studio.

So here I am inside of Copilot Studio. This is how I can configure and enable these different agents. I can have things like instructions and knowledge and actions and triggers. Easy, configurable steps to get the agent I want. And if I look at this co-counsel agent, that autonomous agent. I zoom in here really quick. You can see that it has instructions about the steps and the guidance it should follow to complete its tasks. It even has a description explaining what it's supposed to do. So for example, it helps with that automated document review.

And this is just done in natural language. I don't have to write code, and I don't have to fine-tune a model. I'm able to just speak and describe and document this the same way I would with a new employee on the team. And one of the things which is so powerful about these agents is its knowledge. You're able to easily configure all of your enterprise data sources so these agents operate with the relevant pieces of information you care about, safely and securely, respecting all the security permissions you've already enabled. This way, you can put your company knowledge and investments to work to help automate all these different processes.

And you can even do things like actions or tools so it can go make changes and take additional steps. And all of this, very easily configured right inside of Copilot Studio.

Now, as you can see, Copilot and agents transform how companies work, making them AI first. Copilot empowers everyone, acting as a personal AI assistant. And Copilot Studio is how you create agents to transform your business processes, and it is all secured by the Copilot Control System.

Now, I'm going to turn it over to Nicole to share how companies are using Copilot and agents to drive real business value.

NICOLE HERSKOWITZ: Thank you, Charles.

As IT pros, you are at the epicenter of this AI transformation. And we hear you. It's not an easy place to be. The speed of innovation is unlike anything we've experienced before, and the stakes couldn't be higher.

As Microsoft, we're here to partner with you on the journey. Our job is not only to make tech accessible to end users, but also make it easy for you with a manageable, measurable and secure product experience that you can bet your career on, which is why we're excited to bring Viva Insights to Copilot. Introducing Copilot Analytics.

Let me tell you, this is the tool you've been waiting for. Measuring hours and minutes saved is great, but what you really need is the ability to measure dollars and cents. Copilot Analytics is your go-to tool for measuring your AI transformation across Copilot and all the agents you will create in Copilot Studio. This is what it means to demonstrate AI's ROI.

Now, let's take a look at an organization that's reimagining a process with Copilot. Dow

is a global materials science company that makes everything from the silicon on your smartphone to the foam in your mattress. Every year, Dow spends billions of dollars shipping their products to customers, and any invoicing errors can add up to millions of dollars, which is exactly why they turned to Copilot.

Dow used Copilot Studio to create two agents to tackle this problem. The first is an autonomous agent that reviews invoice data. It's on the lookout for errors, like when the billed amount is different from the contracted rate, or when there's duplicate invoices in the system. When it finds one of these situations, it immediately flags the issue for Dow's employees.

From there, the employees can use a second agent to investigate. This agent lets employees interact with invoice data in Copilot, all using natural language. Here they can easily take action, like approving or declining invoices. Just a few weeks into the proof-of-concept phase, Dow employees were using Copilot to analyze thousands of invoices, flagging issues and highlighting potential savings in real time.

Thanks to Copilot, Dow anticipates they will reduce freight expenditures by millions of dollars in the first year. Now, it's one thing for me to describe it. Let's hear from Dow how Copilot is completely transforming their business.

(Video segment.)

NICOLE HERSKOWITZ: It's not just Dow. Organizations across industries and around the world are transforming work with Copilot. Finastra's marketing team is creating assets in seven weeks that used to take seven months.

Standard Bank's Helpdesk Agent now resolves 99% of employee queries. McKinsey onboards new clients in 90% faster rates. PG&E agents support routine tasks across the business, generating \$75 million in annual savings. Visa is enhancing their technical documentation for customer support, increasing efficiency by 20%, engagement by 30% and customer satisfaction by three points. And Honeywell equates Copilot productivity gains to adding 187 full-time employees.

That's just the beginning. Nearly 70% of the Fortune 500 use Copilot to work smarter, serve more customers and deliver more value. And we're partnering with them as they look across all parts of their businesses to identify opportunities that are ripe for AI.

Now, Microsoft has been one of the earliest adopters of this technology. We're customer zero. Take one of our sales teams, which achieves 9% higher revenue per seller, close 20% more deals thanks to Copilot. Or, one of our customer service teams which is resolving cases nearly 12% faster.

Part of our marketing team is seeing a 21% increase in conversion rates, and our HR team self-service agent is answering employee questions, 42% with greater accuracy. It's only going to get better from here. Copilot Agents will continue to reimagine work and

business processes across every function, to scale employees' impact and ingenuity like never before. We look forward to partnering with you on every step of the journey. Over to you, Pavan.

PAVAN DAVULURI: Thank you. Nicole. Hello, everyone. Windows is the platform companies turn to first to transform their businesses, because we are focused on your top needs. First, delivering you the most reliable, secure platform on Earth. Second, working with you to develop a modern endpoint strategy with Windows 11. And last, helping you empower employees with AI in Windows.

Let's start with job No. 1, reliability and security. Today, we're announcing the Windows Resiliency Initiative to affirm our commitment to make Windows the most reliable and secure platform for all our customers. The latest release of Windows 11 reflects this commitment, building on a bedrock of a hardware security baseline, secure by default posture, and advances that have eliminated security instance by over 60%.

With this initiative, we're partnering across the industry with the U.S. Cybersecurity and Infrastructure Security Agency, or CISA, to set standards for safe deployment practices. We're helping partners move from kernel mode to user mode, delivering new, quick recovery features and continuing to invest in Windows to meet your reliability requirements.

We're also focused on helping you with the three largest security risks we face today. First, you asked for help from phishing attacks. So, we hardened Windows Hello with extended support for passkeys to protect your identities. Second, you wanted to better defend yourself from malicious apps drivers, so we're introducing Smart App Control, new innovations in App Control for Business and announcing Windows Protected Print.

Third, you told us to find a better way to balance user empowerment with local machine controls, so we're excited to announce Administrator Protection.

(Applause)

That's right. Organizations today have a binary challenge that is complex when setting up policies for employees. One option is giving administrator permissions to employees, which is friction-free for users, but grants users and malware automatic direct access to kernel mode and system security. Critical vulnerabilities in Windows have been affected and exposed in this way.

The other option is standard user permission, where malware is silently blocked from changing the machine configuration. This setup is more secure, but it is frustrating for users on a heavy lift for IT. When Administrator Protection is enabled, a user can make a system change that would have previously required admin rights.

For example, take this app installation. A user can now securely authorize this change using Windows Hello. Windows creates a temporary, isolated admin token for the

requested change to get the job done, and once that process is complete, the admin token is immediately destroyed. This means admin privileges do not persist. Administrator Protection is currently in testing with Windows Insiders and will roll out more broadly in 2025.

These immediate steps are just the beginning. We will keep listening to you and deliver more improvements into the first half of next year. We will introduce a new process to verify credentials for any third party to submit drivers to Windows. We will launch a new, secure and connected recovery environment so customers can quickly get their systems back up and running.

With your input, we will keep innovating to make Windows the most secure and reliable platform, continuously earning your trust. Next.

The best way to benefit from these advancements in security and reliability is through modernizing your endpoints with Windows 11. We know this is a significant undertaking, and we're grateful for the over 1 million organizations that have or are actively transitioning to Windows 11, ahead of Windows 10 end-of-support in October 2025.

Customers already on Windows 11 tell us it delivers, with increases in employee productivity and satisfaction, reduction in IT costs and improved management tools, fewer helpdesk tickets, and decreased security incidents, delivering an amazing 250% return on investment for a new Windows 11 PC according to Forester. Thank you.

Today, we're announcing two key features launching this spring to help with the move. First, Windows Backup for Organizations with Entra ID is coming to public preview, enabling smooth transition for your employees. Preferences like pinned apps and desktop backgrounds transfer automatically, reducing setup time and easing IT workloads. And, we are thrilled to announce Windows Hotpatch is coming to Windows 11 and Windows 365, ensuring day-one protection by delivering Patch Tuesday Updates instantly through Intune's Auto Patch with 65% fewer restarts.

(Applause)

Right? Kind of exciting. Windows Hotpatch will accelerate security patch deployment by 60% without interrupting workflows. And it's not just your PC fleet that gets better from the benefits of moving to Windows 11, your VDI endpoints will, too, with improvements in Windows 365.

You asked for a simpler, one-click way to access Windows and the cloud on Windows 11, Mac OS, iOS and now Android, so we built Windows App You asked for support with frontline workers and for creators and developers who need high-performance GPUs and mixed reality. So we added new SKUs to better support your employees' needs.

You asked for us to make it super secure for BYOD and temporary worker scenarios, so we've added Mobile Application Management, or MAM, controls for iOS and Android devices.

You asked for simple, purpose-built devices with no local data that boots straight to the cloud. So, we expanded the Cloud PC categories, starting with Windows 365 Link, that Satya showed earlier.

Just like we did with two-in-ones, we're starting with first-party devices, doing the work to get the experience right, to then scale with our OEM partners in 2025.

Thank you for the work you are doing to modernize your endpoints. In doing so, you and your employees will experience the productivity, security, accessibility and return on investment benefits of moving to Windows 11.

Last, securing and modernizing your endpoint sets the foundation to empower your employees with the best platform for productivity, Windows 11. We will bring Microsoft 365 people, files and calendar to your task bar, so your graph data is just a click away.

And we're building Windows to be the platform for AI. Copilot+ PCs are not only the most intelligent, performant and secure PCs you can buy, they were built from the ground up to extend Microsoft's Copilot Stack down to the client.

We co-engineered silicon with all major silicon partners, including AMD, Intel and Qualcomm, to deliver market-leading neuro processing units with over 40 TOPS of local AI inferencing.

This new silicon, combined with our OS and the Windows Copilot Runtime, offers the world a platform to bring their best AI innovations to Windows. From Anthropic, Hugging Face, Meta, OpenAI, to developers like Adobe and LiquidText to countless more who are already taking advantage of this new platform.

And of course, Windows will continue to be the best platform for Microsoft 365 Copilot. Copilot+ PCs complement AI in the Cloud by adding real-time client capabilities to improve your employee productivity. Copilot+ PCs deliver the best experiences for Teams, WhatsApp and Zoom by pairing AI in the Cloud with local capabilities of Windows Studio Effects with super resolution.

Only Copilot+ PCs will deliver natural language search across Microsoft 365 and Windows, by using the same semantic index to unlock natural language search across OneDrive in the Cloud and files on your PC.

And finally, Copilot+ PCs will accelerate productivity with Click to Do, which connects a deep understanding of your screen and your Microsoft Graph with text and image action.

Windows will continue to listen and innovate to deliver on your top needs by building you the most secure and reliable platform on Earth, by helping you take advantage of the return on investment of moving to Windows 11 and empowering your employees with the best platform for productivity. Thank you.

(Applause)

RAJESH JHA: Thank you, Pavan, Nicole, Charles. We have an extraordinary opportunity together in a world where Copilot empowers all of us as individuals to do more. Agents transform all of our business processes to do more, and together, Copilot and agents empower people and organizations to do more.

Thank you for spending your time with us here today at Ignite. Thank you for your partnership, and most importantly, thank you for the trust that you place in us every day. Over to you, Scott.

SCOTT GUTHRIE: Thanks, Rajesh, and good morning, everyone. It's great to be in Chicago at Microsoft Ignite. AI is going to profoundly change how we work and how every organization operates. Every existing app is going to be reinvented with AI, and we're going to see new apps being built with AI that weren't possible before. Earlier, Satya touched on the three major platforms that Microsoft is building to maximize your opportunity in this era of AI, and in my talk, I'm going to go through the Copilot and AI stack, which is the most advanced platform for creating AI solutions, and share some of the amazing new capabilities that we're announcing here at Ignite.

Let's start by talking about our developer tools. Copilot Studio, which we released at Ignite last year, is now used by more than 100,000 organizations to build AI solutions. The Visual Studio family of products, which includes both Visual Studio and Visual Studio Code, are now the most widely used developer tools in the world, with over 40 million active developers, and GitHub is now used by over 100 million developers and is literally the home of open source. As you've heard from both Satya and Rajesh, Copilot is your UI for AI, and Copilot Studio enables you to build agents that seamlessly integrate within it.

With Copilot Studio, you have a single place to build agents that are tailored to your specific business needs. The agents can work on your behalf and even run autonomously in the background. As you saw in Charles Lamanna's demo a few minutes ago, it's easy to reimagine an entire business workflow using it. Copilot Studio provides more than 1400 pre-built plugins and connectors to integrate all of your data, APIs and other third-party services. What makes Copilot Studio even more powerful is the ability to also leverage Azure capabilities in your Azure APIs and apps as well.

Now, not only can you create custom agents, but you can enrich data further by integrating with your data and custom RAG solutions built with Azure AI directly into your agents. This enables your agents to be more grounded in your business data and provide specific, high-quality responses and business workflows. We're also excited this

week to announce a new way to buy Copilot Studio. You can now use Copilot Studio without having to buy any user licenses. You can instead optionally pay for Copilot Studio agents and workflows using a pay-as-you-go consumptive billing model, available through your standard Azure subscription. This pricing model enables organizations to build and deploy your own Copilot experiences without any upfront costs. IT still remains in full control using this new pricing option, and IT retains the ability to manage which employees inside an organization have access to use it.

Now, let's take a look at some of the deeper innovation that we're driving with GitHub and Visual Studio. GitHub provides developers and enterprises an integrated, AI-powered, end-to-end developer platform that can be used to collaborate, automate and secure DevOps solutions. GitHub Copilot was the first Copilot solution that we built using transformational large AI models, and it provides an AI pair programmer that works with all popular programming languages and dramatically accelerates developer productivity. These statistics speak to the productivity benefits that GitHub Copilot provides.

Developers are now writing code over 55% faster with Copilot through real-time code suggestions and solutions, 85% of developers feel more confident now in their code quality, and 75% of developers say they feel more fulfilled with their job, enabling them to focus on more meaningful tasks. I love this last statistic in particular, and it really shows the true potential that AI can have to improve all of our lives. Over 70,000 organizations use GitHub Copilot today to supercharge the productivity and the developer experience of their employees.

At Airbnb, approximately 70% of developers are using GitHub Copilot on a weekly basis, Shopify estimates that they accept up to 25,000 lines of code written by GitHub Copilot every day, and Microsoft engineers using GitHub Copilot have already accepted over 79 million lines of code that's been produced by GitHub Copilot this year alone.

We're continuing to bring new innovations to GitHub Copilot. Developers in VSCode can now choose from an array of industry leading models, including not only OpenAI's GPT-4o, but also Anthropic's Claude Sonnet 3.5. They can also now use natural language chat to implement complex code changes across multiple files. And when they're on the path to merge a pull request, developers can now use the new code review capability and request Copilot to do the initial code review while waiting for a human reviewer. All of these improvements really provide developers with AI-enabled capabilities that improve the experience across the entire software development life cycle.

Now, for us, security is job No. 1. And one of the best ways to keep applications safe is to ensure everything is secure from the start. GitHub Advanced Security is built for developers and uses AI to help them build apps more securely. GitHub Advanced Security automatically analyzes source code for vulnerabilities, explains to developers areas of the code that look risky and provides the code to fix the problems. Let's watch a video of GitHub Copilot and GitHub advanced security together in action.

VIDEO DEMO: GitHub Copilot is the most used AI developer tool in the world. Since its release, it's been a useful addition to developers everywhere, helping generate code from comments and saving time, but it's grown so much smarter since then. I want to add an about page to the Contoso wireless app, similar to the support page. In VSCode, I can ask Copilot chat where to find the routes that are defined. It uses intent detection, identifying that it needs to look across the whole codebase, and because the code has been indexed on GitHub, it finds the relevant file quickly.

I'll probably need to create some new files, as well as add a few existing ones to add to my about page. So, let's use the new Copilot edits feature to help make changes in several places at once. I'll ask for the changes I want and reference any relevant files. And just like that, GitHub Copilot has made the changes I need across multiple files with just one request. We can see all of the changes requested in one location, and yeah, they all look good to me. I'll save them all and accept those changes, and there we go, a new about page.

GitHub Copilot has long leveraged different large language models for different use cases, but in the past year, we have experienced a boom in high-quality models that excel at different programming tasks. We are now adding developer choice, giving you the flexibility to pick between the latest OpenAI models as well as a range of industry models. Here I can ask Anthropic's Claude Sonnet to explain some of the code in our project, and it can quickly get me up to speed. Perfect for when you're working with existing code bases.

We are embedding the power of Copilot all across the development workflow to help your team at every step. With GitHub's advanced security features, you can automatically find vulnerabilities in your existing code thanks to code scanning. But better than that, you can now ask Copilot to turn found into fixed. Here, we could see where GitHub Advanced security has automatically identified a number of security vulnerabilities. Copilot will look into the issue and the information in the GitHub vulnerability database to generate a proposed solution. When I click "generate fix," Copilot generates the suggested code needed to fix the issue for me, dramatically reducing the time and effort needed. You can then submit that code for your engineering team to review and test, dramatically reducing the time it takes to find and fix security issues.

SCOTT GUTHRIE: The combination of GitHub and Azure is incredibly powerful, giving developers an end-to-end experience to build, scale and deploy secure software to the cloud. With GitHub Copilot for Azure, we're making the entire process even more intelligent and removing friction across the entire development lifecycle. Let's watch a video of GitHub Copilot for Azure in action.

VIDEO DEMO: Let's look at how GitHub Copilot for Azure can help us build a new AI application on Azure. GitHub Copilot for Azure is an extension for GitHub Copilot chat that you can use in VS Code. I want to build a chat application that uses GPT-4o. First, GitHub Copilot for Azure can help me learn about the AI models that are available on Azure. Next, I'll ask Copilot to help me build a chat app on Azure. Copilot recommends

a template from the template gallery. The new AI App Template Gallery is a great way to discover new resources for building AI applications. Each template includes application code and infrastructure as code files needed to deploy an app on Azure. You can find a range of templates across multiple programming languages and frameworks. GitHub Copilot for Azure gives me guidance on how to use and deploy the template to Azure quickly using the Azure Developer CLI with its developer friendly commands, keeping me focused on the code.

Now, even developers that are new to Azure can access the full power of Azure and give their app the best environment to run their app. GitHub Copilot for Azure understands the context of the code I'm working on and of my Azure environment and subscription. The Azure extension can answer questions about my Azure resources and help me troubleshoot issues. Here we can see the resources that are created when I deploy this app in the Azure portal, including the Azure OpenAI resource and the container app where my app is deployed. In just a matter of minutes, I have provisioned the infrastructure and deployed application code using built-in Copilot intelligence and automation to simplify my path to Azure. And after the application has been running, I can also use GitHub Copilot for Azure to ask questions about my costs to find ways to ensure my app is running optimally in production.

SCOTT GUTHRIE: Now that we've covered some of the innovation happening with our developer tools, let's now talk about Azure AI. The innovation in Azure AI is already powering all of Microsoft. All of our Microsoft Copilots are built on top of the same AI platform that's available to all of you, and this means that you benefit from the battle hardening that we've done to support tens of millions of users and organizations with AI around the world. Customer adoption of Azure AI is accelerating, and we already are helping more than 60,000 companies around the globe achieve real business impact with it. These are just a few of the great companies with solutions live on Azure AI today.

Now, the data in the AI landscape is incredibly crowded and pretty complex, and you can see this eye chart of some of the different solutions out there. When you see the huge number of technologies and vendors and the huge number of options that organizations need to grapple with, it can be challenging, and the complexity is only increasing. We see that 93% of use cases are now using three or more AI model providers, and 80% of early AI projects often fail to meet expectations due to this complexity. To help organizations simplify and scale their AI workloads, we're excited to announce the availability of Azure AI Foundry.

Azure AI Foundry enables organizations to design, customize and manage the next generation of AI apps and agents at scale. Integrated with the most popular developer tools in the world, GitHub Visual Studio and Copilot Studio, Azure AI Foundry brings together Azure AI models, tooling, safety and monitoring solutions to help you efficiently and cost effectively design and scale your AI applications. And as Satya mentioned, we really consider this a first-class app server for the AI age.

An important part of AI development starts with having access to the latest and greatest models. Azure OpenAI service is one of our flagship products. It's available in more than 28 regions around the world and offers data residency in every one of those regions. We now provide limitless throughput with the latest OpenAI AI models, enabling large customers to scale as high as they need, and we offer flexible deployment options, including batch off-peak capacity, which enables you to use AI in off-peak times and benefit from an additional 50% price discount. Our Azure OpenAI AI service delivers advanced enterprise security controls and compliance guarantees, and we back the service with industry leading SLAs.

Azure AI Foundry offers the most advanced and complete set of AI solutions of any cloud provider. We support a comprehensive set of frontier and open-source models, so you can use a variety of different AI models for your solutions, and the Azure AI model catalog now includes more than 1800 Foundation open-source task and industry models from OpenAI, Meta, Mistral, Cohere, Nvidia and many more.

And as AI continues to transform industries, the ability to fine tune models and customize them for specific use cases has become more critical than ever. We already offer end-to-end fine tuning for Azure OpenAI service, Meta, Mistral and Phi-3 models, and today we're excited to announce fine tuning collaborations with Weights & Biases, Gretel, Scale AI and Statsig, which makes fine tuning, data preparation and experimentation easier for organizations of all sizes.

We're also excited to announce the agent service to help developers build, deploy and scale enterprise AI-powered agents to automate a variety of business processes. And with a rich ecosystem of models from the model catalog data sources such as Fabric, SharePoint and Bing, and more than 1400 action connectors in Azure Logic apps, the Azure AI agent service simplifies how you build agents and microservices.

We're also making a bunch of advancements to our Azure AI content safety service, including adding risk and safety evaluations for images. This adds to the extensive security and safety capabilities that we've already released, including the ability to detect and correct hallucinations and mitigate security threats. This integrated tooling makes it easy for you to test and evaluate your solutions for safety, security and responsibility. We're integrating all of this inside GitHub. We're bringing the power of Azure AI Foundry directly to developers using the tools that they live and code in every day. Let's take a look at this comprehensive end-to-end AI developer toolchain now in action.

VIDEO DEMO: We've put our models directly where 100 million developers live, in GitHub. Let's have a look. Developers can explore a library of the most recent and popular models from the Azure AI Foundry. Every model is unique in its capabilities, performance and cost. You can click into the model cards to see what each model specializes in. Mistral's new Ministral model offers low latency, for example, while GPT-4o is currently the most powerful and widely used multimodal model. You can now test drive these models for free with the interactive playground and try out your prompts before you spend a dime in production.

Let's check it out by taking another of the latest models from OpenAI, GPT-4o mini. You can just upload a photo and see how the model will respond. You can even compare one model to another with this compare button. Say you're building a multimodal app using text and vision. At no cost, you can easily compare GPT-4o mini with Meta's new llama 3.2 vision model to help you decide which one is better for your use case. But for any real application, the model alone is not enough, it needs your data. That's where RAG comes in.

Now, for the first time, you can ground a model to your data right here in GitHub just by clicking the "Upload Files" button. Then you can compare the grounded model to the ungrounded one to see how adding data will improve your app's ROI. When the model's answers are in line with your business goals, you can just click the code button and get the code right here in GitHub, copy it into VS Code, and with your GitHub credentials, you can run it straight away. The key to success with your AI app is choosing the right model and the right prompts to serve your users.

Until now, you can only evaluate those variants pre-production, but now, for the first time, we're offering model experimentation in GitHub Copilot, which lets you run live prompt and model experiments on apps in production to improve your business outcomes. Here, Copilot has made a thumbs up and thumbs down button, and it's stubbed out a few variants, which can be different prompts or even entirely different models. You can assign each variant to a percentage of your user population. When users use the app, they'll get one of these prompt variants and the results will be recorded right in GitHub.

Now you can see how the variants stack up to each other on metrics like the number of chat users, the cost to run each variant, how long it took each one to respond, the safety and error rate of each variant, and direct user feedback. With model experimentation, you can make a data-driven decision on how to best serve your users. Once you've tested your app in GitHub and are ready to scale, you can get a production key and that brings you to Azure AI Foundry.

Here you can continue to work with models and prompts. You can customize models for your business with fine tuning, you can build and test AI agents, and you can provide your agents knowledge from sources like Bing web search, unstructured data in Azure AI search, or analytics data in a Fabric lakehouse. And you can create agents that take actions with Azure Functions and Azure Logic Apps. Then you can evaluate your AI application to ensure it's behaving in a responsible and trustworthy manner. Finally, you can experiment with the new real-time models in Azure AI Foundry.

You sure can!

So, what do you think of Azure AI Foundry?

Azure AI Foundry is the platform to deploy your AI applications at scale.

You said it GPT!

SCOTT GUTHRIE: Industries around the globe are already using Azure AI. From automotive to retail to manufacturing and finance, Azure AI is helping companies redefine processes and services to their customers.

A great example is the NBA. The NBA collects an incredible amount of data on players' on court performance like speed, dunk, height, and even injury risk. And with Azure AI Foundry, the NBA can now synthesize all this data and create experiences that delight fans both online and, in the arena, with real time stats and insights directly on their devices. And we're working with healthcare company McKesson to drive breakthroughs in cancer care with the power of data and AI. Let's take a look at some of the amazing work that they're doing today.

VIDEO DEMO: Early cancer detection is critical to improving treatment outcomes and increasing patient survival.

Ontada, a McKesson company, is providing innovative technologies at the point of care.

Manually dealing with unstructured healthcare documents is time-consuming and error-prone, but are where important insights can sit.

Azure OpenAI Service batch API was the solution and through this partnership we were able to reduce our processing time by 75%.

We have implemented the latest GPT models to target nearly 100 clinical oncology data elements across 39 cancer types.

This allows healthcare providers to identify key cancer attributes like tumor site, comorbidities and clinical stage. Azure AI opens up computing power and transforms an estimated 70% of previously unanalyzed unstructured data.

Our strategic partnership with Microsoft allows us to advance AI-driven oncology research and drug development.

We have access to richer data at much faster speeds, which allows our life science partners to obtain meaningful insights much more quickly, driving treatment adoption and positively impacting the lives of cancer patients.

SCOTT GUTHRIE: Now that we've covered Azure AI, let's talk about our data platform. All of the AI innovation that you've seen so far is built on a foundation of data. Data really is the fuel that powers AI. And the Microsoft Intelligent Data Platform is an integrated suite of Azure data and analytic services, designed to help organizations manage and utilize their data effectively. The platform supports a range of different deployment options, including infrastructure as a service, platform as a service and fully

managed SaaS environments. Integration with Azure AI Foundry facilitates the development and deployment of AI applications. And Microsoft Purview provides data governance and compliance, ensuring consistent management across the entire data estate.

As you heard from Satya, we're investing in data innovation across this entire platform. SQL Server 2025 now in preview is an enterprise AI-ready database from ground to cloud. This new version of SQL Server has AI built in. It simplifies AI application development and RAG patterns with secure, performant and easy-to-use vector capabilities, leveraging the familiar T-SQL language. SQL Server 2025 continues its best-in-class security and performance, and in this new version, you can now automate credential management and reduce potential vulnerabilities with Microsoft Entra managed identities. You can also boost workload performance with enhanced intelligent query optimization and query execution.

Additionally, SQL Server 2025 is fully integrated with Azure and provides cloud security management agility anywhere using Azure Arc. Now, one year ago today, we GA'd Microsoft Fabric, our end-to-end developer or data analytics platform designed to help organizations power their AI transformation and reimagine how to connect, manage and analyze their data. Fabric is an all-in-one SaaS platform with AI-powered services to accomplish any data task, all in a pre-integrated and optimized environment so your data teams can work faster together.

Fabric enables you to integrate and work with your data regardless of where it's produced or stored. It even lets you seamlessly work with data stored in Snowflake, Amazon and GCP. Fabric is unified at the most fundamental level with unified compute and storage, unified governance, a unified data tooling experience and a unified business model. This unification dramatically accelerates your time to value and helps your organization save tremendously on costs. Companies like Chanel, KPMG and others are betting on Fabric because it utilizes unified tools for every data professional into a single platform, simplifying data management and analytics while reducing costs and effort.

Now, Fabric has shown the power of integrating multiple data analytic workflows together, and today we're thrilled to announce a major leap forward with Fabric databases. This is really a new class of cloud databases. And with the addition of Fabric databases, Fabric now brings together both transactional and analytic workloads, creating a truly unified data platform to enable scenarios that were previously very difficult.

The first of these Fabric databases, now in public preview, is our SQL database, which is built on top of the familiar SQL Server engine. SQL database and Fabric can be provisioned in seconds, is secure and governed by default, and is integrated with Azure AI. With SQL database and Fabric, you're now going to be able to build insight-driven intelligent apps faster with the data that's automatically replicated to Fabric's OneLake, making it instantly available to Fabric's analytic engines. Build AI solutions using Fabric databases, native vector search capabilities in Azure AI Foundry integration, and boost

productivity with auto optimizing and auto-scaling databases that provision in seconds. With that, let's watch a demo of SQL database and Fabric in action.

VIDEO DEMO: I'm a developer at Contoso Wireless, a multinational cellular company. I've been tasked to add chat to the application, so customers can get product recommendations and support. I've got the UI piece built, but it needs the latest AI models and access to our data. We have everything we need for our whole data platform in Fabric, and this now includes SQL database, which is built on top of the SQL Server and Azure SQL Database engine. One of the great things about Fabric is it makes getting started super easy.

Let me show you how quickly I can set up a database. I just write a name and click create. Provisioning takes seconds. I can use Copilot to help jumpstart my work. It gives assistance designing my database and writing queries both inline and in the chat. Here I'm asking Copilot how to automatically generate values for a column so I can develop my app. It explains several solutions like identity, column, default constraints and computed columns, as well as the code needed to make it work. Now that you've seen how easy it is to get started, let's switch to the Fabric database we're using in production for our app. We're storing a lot of information related to the products that we sell, along with descriptions, inventory, order details and more.

Remember, since this is built on top of SQL Server and Azure SQL database, it's got the T-SQL support you know and love. So, I can write queries just as I'm already familiar with. In addition to T-SQL compatibility, managing these databases is easy. SQL databases in Fabric are zoned redundant by default to keep them highly available. They're also secure by default and autonomous. Not only is my database auto-scaling to give me the best price performance, both from a compute and storage perspective, but it's also optimizing my indexes and applying the best practices by default. Now I need to implement the RAG pattern using our data to generate product recommendations for the chat app.

With new native vector type support in SQL database, I can easily make use of the latest AI models or any model in the Azure AI Foundry model catalog. I can generate embeddings, which are numerical representations of the data for all the product data, and store the vectors in my database. I can then use built-in functions like vector distance to search for results. SQL database works with popular frameworks like LangChain and Semantic Kernel, which made it super easy for me to build a natural language chat experience with Vector Search.

You can see here I can search for similar phones to the Alpha Phone X1 with a better camera. All the chat history is also stored in a SQL database, and all the data is fully integrated into one link in Fabric as delta parquet. This means my data is available to seamlessly use in any other Fabric workload and anywhere across my organization. For example, I can connect Power BI directly to the data and even use Copilot to quickly generate a report. Let's ask for a breakdown of chat interactions by sentiment and number of terms.

And just like that, Copilot has generated the report and I can see that the more terms users are doing in the chat, the more often they're ending in a purchase. And since this is all integrated in Fabric, data updates in the SQL database are instantly available in my report. Switching back to look at my solution in Fabric, you can see how Fabric has become the data layer for my AI app. From my database, and the report you just saw; to services related to event streaming and data pipeline, to the APIs my app uses and more. With Fabric, you can build AI apps faster and easier than ever. Try it for free today.

SCOTT GUTHRIE: Now that we've covered some of the new capabilities of our Azure Data Platform, let's now dive deep into the infrastructure that's powering all of this, which is Microsoft Azure. And we've invested heavily over many years to make Azure the place to do cutting edge AI innovation. We now have more than 60 Azure regions around the world, and we're bringing all of that scale and power to you. Azure is purpose built to support all workloads, and that includes those spanning on-prem, edge and multi-cloud environments. We know that enterprises value the modularity, integration and simplicity that Azure provides, but they also want the flexibility to extend that to wherever they need.

We're addressing these needs through an approach that we call the adaptive cloud. The adaptive cloud enables you to have consistent operations, applications and data models on the infrastructure of your choice. You can bring Azure anywhere while gaining control over your IT estate and investments, with a single management pane of glass. Underpinning this approach is Azure Arc, which supports running Azure services and workloads wherever they're needed, regardless of location. And this enables organizations to unify siloed teams, distributed sites and sprawling systems.

Over 39,000 organizations are using Azure Arc today to extend the Azure platform across hybrid environments to manage the resources consistently. Building on this strong customer adoption, we're continually expanding the capabilities enabled through Azure Arc. Generally available today, Windows Server Software Assurance customers can now access over 20 Azure management services at no additional cost. You can use favorites like the Azure Update Manager, Azure policy, change tracking and more for your hybrid machines, all for free.

As you heard in Rajesh's keynote earlier, Windows Hotpatch is coming to Windows 11 and Windows 365. Windows Hotpatch is also the No. 1 most requested ask from the Windows Server community, and we're delighted to share that HotPatching is now available for Windows Server 2025. This is going to enable you to patch security updates on your servers without reboots. For your Windows Server machines running on premises, or even in other clouds, Hotpatching is available through Azure Arc and also for VMs inside Azure. Now, let's see some of the new Azure Arc management capabilities and Hotpatching in action.

VIDEO DEMO: Azure Arc lets you manage Windows Server, SQL Server and Linux on physical servers and virtual machines hosted outside of Azure on your corporate network

or other cloud provider. In addition, Azure Arc simplifies security and governance by delivering consistent hybrid and multi-cloud management. Here we have Arc-enabled machines running around the world in Asia, Australia, Europe and the U.S. on physical, virtual and multi-cloud, all outside of Azure.

Since these resources are Arc-enabled, I can use Defender for cloud to review and improve my security posture. With Secure Score, I get actionable intelligence to protect my assets, even those running in multi-cloud. Next, we have Azure Update Manager, a unified service to help monitor and manage Windows and Linux update compliance across hybrid and multi-cloud from a single pane of management.

Update Manager is also how we're delivering the No. 1 customer requested feature in Windows Server 2025, Arc-enabled Hotpatching. Hotpatches are surgical security-only fixes that work by patching in-memory code of running processes without requiring a reboot. The result is that the moment the patch is applied, the workload is secure. In this example, we have a front-end application communicating with a back-end AI chatbot server, and we're hot patching these workloads without disruption. With Hotpatching, you go from rebooting your servers every month, 12 times a year, to rebooting your servers once every quarter or four times a year. A huge reduction with Windows Server 2025.

SCOTT GUTHRIE: Earlier this month, we released Windows Server 2025. Windows Server 2025 features advanced multi-layer security with delegated managed service accounts, the easy ability to integrate with Azure Arc, and lastly built in support for GPU partitioning, designed to handle AI and machine learning applications. Windows Server 2025 brings the best of Microsoft security, hybrid management and AI to you. Let's watch a demo of Windows Server 2025 in action.

VIDEO DEMO: Windows Server 2025 is the most secure, highest quality release with the easiest upgrades ever. It offers a modern look and feel with the tools to make you productive, like Windows Terminal and Winget package management. With Winget I can search for critical apps like SQL Server or tools for diagnostics and management like Sysinternals and Windows Admin Center, and install directly or script with automation. Speaking of Windows Admin Center, we're announcing the release of version two that can be used in Azure and on-premises. One popular use for Windows Admin Center is managing Hyper-V virtual machines. Here's a Hyper-V host running a variety of VMs, including Windows 11 VMs with GPU partitioning.

GPU partitioning allows you to divide a physical GPU among virtual machines and ensures predictable performance and secure isolation. Here I have a VM running at the bottom with a partition of the GPU assigned to it, and it's achieving about 40 frames per second. The VM is currently running on the server in the upper left corner, and I'm in the process of live migrating it to the server in the upper right corner. You can see the CPU utilization drop on the left as the VM seamlessly live migrates from server to server without disrupting the video. Windows Server 2025 delivering GPU partitioning and seamless live migration without missing a beat.

SCOTT GUTHRIE: As you've just seen, Azure and the Copilot stack is empowering every developer and every IT professional on the planet to innovate using AI. I'm happy to now hand off to Charlie and Vasu to talk more about how we're putting security into the foundation of everything that we do at Microsoft, and helping you secure your entire environment. Thank you and have a great Ignite.

CHARLIE BELL: Well, thank you, Scott, and hello, everyone. It's really great to be here.

VASU JAKKAL: I'm thrilled to be here as well. Love this energy.

CHARLIE BELL: Yeah. We're at the start of what could become one of the most transformative technological eras in modern history. AI has the capability to elevate human potential in all facets of our lives. And AI must be trustworthy in order to achieve this promise.

VASU JAKKAL: That's so right, Charlie.

And in security, we believe one of the most critical use cases of AI is protecting our world from cyber threats because the economics and sophistication of attacks are changing rapidly.

Every single day, our customers are facing more than 340 million attacks across cyber, criminal and nation state. Software attacks are occurring at a rate of 7,000 attacks per second, and the number of threat actors that Microsoft tracks has increased by more than five times in the last year alone. What's more, global losses from cybercrimes are expected to reach \$9.2 trillion this year. It's no wonder that cybersecurity is a top priority for businesses of every size across every industry.

CHARLIE BELL: Yes. And for everybody in this room, the good news is that while there will always be attacks, there are also defenders fighting tirelessly to stop them across the world and in all industries. They're performing heroics that make us all more secure. And they could use some backup, which is where AI will be a game changer.

We are a firm believer in the incredible audiences being created through AI innovation, but there is no doubt that security must be its foundation. AI transformation requires security transformation.

VASU JAKKAL: Absolutely right. And that is why security is job No. 1 for us in the age of AI. We launched the Secure Future Initiative, or SFI, in November 2023. And this year, we increased our investment in it with the equivalent of 34,000 engineers working on it and making it a core priority for every single employee at Microsoft. It reflects our commitment to resiliency and transparency and our role in the global security ecosystem.

In September, we published our first report, which gives insight into every single area.

And today, at Ignite, we are sharing our latest insights, our implementation strategies and guidance, including key learnings from security incidents to help us and our customers stay secure. The full report is available on our website, and we would love for you to read it.

CHARLIE BELL: Anyone in the audience who has been through a cyber-attack, and that will be many of you, knows what it means when I say it's not something you look forward to. What we at Microsoft experienced with Midnight Blizzard was about as uncomfortable as it comes. That said, the attack also provided a watershed of understanding and how an A player bad actor will try to exploit an environment. And what's needed to defend today's world of pre-planned, pre-staged and high-speed lateral movement.

We observed Midnight Blizzard working from many time zones effectively around the clock. They were moving quickly. In addition to the information we were getting from all the security tools we have, we found it incredibly useful to layer on our red team's attack graph, a database we kept to support our own simulated attack operations. By combining our understanding of this attacker with the graph, we're able to execute new, proactive defenses in real time as we rapidly closed off the edges of the graph that could provide potential future paths for the attackers.

We also employed generative AI tools from our research labs. One example was to find potentially exposed credentials that could lead to further attack progress. Every day, we improve our security offerings to customers with what we learn as a cloud provider. We see a lot of both normal and nefarious activity, and AI training feeds off of good and bad examples. The grounding we can provide through the graph, however, takes this to a new level.

Also hugely important is how the graph breaks down the silos between disciplines, organizations and technology piles. We talk a lot about the importance of end-to-end security. This is how we make it real. We also know that innovation comes from many places, and that islands don't work in security. That's why we work with more than 300 partners in our Intelligent Security Association.

As we work on the graph, we have a few early ecosystem partners, including Cisco, Rubrik, Illumio, OneTrust and Wiz. We want to make sure that we, our customers and other security companies, can all work as a unit. Open standards and interoperability matter because, as my friend Vasu always says, security is a team sport.

In just a bit, we'll talk about some of the early fruits of all this work.

VASU JAKKAL: Yes. Security is a team sport indeed, and SFI is all about evolving how we design, build, test and operate our products to achieve the highest possible security standards. SFI and the graph that Charlie just talked about are turbo charging a flywheel of security innovation, helping us take these learnings to our customers to provide them with a comprehensive platform so they can protect their organizations

across their heterogeneous environments, integrating Microsoft as well as third-party tools. A Microsoft security portfolio reflects our SFI journey.

Now, four years ago, we set out to unify and simplify security, making it easier for our customers to protect their digital estates. We brought together threat protection and cloud security and security ops with Defender and Sentinel, data security and governance with Purview, identity and access with Entra, and device management with Intune. And as we implement SFI, we continue to enhance our security products and use them to protect Microsoft.

For example, we protect nearly 1.5 million endpoints with Defender. They ingest 225TB of data into Sentinel every week for security operations, and we protect data on more than 72,000 SharePoint sites with Microsoft Purview. And with this advent of GenAI, we introduced Security Copilot last year, the first GenAI security product in our industry, giving security teams a powerful tool to defend at machine speed and scale.

And as Charlie said, we've seen firsthand how important it is to proactively manage your security posture. So I'm so proud to say today that you can experience some of the innovation and the prevention enhancements powered by the graph in Microsoft Security Exposure Management, which is now generally available. Exposure management is critical for enabling teams to understand the posture of their organization, and it helps security teams see all the potential attack paths to critical assets as if they were looking through it through the eyes of the attacker.

To show you what this looks like, let's go to Sherrod, who is on the front lines every single day watching and defending against threat actors. Sherrod?

SHERROD DEGRIPPO: Thanks, Vasu.

So what we've done here is we've learned that in the security world, it's all about increasing your security posture to prevent successful attacks and get proactive. So the reason that this matters so much is that advanced threat actors tend to gain access silently, maybe from a stolen credential, and then they just quietly sit in your environment waiting for new opportunities to arise.

What they're trying to discover is a path that will lead to some kind of valuable asset. One of the things I'm most excited about in exposure management is this view, where you can see the potential path that an attacker could take to get to your most valuable assets. This is proactive security. By finding these potential paths, you're identifying where you can prevent a future attack by changing a configuration or addressing a vulnerability.

The other thing we can do with these attack paths is find what they have in common. If you have several attack paths that go through a particular identity or a particular device, we call that identity or device an attack, a choke point. These choke points are some of the most leveraged opportunities to prevent potential attacks. If you address one of them,

you can cut off multiple attacks.

Because your environment is dynamic, new, potential attack paths and new choke points can pop up regularly. With exposure management, it can alert your security team when that happens so they can stay ahead of the attackers. We're enabling them to play offense instead of always playing defense, and thus, preventing an attack before it even happens.

Visu and Charlie, back to you.

CHARLIE BELL: Thank you, Sherrod. Prevention is definitely the best defense.

Now let's talk about two important strategies in the AI era. First, how we put GenAI in the hands of our security teams to help them be more efficient and more effective with Microsoft Security Copilot. And second, how we secure and govern AI itself, especially the data that is used by AI.

With Security Copilot embedded in Defender, Purview, Entra and Intune, members of the IT team and the security team can be more efficient and more effective. When we looked at how Copilot helped Entra and Intune admins in a controlled study, we saw they were 30% faster and 35% more accurate in tasks like troubleshooting and policy creation.

In another study, we saw security teams were 30% faster at resolving incidents. And in a Forrester survey, decision makers said they experienced or expected to see 17% fewer breaches. Let's see how Copilot has transformed security for Eastman Chemical.

SPEAKER: Eastman is a global specialties material company that produces a broad range of products people use in their everyday lives. Our scientists and researchers are on the leading edge to deliver innovative products and solutions.

The most important thing for us is our intellectual property because that's how we serve our customers.

We have sites around the world, and with that global presence comes a series of global threats.

It's important for us to have integration with Microsoft security solutions to pull data faster in a time-sensitive situation. Microsoft Copilot for security gives us the ability to be more effective when an incident occurs.

The next generation abilities benefit both junior analysts as well as experienced analysts.

Now they can concentrate, which alerts the most important one. That's a game changer for us, and we can take the remediation action very much faster than ever before.

We see Microsoft Copilot for security as a comprehensive solution.

We haven't found anyone else that can do that for us at the scale we need to do it, or the speed we need to do it or the accuracy we need to do it.

The speed and efficiency provides a huge advantage in not only responding to incidents, but ensuring if we're vulnerable in certain areas, to help better protect everything that we're here for.

VASU JAKKAL: All right. So now let's look at this. Thank you.

Let's look at the second side of the AI coin: securing AI. Since May of this year, we have delivered more than 30 features to help you secure and govern prebuilt and custom-built AI. We are the first end-to-end security solution to cover these five critical aspects of AI security. And at Ignite, we are announcing even more. We don't have time to go through them all, so we'll highlight just a few.

Charlie?

CHARLIE BELL: In every enterprise, the volume of data and the level of activity around it is hard to get a handle on. And this is only getting amplified with the adoption of generative AI. 84% of organizations agree they need to do more to protect against risky use of AI tools, making data security top of mind.

This is where Microsoft Purview Data Security Posture Management comes in. It gives you prioritized visibility into the data security risks impacting your organization. As you saw earlier, posture is critical to understanding potential attack paths, and this data posture management now extends comprehensively across the entire data estate, including generative AI.

Further, with built-in Security Copilot, you can easily uncover risks that might not be immediately apparent. Here, Copilot is helping the data security admin to discover unprotected data that is at risk.

VASU JAKKAL: That's great. And now let's look at how we protect data usage by AI applications.

New controls for Microsoft 365 Copilot help you prevent data oversharing. Here, Purview has detected that there's sensitive information in files that are not labeled and not protected, which creates an oversharing risk. Purview will walk the admin through the steps to get that data protected, and to create a policy that prevents unauthorized users from accessing that information via Copilot.

CHARLIE BELL: Right on, Vasu.

Purview can also detect risky usage of Copilot and other AI apps. For example, imagine a user has just resigned but hasn't left yet. They're using Copilot in ways that are significantly different from past activity patterns, like attempting to access sensitive data

from an over-permissioned SharePoint site. The same insider risk management and data protection controls that you just saw are also natively integrated into Copilot Studio, so you can protect your custom agents.

VASU JAKKAL: That's pretty amazing, Charlie.

And now let's talk about governing custom AI apps you build in that wonderful Azure AI Foundry that Scott and Seth showed earlier.

The security team can set organizational policies that govern which AI models your developers are allowed to include in their apps. So now you can see how we're innovating to secure AI and at the same time using AI for security. More so, we're building out a platform which can protect your organizations end-to-end.

I am so excited about the possibilities of what we all can do with these innovations, and I'm so optimistic about how we will build a safer world for all. I hope you will take this time to explore our 80 security sessions across Ignite, including our general session tomorrow morning, which is going to be packed with demos.

CHARLIE BELL: AI will elevate human potential. Those of you listening are the people who are going to make it work and make it trustworthy. You are going to bring the benefits of AI to the world. We are excited to be on the AI transformation journey with you. Thank you all for your incredible energy and engagement today.

Now let's keep the momentum going as we welcome our host in the Hub to share more about the amazing content we've lined up for the week ahead.

VASU JAKKAL: Thank you.

END