

11182025 Microsoft Ignite 25 Keynote

Microsoft Ignite Keynote

Judson Althoff

Tuesday, Nov. 18, 2025

VOG: Please welcome chief executive officer, Microsoft commercial business, Judson Althoff.

JUDSON ALTHOFF: Well, good morning. Good morning, everyone, and welcome to Microsoft Ignite. It's fantastic to be here with more than 20,000 of you live here in San Francisco, and more than 200,000 of you watching online around the world. So good morning, good afternoon, good evening, no matter where you may be.

We are all very excited here at Microsoft to talk to you about the more than 70 announcements we have for you this week. For today's keynote, we're going to focus on the major ones that empower the Frontier Firm. I want to thank you all, from the bottom of my heart, very sincerely, for spending time with us today. Your time is valuable. Once you spend your time, you never get it back. So we're grateful that you've chosen to spend it with Microsoft today.

Let's get started. I want to begin with this notion that Frontier transformation is distinctively different from AI transformation. But like any good journey, it's important to reflect upon where you've been so that you can plot the course forward appropriately. So, like for any event, I used Copilot to prepare for this one. I asked Copilot, hey, what is the state of the union of AI. It gives an honest response. Hey look, there's been no wave of technology ever in the history of the world that has ever been adopted faster than AI. But on balance, the success rates of AI projects are not what we'd like them to be.

Depending on the report that you read, there's a range of failures out there, and you have to look into the data to understand why. There's four very common reasons why AI projects failed. First, the alignment between business and IT professionals is very inconsistent. Second, there's a sea of data quality issues out there that need to be rationalized by AI developers in order to get meaningful output from the efforts. Third, governance and requirements in the markets, whether they be regulatory or otherwise, require AI to be kept on the sidelines from real, major efforts that can impact business. And finally, there's an overemphasis on experimentation. Random acts of innovation, rather than AI being put to use at scale across real business scenarios that can make an impact.

But fear not, the ROI in AI is not somehow elusive. At Microsoft, we have the pleasure of working with thousands of customers around the world, across industries, and they are making a difference with AI in their business. Many of them use the success framework in front of you, what I call the Frontier Success framework. And there are some common patterns that emerge.

These companies are deeply focused on enriching the employee experience, hiring the best talent, nurturing that talent, equipping them with the world's best tools, AI-infused into the way

they work, tying it back to measurable KPIs on the business so that they can see real impact in the results that they strive to achieve.

Second, Frontier Firms are ruthlessly focused on making sure that they see the results in the customer experience journey. It's about real time engagement with customers. It's about making sure that they have the connections, better relationships with customers so that at the end of the day, they achieve better results with a better cost structure.

Third, you have to reshape business processes. You can't simply apply technology to an existing process and expect revolutionary results. You have to comprehensively look left to right to see how a business process can become AI-first and truly be changed to impact the business.

Last but not least, the Frontier Firm focuses deeply on innovation and bending the curve so they can drive real competitive advantage in the industries that they serve. Whether it's advances in material science or drug discovery, or simply writing code faster than you've ever written it before. There's no question that putting AI to work for your competitive advantage is the right approach.

Look, the more you study it, the difference between successful projects and those that fail are this mindset shift from technology focused efforts to business led transformation empowered by AI. And in that order, look, the Frontier Firm needs to holistically reimagine the business, putting AI deeply aligned with human ambition to achieve the organization's most aspirational goals and growth potential.

At the heart of every Frontier business flow is this notion of democratizing intelligence. Again, human ambition at the core, coupled by your AI assistants and agents to get real work done. At Microsoft, this is Copilot as an agent ecosystem aligned to your business, infused into the way you work. Simply stated, AI transformation needs to do more for humanity.

We need to evolve the approach to become Frontier. We have to obsolesce the mundane, unlock creativity and innovation. The world's greatest filmmaker. The world's greatest musician. Educator. The cure for the incurable. Advances in climate science and basic research. All of these things that can move humanity forward can be unlocked by democratized access to intelligence.

So there's no question that aligning to the business aspirations and putting AI to work is the right approach. But beyond that, if you study successful Frontier Firms, there are three common traits to their approach.

First, they put AI in the flow of human ambition. And this is contrary to somehow cutting and pasting your life into a chat window and expecting insights. It's about AI infused in the tools that people love to use each and every day, seamlessly accessible to get real work done.

Second, it's about ubiquitous innovation. There's a maker in every one of us, and the Frontier Firm has a maker in every room of the house. By putting tools into the hands of people that can create AI artifacts right in the flow of work. The person often closest to the problem and the

challenge and the opportunity is the best person to solve it. And that's what we do here at Microsoft.

Third, the Frontier Firm needs to govern and manage the AI they deploy, and it requires observability at every layer of the stack. What agents do you have? How many people are accessing them every day? What are the workloads they're assigned to? What results are you getting? Governance, manage, secure and optimize the AI that you put to work in your business.

Beyond these three traits, there are two fundamental building blocks: Intelligence and trust. Now you might say, hey, Judson intelligence is kind of obvious. But I will challenge you that we need to put the I back into AI, because if you look at what's happening today, developers have to navigate a sea of data, thousands of connectors, thousands of APIs, all while sipping data through a thousand tiny straws and trying to inference over it, which is just a fancy way of saying guessing. Folks, we can do better. And today, we're going to show you how.

The fundamental building blocks of AI need to know how you work, with whom you work. The content over which you reason. They need to provide real intelligence to developers, to help them scale and build better solutions than ever before. Finally, trust. We're going to talk a lot about observability today, but it's about aligning AI to your purpose and your principles to get the outcomes that you achieve. Want to achieve. At the end of the day, it's about your IQ and your differentiation in your business and in your industry.

Today, everything we'll show you will be model diverse, open and heterogeneous at every layer of the stack. Because again, it is about your ambitions, what you aim to achieve, and Microsoft here is to empower all of those things. The rest of our keynote this morning is organized around these three successful traits of the Frontier Firm AI, and the flow of human ambition, ubiquitous innovation and observability at every layer of the stack.

At Microsoft, we're at our best when we put our portfolio of technology and capabilities to work for your business needs. So accordingly, everything you'll see today is an integrated flow of our capabilities aligned to real business outcomes. Let's get started with AI in the flow. And to do that, please join me in welcoming to the stage Ryan Roslansky.

RYAN ROSLANSKY: All right. Thanks, Judson. It is an honor to be here today with all of you shaping the future of work.

You might be wondering, why is the CEO of LinkedIn, a company owned by Microsoft, here working on and talking about M365 Copilot? Well, the reason is simple: the future of work and the future of AI are inseparable. LinkedIn is the largest AI driven professional product and community in the world. We know what it takes to build AI that works for 1.3 billion people.

And now we're bringing that expertise together with Microsoft's enterprise scale. We're bringing Copilot to put AI in the flow of human ambition. Over the next 20 minutes, you'll see the future of work unfold, how we take the most powerful AI innovations and make them yours. Built into the rhythm of every job, every role. Every day. From developers, security analysts, knowledge workers, Copilot is becoming essential to how work gets done.

Today, more than 90% of Fortune 500 companies turn to Copilot as they journey to become Frontier Firm. Why? Because Copilot is built for work, and there are three things that make Copilot special. First is something that we call Work IQ. Second, Copilot is embedded in the Office millions of people use every day at work. Third, it's your window into the world of agents.

Let's start with Work IQ. This is the intelligence layer that enables Copilot and agents to know you, your job and your company. Work IQ isn't just Copilot brain, it understands your business, workflows and relationships. Work IQ connects everything that matters inside your company and turns it into an experience that feels effortless, powerful, personal. Made up of three things: data, memory and inference. Every organization is built on data. Think about it. Your emails, documents, databases, line of business systems. Data is what codifies how work gets done. And it serves as your biggest differentiator and your most valuable asset. Think about it.

So much of that data is untapped. It's tied up in disparate systems and files. But with Work IQ, Copilot enables you to unlock that data, putting it to work to deliver real value for you and your business. Memory is what gives Copilot the ability to deliver personalized responses based on your style, your habits and your preferences. Tracks and understands workflows. Recognizes how people, task and tools all fit together. And then, last, inference uses all that rich signal in the data combined with memory to make valuable connections, to unlock insights and predict the next best action, or even the next best agent for the job.

Most importantly, Work IQ is the AI that you can trust your most valuable information with. For five decades, organizations like yours have trusted Microsoft to securely handle that data. Now, you might be wondering how is Work IQ different than just using connectors? Work IQ works across all of your data in real time and has full context, which allows you to find patterns and insights, while connectors only pull fragments of the data with no real understanding of the context, which means they can miss critical information, they can ignore relationship and give incomplete answers.

But Work IQ is built into your trusted M365 environment, so it respects all your permissions from the start. With Work IQ, Copilot understands the whole picture securely and instantly. Delivering faster, smarter and more accurate results in the flow of work in ways that connectors just can't replicate. Now look, many of you are asked by your CEOs, your board of directors, what's your AI strategy? Here's the good news, whether you know it or not, inside each of your companies today, you have what you need; the data, the memory, the inference, and it's all wrapped up in Work IQ.

Now, we all know that AI works best when it's woven into everyday workflows. So then you ask yourself, where is work happening? Turns out it's happening in these apps. Literally. Billions of Word docs are created monthly. Hundreds of millions of PowerPoint presentations in team meetings are created weekly. Apps like Word, Excel, PowerPoint, Outlook and Teams. That's where work happens and they have changed the way that people get their jobs done. We've seamlessly integrated Copilot into the Office apps, enabling productivity directly in the flow of work.

And while we all know that Copilot can do a lot inside of these apps, we also know that there are just so many specific use cases, processes, tools that are unique to each company, which is why we built Copilot to be your window into the world of agents. To handle custom and specific use cases. Agents can take on tasks, handle routine workflows and automate really complex processes.

Now, many of you already know this. You already embraced it. But every organization will run on a distributed ecosystem of agents moving forward, and Copilot is your front end for that. You can now tap into that Work IQ intelligence layer and build custom agents tuned for unique workflows and business needs. Because Copilot gives you the ability to build, direct and coordinate that network of agents, whether they come from you, they come from us, they come from any other partner. Together, Copilot agents and Work IQ amplify what your people can achieve.

So to recap, you've got your intelligence layer with Work IQ. We've got Copilot embedded in the flow of work. And we have an ecosystem of agents where Copilot is the orchestrator, putting agents to work for you so they can execute any task or business process. OK, that's enough talking. Let me show you a little bit more about how we're bringing Copilot and agents to life for every employee in your organization.

I start every morning by pulling out my phone, and asking Copilot: Copilot, what's on my agenda for today?

COPILOT CHAT: You've got the Ignite 2025 opening keynote organized by Judson Althoff, starting at 9 a.m. and going until 11 a.m. today.

RYAN ROSLANSKY: Thanks. Has Judson sent me any emails today about Ignite?

COPILOT CHAT: Yes. You have an email from Judson this morning about Ignite. He's suggesting giving everyone a T-shirt during the keynote to boost excitement and make the event memorable.

(Cheers)

RYAN ROSLANSKY: Yeah. It's easy for you to clap, but I got to figure this out now. So Judson wants 20,000 Copilot T-shirts delivered to this audience. Now, OK, yeah. Lucky for you. We don't have a lot of time, but I'm going to hand this over to my friends at Zava. This is a fictitious company we created just to show you what a Frontier Firm can look like. We don't have a lot of time. I'm going to give them six minutes to fulfill this order. So, Zoe, over to you. Hope you can pull off a little magic.

ZOE HAWTOF: Thanks, Ryan. I'm the CEO of Zava, a fictional retail company that specializes in intelligent athletic apparel. Like a lot of you, my day is busy juggling lots of competing demands. I'm gearing up for the holiday demand, I'm dealing with security concerns, and managing urgent customer requests. I use Copilot to stay on top of it all. I'm going to ask it

to review all my emails. And with GPT-5, but powered by Work IQ, Copilot knows who I am, what I work on, and who I work with. It's actually why I use Copilot. It has the context it needs to give me real-time and actionable results.

Like I can see here, there's a request for 20,000 branded T-shirts from Microsoft, and I heard the crowd's pretty excited about this, and I only have a few minutes to get it done. I can already see the timer going, so I got to be quick. I'm going to get a quick overview of Microsoft as our customer, that way I have full understanding.

And Copilot is going to again tap into that Work IQ to pull my meetings, my emails, my chats, but it's also starting to infer what I'm trying to accomplish. It understands that I'm trying to dig deeper into my customer. And so, while it gives me these great results, it also recommends an agent that's specially built by my company for customer insights.

Now, in this case, the Copilot response looks great, and I need to quickly mobilize my team because work doesn't happen in a silo. I'm going to start a group chat with my analyst Anne, and my sales lead Nirja, and Copilot is going to bring this conversation over into Teams in the same place that I already communicate and collaborate with my coworkers today, and they have all of the context that Copilot just showed. When I reach out and ping them, they have the same context I do.

Oh, they're super excited. I can already see they're replying. And now we can continue to use this to work together and with Copilot as we get this amazing opportunity. But I want to quickly check in with my analyst Anne. Anne, how's inventory looking? Are we going to make it?

ANNE KRUPKE: Well, this is a pretty big rush order, so we're going to have to get creative. We'll need to break it up into smaller pieces and allocate it across multiple production facilities if we have a chance to get this done in time. As soon as Zoe pinged me, I turned on Agent Mode in Excel and I set it to work, asking it to whisk us up a plan of action, and this is what we're looking at in real time. I can follow the step-by-step process of the agent as it works through the problem, and we can see how it's building out inside the spreadsheet right in front of us, a solution.

And it looks like it's just now finishing up. Because the agent is built into Excel, it speaks Excel natively, which means it can do everything from simple to complex. In the plan here we can see the scope of our problem, the total quantity, the due date and the question, and I get the detailed breakdown of the proposed plan in table format. And this is not just static fields; this is dynamic formulas, so as I make tweaks to the plan, everything updates in real time.

And not just tables, I also got a chart here from the agent, which shows me the impact of the proposed plan to overall production capacity. And finally, I get this beautiful conditional formatting gives me a really clear green flag that yes, we can meet the order by the due date. And if Zoe wants to know how I got that answer, I have the entire history of the agent's calculations and reasoning that I can share with her. So, Zoe, looks like inventory is good to go.

ZOE HAWTOF: Excellent, Anne. That's great news. We have about two minutes. Nirja, how's the sales quote coming along?

NIRJA TREVEDI: I'm on it, Zoe. All right, so here I have an email with details about the opportunity from the customer. Now, I need to generate a sales quote, and I need to do it fast, so I'm going to jump over to our sales quote template where I already have Copilot opened up. Here, I'm going to activate Agent Mode and then give it a bit more direction to get started.

Now, this type of work is very time-consuming, but the agent helps lift things off my plate. What I want the agent to do is I want it to pull in information from the email, specific information like the address and the number of T-shirts. But because of Work IQ, it's going to look across my entire work context, pulling in information like the shipping costs and the pricing that didn't exist in my email.

And I could see here that it's getting to work. I can see it co-writing with me as it fills out the form, but I am still in the driver's seat. I have the ability to edit as needed or have Copilot take another path. But this looks great, so I'm going to keep it all, and done. All right, Zoe, here's the final quote for you to send to the customer.

ZOE HAWTOF: Thanks, Nirja. OK, let me let Ryan know we have less than a minute. I'm going to have Copilot help me quickly draft a response saying that we can fulfill this order. I'm going to keep that, and I'm going to attach the sales quote that Nirja added, and then I'm going to press send. Now we can stop the clock. We made it just in time with 27 seconds to spare, so everyone gets a T-shirt. And that's the magic.

(Applause.)

It's the magic that comes from Copilot powered by Work IQ. It goes beyond connectors to truly understand the context of how work gets done at Zava. But of course, Microsoft isn't our only customer. We're getting ready for the holiday rush, we're pulling out our aprons, we're rolling up our sleeves, and we need to be sure our website is ready to deal with the surge. Let me hand it over to Dean, our cloud architect. Dean, how's it going?

DEAN KROKER: Hey, Zoe. Everything's going great over here. Here's the scoop. I just recently scaled our Kubernetes environment, and I added a few new SKUs to our website, so now I'm going to test our checkout flow with this new water bottle. We'll go ahead and add her to cart, head over and proceed to checkout. It's saying that checkout is currently unavailable, which is odd because I know we have inventory on hand. So let me head over to the Azure portal to investigate our Kubernetes environment.

Azure Copilot agents makes this sort of thing super simple. I'm going to prompt Copilot to investigate the health of my Kubernetes cluster, we'll go ahead and tag the cluster and let Copilot get to work. Copilot understands the full context of my AKS cluster. It checks pod health, looks for anomalies and pinpoints which pods are unhealthy. Here we can see that access to the Cosmos database is blocked by the account firewall, so there might be an API issue, but let's use Copilot to help with this.

We'll go ahead and ask Copilot to help me troubleshoot my Cosmos Database and go ahead and tag the Cosmos DB and let Copilot get back to work. So this time Copilot is checking firewall signals and service health checks and firewall diagnostics here, and it's saying that network access is disabled within the firewall. Actually, here's my favorite part. Rather than integrate the code, which we can do using GitHub's integration, we can approve the change and let Copilot do it for me.

OK. That was really, really quick. But go ahead and head back to the website. We'll double check the work by trying our checkout again. We'll proceed to checkout. OK, great. Looks like we've got a fully operational checkout page. And, wow, that used to take my infrastructure team hours. We just did that in I think just a couple of minutes. Hey Zoe, the website is good to go for the holiday rush.

ZOE HAWTOF: Awesome. Thanks, Dean. For a retailer like us, the volume of holiday business can also expose us to heightened security threats. Our security team is working around the clock, so let's check in with Scott, our security analyst, to see how things are going. Scott?

SCOTT WOODGATE: Thanks, Zoe. Now that many orders are coming in, my job as the security analyst is to keep Zoe and team safe, and that's why I use Security Copilot directly inside Defender to stay ahead of constantly evolving threats like phishing attempts. Now, every day, users submit hundreds of suspicious emails to me, and most of these turn out to be harmless spam. The phishing triage agents helps me cut through the noise at scale and identify malicious emails. Let me show you.

It sits right here in my flow of work, and it autonomously filters out the false alarms so I can focus on what truly matters and what really requires my attention, and that is, of course, the true or real active alerts such as these. Let's drill into this one. You can see that the agent explains what it's thinking in simple natural language, which is awesome. It reasons over content, headers, URLs and attachments, and then visualizes its entire process in this intuitive decision graph.

And if I open one of these cards, I can see what it's thinking at any step and the evidence it used. As you can see here, this looks like it is a real order from Zava, but unlike your T-shirt order, this one is actually a phishing attempt targeting Zoe. And if I continue to follow the flow from the agent to the end, it does in fact confirm that this is a true phishing email. With Security Copilot, I'm able to keep Zoe and Zava safe faster than ever before. Back to you, Zoe.

ZOE HAWTOF: Thanks, Scott. As a Frontier Firm, every one of our employees, including our most specialized teams, they rely on Copilot to get work done. We can't imagine working without it. Back to you, Ryan.

RYAN ROSLANSKY: Awesome. Well done, Zoe and Zava. A round of applause for them. So that was a whirlwind tour, and what you just saw was M365 Copilot, Azure Copilot agents and Security Copilot agents bringing AI into the flow of work for every knowledge worker and every domain expert. M365 Copilot is AI built for work. It delivers our richest, most powerful

experience, and Copilot Chat is secure AI chat for every employee in your organization. It has enterprise data protection, it gives you access to agents in the flow of work and it's available to every M365 subscriber at no additional cost.

Today, we're announcing two new updates to Copilot tied to some of the innovations that you just saw. First, we're excited to announce we're bringing Word, Excel and PowerPoint agents into Copilot Chat. Whether you start in chat or start in the Office apps, Copilot can actually do work for you, and you stay in control the whole time, like you saw, reviewing, editing or guiding Copilot to rework as needed without ever having to copy and paste from chat.

Also, Copilot Chat can now reason over and make sense of your entire inbox. It can prioritize and summarize your emails, help you prepare for upcoming meetings and even highlight action items from previous meetings. All of this is coming to every employee in your organization. Of course, security is the foundation of everything we do at Microsoft. I'm excited to announce we're bringing Security Copilot into M365 E5, putting AI-powered protection and response right in the flow of work.

(Applause.)

And if you only take one thing away from the past 20 minutes, let it be this: Work IQ is the intelligence layer that knows you, your job and your organization in a way that's simply not possible with connectors. This is the new era of work where every employee across the organization has an AI thought partner, and there's just no other solution on the market that comes close to the value that we're able to deliver with M365 Copilot.

JASON GREGORY: Excuse me. Excuse me. I hope I'm not interrupting anything important. I have a delivery for Ryan. 20,000 T-shirts.

RYAN ROSLANSKY: Excellent. That's right on time.

JASON GREGORY: Would you like me to leave them on the counters by the exits, or do you want me to hand them out?

RYAN ROSLANSKY: Why don't we go ahead and put these at every exit so that when the keynote's over, everyone can grab one on the way out? But everyone, please don't get up yet, we've got a lot more to show you. Because we're not the only ones bringing AI into the flow of work, it's happening across all industries. Teachers, farmers, factory workers, even doctors are transforming their day to day with AI. To talk more about that, please join me in welcoming Seth Hain from Epic.

(Applause.)

SETH HAIN: Thanks, Ryan. The real value of healthcare tech should be measured in the human possibilities it unlocks for both patients and their clinicians. In the context of medical care, embedding AI into the flow of work requires disciplined acceleration, combining the immediate assistance of AI and workflow with the increasing familiarity and confidence needed to quickly

adopt the next wave. The importance of getting this right has never been greater, given the workforce shortages and aging population needs facing our society.

At Epic, we create technology to help people get and stay well. Our clients call it their operating system, informing mission critical decisions from the emergency room to the ICU, and importantly, connecting with you, their patients through Mychart chart. In March of 2023, just as generative AI was gaining momentum, doctors started using it in our software to respond to questions from patients faster. Our first AI agent was released in the fall of that year to help study patterns in medical data.

Fast forward to today, and AI is embedded throughout our applications, from the clinic to the bedside to the back office. And worldwide, hundreds of thousands of clinicians are using it to improve care. Central to this rapid global rollout was our established AI foundation on Azure and the secure, reliable environment it provides for sensitive healthcare data. AI in the flow of work both saves time and improves quality.

AI-generated end-of-shift notes, discharge summaries, those things help nurses wrap up their day quickly, and with Epic integrated with Microsoft DAX Copilot, doctors' notes are prepared for them based on the conversation in the exam room. Health systems are finding a 30%–40% reduction in the time it takes to generate medication prior authorizations. Let me translate that for you. It means you get treated faster.

AI also helps make sure no handoffs in your care are missed. Last month alone, using our software, 6 million summaries of medical records were generated for clinicians. Doctors reported discovering information 28% of the time that they would have missed with manual review. AI also assists in identifying patients needing follow-up imaging. Using our AI capabilities, the Christ Hospital reached a 69% rate of early-stage cancer diagnoses, compared to the national average of 46%.

Approaching AI with disciplined acceleration helps physicians care for more people today, while building toward a future where AI unlocks new possibilities in drug discovery, precision medicine and potentially understanding disease itself. My colleague, Dr. Jackie Gearhart, is here to share what this looks like in practice. Over to you, Jackie.

JACKIE GERHART: Thanks, Seth. As a family physician, AI is transforming the way I practice. Today, I'd love to show you how. I'm ready to see my first patient, Oscar. So before I enter the exam room, I'll pull up his chart. On the right side of the screen, I can see that generative AI created a summary. This provides me with the most relevant details of his health history. Since I last saw him, he's had genetic testing done because of his family history of high cholesterol.

I can scroll down to that section of the summary and click to see his test results in the progress note. Now that I'm up to speed, let's head in. The notes panel on the right side of Oscar's chart is where I would normally type my observations during the visit, but instead I'll use Epic's AI charting to capture the conversation. I can place a mobile device between the patient and me that will securely record our interaction and turn that transcription into a clinical note.

It's sent right back to me here in the chart, and in fact, I can see it right here on the right side. I love that this feature enables me to focus on face-to-face care with my patients, but the AI didn't just produce a note, it did so much more. The AI picked up on my care plan for Oscar and already teed up his orders and follow ups, like a referral for physical therapy. This makes it so easy for me to review, edit as needed and sign.

Now, one of my patients' biggest frustrations is when their prescriptions are delayed often because insurance still needs to gather more information. Luckily, AI could help me there, too. It's noticed that one of the medications I prescribed Oscar might not be covered by insurance, but the prior authorization feature can search through his chart and fill out the answers to his questions from insurance company to make sure he gets approved faster.

But the story doesn't end when I close his chart. My care plan becomes actionable for Oscar right on his phone, in the MyChart app. AI can provide him with a set of to-dos post visit. In the app, Oscar can see suggested tasks based on my reminders from the visit, like this one to ice his shoulder. And what's really cool is I can check off when he's completed that task, so I can follow along with his progress.

AI lets me focus on what matters most, my patients, deepening the purpose and joy that brought me to medicine in the first place. Back to you, Judson.

JUDSON ALTHOFF: Thanks, Jackie. How about a round of applause for Jackie and Seth?

(Applause.)

Look, the work that Jackie and Seth do is deeply personal to me. My father was a surgeon for over 50 years, and even as a child, he would lecture me that there was a mantra about operating at the top of your license as a doctor. The work that EPIC is doing to infuse AI into the flow of how health care providers deliver care allows physicians to go one step beyond, not just performing at the top of their license, but performing at the top of their potential.

We're really privileged at Microsoft and honored to be able to provide the platform and tools that EPIC users to build their AI solutions. We're actually very excited to share more with all of you today about how to leverage the same.

Like I said in my opening, there's a maker in every one of us, empowering citizen developers and professional developers, and enabling them all to collaborate with IT, to create amazing AI solutions, and then allowing them to build upon new layers, new fundamental building blocks that embody context, and intelligence is what it's all about.

Here to share more about ubiquitous innovation, please welcome to the stage Asha Sharma.

(Applause.)

ASHA SHARMA: Every person in every role now holds the power to create right inside of the tools they already know and love. From the soil of a Nebraska farm to the code running a global bank, innovation is no longer the domain of few. It has become the superpower of many.

Work will be defined by human creativity and agents that help it scale in this era. And at the center of this shift across individuals, domain experts and developers is Microsoft. Let's start with the individual maker inside of all of us.

Every one of us has had an idea for an app that would make our lives just a little bit easier, but today, turning that idea into something real takes someone technical and often months or years of work. Not anymore.

Today, we're announcing App Builder so anyone can create an app in minutes, right inside of M365 Copilot. All you do is start with your data you already have, your charts, your files, your meetings, all powered by Work IQ. You design the experience you want in natural language, and you share it with your coworkers, just like you would a doc. This is the beginning of something new, software evolving from being made for people to being made by people.

Now, innovation doesn't just stop at the individual builder. Some ideas require more domain expertise, more data and deeper connection across your business. Copilot Studio is the only low-code agent builder that understands your business right out of the box. It uses Work IQ, workflow, security, identity, to take real action and not just generate basic prototypes. Today, more than 230,000 organizations use Copilot Studio, including 90% of the Fortune 500.

Now, for decades, developers have lived at the frontier, and they have been the ones to solve the world's hardest problems, and that doesn't change in this era. What changes is the leverage. We're bringing them new AI-powered tools to continue to build the future of what's possible, and they do this on GitHub.

This year, someone new joined GitHub every single second, which is the fastest growth we've ever seen. GitHub Copilot started as a pair programmer, and now it's become an entire fleet of coding agents. And now it is also the single largest contributor to GitHub's codebase.

But we know that developers want choice, enterprises want control and everyone wants fewer subscriptions. That's why, at GitHub Universe, we announced Agent HQ, a single place to pick agents from GitHub, OpenAI, Anthropic, Google Cognition, xAI and more, all with consistent governance and observability as part of Copilot subscriptions.

Now, when innovation becomes ubiquitous, everyone gets a new set of powerful tools to drive progress forward. And in a retailer like Zava, that means the frontline associates in the aisle, the regional managers are running performance across the stores, and developers are writing the system to code every single line. That's exactly what we're going to show you.

Please welcome Miti, Lydia and Kyle to the stage.

(Applause.)

MITI JOSHI: Today at Zava, the store is packed and sales is strong, but as the store manager, I am scrambling behind the scenes. Every shift change means tracking down who's free, having multiple forks to negotiate swaps and calling in favors just to keep things running. I really need some help, so I reached out to Lydia over in the Store Operations Team, who says there's no bandwidth for this right now. There has to be a better way. Lydia says I should check out Microsoft 365 Copilot, where I can build apps, agents and workflows.

Now, I want to build an app to solve my staff scheduling challenges. I'm going to go ahead and ask it to build me one. I'm going to ask it to create an app with a dashboard and rich visuals, so I can easily manage staff scheduling.

Now, App Builder understands what I'm asking for, and it immediately begins generating the app right here inside M365 Copilot. I'm really hoping it uses all the documents, spreadsheets, notes and emails I use today to manage schedule.

And here is the app it built. And I can see here a dashboard for shifts across time, a manager view just for me, a visual for store coverage, and even a table for our team members and their shifts. Most importantly, App Builder builds on top of Zava's Microsoft 365 data, and it understood the business the way I do. I did not need to learn any new tools. I just described what I needed. No more frantic calls, no more guessing who's free because I've got an app to do it for me.

Now, I'm so excited with what I built here, and I'm going to share it with Lydia to see where it can go.

LYDIA WILLIAMS: And the second I saw it, I knew this wasn't just going to help one store. All of my stores could use this. Plus, I'm going to address another challenge, when someone can't make a shift and chaos ensues. Enter Microsoft Copilot Studio to help me build an agent to fix that.

I say to Copilot Studio, when someone emails that they're sick, I want an agent to secure a replacement team member. Instantly, Copilot Studio jumps into action and creates the full agent framework. Copilot Studio has native access to Work IQ, including Microsoft SharePoint and the Microsoft Graph, so it's grounded in rich context.

First, I need to determine when I want the agent to be triggered. When a team member emails our scheduling inbox to say they're sick, I want the agent to jump into action. We'll go ahead and choose that, go ahead and give it an appropriate title. And just like that, my trigger is added.

Next, I want to go ahead, and I want the agent to interact with a staff member to confirm their availability. To do that, I'm going to go ahead and add a Copilot Studio Flow. This one's called Contact Replacement Team Member, so we click on that. And just like that, that's added to my agent.

Finally, I want the store manager in the physical store to be updated about which staff member will be covering the shift. I'm going to use an existing agent that my colleagues built in Microsoft Foundry, which drastically speeds up my efforts.

All we need to do to integrate Copilot Studio with a Microsoft Foundry agent is give it a name, add the description and the ID.

Let's play this out.

It's 7:45 a.m. on a Saturday, and William, a team member, says they're sick. Before anyone sees the email, the agent jumps into action. It starts by suggesting a replacement team member based on skills and availability, then contacts them to confirm they can cover the shift. Finally, it goes ahead and contacts the store manager and lets them know that in this case, Luca has confirmed he's going to be able to cover the shift. And now my Zava pop up stores have an app and an agent working together to ensure every shift is covered inside Microsoft 365 Copilot and Copilot Studio.

And now on to the next task on my to-do list. I want to handle this message that Ryan sent. It looks like the new ecommerce site isn't currently optimized to work for mobile devices. Let me go ahead, file that as an issue. We've got GitHub, and let Kyle from our Development Team figure that one out.

KYLE DAIGLE: Thanks, Lydia. I am a developer at Zava, and I work on our online retail store. We use GitHub for our entire software development lifecycle. While I'd normally be making my coffee right now, I'm going to start by burning down issues in our backlog. And I've asked Copilot to tell me what's on my plate.

GitHub Copilot gives me a ton of choice with models from all the Frontier labs, and as of earlier this morning, that includes Gemini 3 Pro. And we even have a custom, finetuned model deployed to Azure Foundry.

Here's the issue from Lydia, improving the website to make it more mobile friendly. OK, honestly, this is important, but it's kind of boring to do, so I'm going to ask Copilot to start work on this. And because this is a UI task, I'm going to use the custom coding agent that our team built focused specifically on frontend changes based on our frameworks and styles.

Let's also ask Copilot to include screenshots so I can easily see the difference without having to run the app. Copilot gets to work right away, planning and writing code, tests, and opening a pull request for my team to review and merge later. I can manage this task in all of my agent sessions in one place, Agent HQ.

OK, Copilot has that task in hand, so let's go be a team player and work on some security remediation for our project. Our security team uses Microsoft Defender for Cloud to monitor our security posture and find risks in production. Today, we are announcing a new code-to-cloud integration between Defender and GitHub Advanced Security so we can filter for vulnerabilities using runtime data, like sensitive data runtime issues.

Our security team has created a campaign to tackle these vulnerabilities. They choose which vulnerabilities to focus on and a deadline to get it done, so we have an achievable target. I'm going to select all of these vulnerabilities and assign them to Copilot. Copilot will generate these fixes all in a single pull request for our team to review.

Now, I also spend a bunch of time in VS Code, and of course, I use coding agents to help with my work here, too. Here, I can see the agent sessions related to the code I'm working on across local chat, CLI sessions and even third-party agents like Codex. I can even start new tasks from right here, and they'll run in the cloud.

Now, I realized I haven't written any tests for this code, so before my team calls me out in a pull request, let's add some tests so we can make sure that we're using off on all of our pages.

At Zava, we have developers using Copilot in a lot of different ways. And now all enterprises can see how Copilot is being used within their organization with Copilot Metrics. I can see usage statistics, what models my team is using or not using, and even the programming languages that Copilot is assisting with most.

OK, let's go back to where we started with Agent HQ. We kicked off a bunch of coding tasks, and I can see all of that work here.

Now, I want to take a look at the mobile task I created earlier to resolve Lydia's issue. I want to give it a little more context, and I can just do that while it's running and let it know that I only care about phone sizes, not tablets right now.

In only a few minutes, I got Copilot to remediate security issues, start making our online retail store more mobile friendly and improve our testing, so I stay off the naughty list this year, all while keeping me in control.

What we just saw was Miti using App Builder to solve a scheduling problem, Lydia using Copilot Studio to scale it with an agent workflow across their region. And I used GitHub Agent HQ to tackle three tasks in two minutes. That's the power of ubiquitous innovation for each of us.

Back to you, Asha.

(Applause.)

ASHA SHARMA: Thanks, team. Incredible work.

Now, with this new work and this new world of creation, people that build AI-driven products need a very different foundation, one built from models, tools and intelligence. That's why we built Microsoft Foundry, an open and modular app server for the AI era.

Foundry has powered nearly every single product that you've seen on the stage today. And as we mark our first year, the momentum is unmistakable. Over 80,000 organizations are building on Foundry, and today, we are the No. 1 AI platform in the world.

(Applause.)

Now, we find ourselves in a multi-model world. Foundry has become the home for the leading models like OpenAI, Mistral, DeepSeek, LLaMA that your teams use every day. But there's been one Frontier provider that we've been missing.

This morning, we changed that. Anthropic models are now live on Foundry.

(Applause.)

Please join me in welcoming Mike, the chief product officer of Anthropic.

(Applause.)

Hey, Mike.

MIKE KRIEGER: Hey, Asha.

ASHA SHARMA: Thanks for being here. What a year it's been together.

MIKE KRIEGER: It's been an amazing year together. I am so excited to be here and so excited for the news.

ASHA SHARMA: Amazing. Do you want to – let's tell everybody what we launched this morning.

MIKE KRIEGER: Absolutely. Our partnership together is really about enabling developers and enterprises to build intelligent agents with cloud, no matter where they are. And that's why, today, we announced the availability of Claude Sonnet 4.5, Haiku 4.5 and Opus 4.1 models, all our Frontier models in public preview on Microsoft Foundry. Cloud on Foundry enables customers to build, scale and secure production-ready applications and enterprise ready agents in one trusted platform.

ASHA SHARMA: Now, we've been finding ways to work together over the last year, but we've been cooking up some exciting things. Do you want to show everybody what we're going to be rolling out over the next few months?

MIKE KRIEGER: Absolutely. We've already partnered together to bring the cloud family of models to GitHub Copilot's 26 million users. We've been collaborating closely on the MCP standard for tool calling, and we've been offering Claude as part of M365 Copilot Chat for billions of users.

Now, we're taking the next big step by bringing cloud models to all of Microsoft Foundry. In this example, we're using Sonnet 4.5 and skills to build a product plan for our favorite line of socks for Zava. Claude builds the plan by using Work IQ to find context about market details and potential pricing. And then it uses the Zava branding skill to create potential tag lines that feel on brand.

Next, we use Claude to create a PDF by using the PDF skill, and this quickly creates a beautiful, fully designed executive summary to share with the team. And then finally, and this is cool, we'll use M365 to send the PDF back to the two of us. This is Microsoft's enterprise knowledge combined with Claude's agentic capabilities and skills.

ASHA SHARMA: The thing I love the most about this is Claude's skills working together with Work IQ. I feel like we're the only two companies in the world that can bring that type of use case together for enterprises.

MIKE KRIEGER: Absolutely. And this goes beyond knowledge work to coding as well. thanks to our partnership with GitHub, Claude has become a first class agent, which means your developers can collaborate with Claude as part of GitHub's Agent HQ. For example, you can start by assigning Claude an issue, just like you would a teammate. Claude instantly picks it up, creates a branch, opens a pull request. And then from there, you interact directly with that PR. You can ask questions, request changes, iterate together. When it's ready, you can merge, and that's it.

This is Claude working alongside you, like a coworker, directly in your team's workflow. It'll write code, learn from your feedback and help you move faster.

ASHA SHARMA: Amazing. You have been so thoughtful about your model advancement, safety, everything you're doing in coding and enterprising and everyone you partner with. I'm sure everybody here and everybody at home would love to know, why is our partnership different than anyone else's?

From the very beginning, there just felt like there's been a lot of shared DNA and trust across both companies, and I'm really excited about what happens when you combine the power of our trusted models with Microsoft Foundry and M365. So we're making cloud available to enterprises that have already invested in Fabric and Foundry in M365, and removing barriers like navigating separate vendor contracts and building systems so developers can get started with cloud right away. I can't wait to see what you all build. Stay tuned for much more to come!

ASHA SHARMA: Thanks so much, Mike.

MIKE KRIEGER: Thanks, Asha.

ASHA SHARMA: Now, what's unique about this moment is that now Azure is the only hyperscaler to offer both OpenAI and Anthropic models to fully empower our customers to build the best AI applications on the planet. So, with all this choice, picking the right model for the right job to be done becomes even more critical. That's why today we're also announcing the

general availability of our new model router. This model router automatically selects the best model based on accuracy, performance, cost or balance. It enforces U.S. and E.U. data boundaries by default, and customers are already seeing 50% lower latency and 15% higher quality.

So with models multiplying tools everywhere, we've never had so much raw power. But raw power is not enough. Models don't understand your margins, tools don't understand your customers, and most solutions only understand the what of the work, not the why. The why lives in documents, lives in chats and meetings, and web pages, and files. And in every organization. This sits in two oceans of data, one that's structured and one that's unstructured. And between them lies this intelligence gap. Agents should not sit on top of these systems. They need to live in between them. They need to be able to see across silos, reason through context and act on behalf of people.

That's why today we're also announcing a new intelligence layer, one that knows how your company works, understands your business logic, and can take real actions in the real world. It is made up of three connected planes Work IQ, Fabric IQ and Foundry IQ. Now, Ryan talked about Work IQ. So let's start with Fabric IQ.

Fabric IQ builds on 20 million semantic models and Power BI that encodes your business logic. Now those models are going to extend beyond analytics and into operations. So humans and AI can share the same understanding of your organization in real time. And because Fabric IQ is integrated with Microsoft 365, every single user benefits, in Power BI, in Excel and Teams, and Copilot. All with security flowing automatically. To take us deeper, please welcome Zia.

ZIA MANSOOR: Thank you, Asha. Imagine this. You need to understand that Zava should open a new pop up store. Our store operations dashboard has an agent built on Foundry. Taking advantage of our data and Microsoft Fabric. You start by asking for markets with high demand signals. But the agent didn't understand what demand signals meant or where the more relevant demand data is. Results by just dropping an agent on top of your data can be mixed at best.

AI needs an understanding of what the data means and how things relate to take further action. With Fabric IQ, you can adopt the language of your business, defining entities and connect them across all your data in one lake. From real time data, semantic models and even mirror data from sources like Snowflake, Oracle or Google BigQuery. You're not just analyzing data. You're giving AI and your people a shared understanding of what it means so they can act on it in real-time, similar to what semantic models did for business intelligence, Fabric IQ is doing for AI. Making results consistent, accurate and transparent.

With the model built, you can now see a live, unified view of your organization and visualize how everything connects, from relationship graphs to supporting reports and other business data. You can also trigger rules to take action and leverage agents to help run your business through automated decision-making. Using Fabric's geospatial capabilities, you can instantly see how shipment delays impacts demand across regions and where opportunities are emerging. After adding Fabric IQ as a knowledge for our agent and Foundry, it looks like Boston is the strongest candidate for the next pop up store.

We now have confidence in the results provided by our AI solution. It has a deeper understanding of our business and provides much more relevant guidance. And look at that jump in quality. Same question, same agent. But now it understands us. Adding our business knowledge turns a generic response into something actionable and tailored. Now that we publish Fabric IQ, it's ready to fuel other solutions and any agent, including Foundry. From unified data to unified intelligence, this is IQ in Microsoft Fabric.

ASHA SHARMA: Thank you so much, Zia. Now, if Fabric IQ is shared business logic, Foundry IQ is the intelligent connection point across all of your structured and unstructured knowledge that your agents rely on and take action on. This is the first large-scale implementation of context engineering, extending RAG built for agents. And unlike traditional RAG, Foundry IQ doesn't just retrieve, it plans, it reasons, and it iterates across Work IQ, Fabric IQ, Blob Storage, the web and more.

So to put it to the real test, let's build a multi-agent system end-to-end, live on stage. Please welcome Elijah to bring it all together.

ELIJAH STRAIGHT: As a developer at Zava, my job is to make sure that our store ambassadors can focus on customers while our AI agents watch the store. Let's dive in. Here we can see our workflow with three different agents. We're going to talk about all of them today. But first I want to talk about our Zava store operations agent. This is the agent that knows everything a store manager could possibly need. Let me show you how I built it.

Every agent and Foundry starts with a model from the Foundry catalog. Foundry gives us same day access to the latest open source and frontier model releases, so we're always working with the best, the moment it's available. The Foundry catalog even lets me compare two models against each other to see which one is best for my use case. Today, I'm going to be using DSE v.3, which gives me the perfect blend of speed and accuracy for our low margin retail environment. Now, let's go ahead and create our agent. I'm going to name it Store Operations and I already have it ready. So let's go check it out.

So you can see here, first I gave the agent some instructions and then I connected a knowledge base using Foundry IQ. We're going to talk a little bit more about Foundry IQ here in a second. But first, let's give it a real task and help plan a sale for the upcoming SIM marathon in Sacramento. When you use the starter, prompt and kick that off.

While that's running, I'm going to use Foundry IQ. Knowledge Base and Foundry IQ contains a bunch of different pieces of information, including our product catalog, store inventory, employee handbook and local market intelligence. It's possible to do this from using a wide variety of data sources, including Azure Search Index, Blob Storage, Bing, SharePoint OneLake and Fabric IQ.

Fabric IQ, which Zia just showed us, gives our agents real time business context on our store operations. Now, if we go back to our agent, we can see here that it created a full plan for our marathon sale, pulling in info from previous events using Foundry IQ. But let's take it a step

further. I'm going to ask, what should I do if there's an item on the floor in the store? And what's going to happen is our agent will then pull from Foundry IQ to pull in information such as our safety policy documents and other information to give us the answer.

Foundry IQ supercharges our agents to be able to make critical business decisions quickly. So not only does this agent help us plan, but also react to what's happening in the store. Speaking of reacting, let's go to our next agent, which are video processing agent powered by the Foundry Video Indexer. This agent watches our store in real time and emits inventory events like if there's low inventory or if, whoops, somebody knocked something on the ground. I don't know who knocked that shoebox over, but we'll come back to that.

Zooming in on this agent, you'll see that Foundry's open ecosystem lets me run Microsoft Agent Framework, Landgraf, which I'm actually using here today, or other frameworks on agent service. This is all possible through our new hosted agents, which handles the deployment, lifecycle management and auto scaling for you.

Foundry's open ecosystem also gives me access to over 1400 enterprise-ready tools and MCP servers. And because it's all a-to-a compatible, it interoperates seamlessly with agents built in Foundry, Copilot Studio or any other compliant system. Let's go ahead and check out our store camera feed. And we see here there is a shoebox on the ground, and it says right here. And our agent picked it up, saying hazard shoe box on the floor. Pretty cool.

So now finally, I'll add our policy and notification agent, which uses our new agent memory to remember previous conversations and personalized messages to store managers. And so here we can see our agent memory, pretty great. And so, returning to our workflow, if the video agent deems an event requires action, and I just got a Teams ping, we're going to get back to that here in a second.

So, first our video processing agent will monitor the situation. And if it has an event that it feels like it requires action, it'll pass that to the store operations agent, which, using Foundry IQ, will add context around store layout, safety policy and risk level. And finally, the policy notification agent notifies the store manager.

And I don't think you sent me that ping. So let's go see who it is. And if we flip over to Microsoft Teams, there it is. Clean up on aisle six. Our workflow just notified us that there's a white cardboard box currently lying on the floor. Fully automated, real time, end-to-end. We just saw how easy it is to build multi-agent workflows with the all new capabilities of Microsoft Foundry.

Only Foundry brings together multi-agent fleets, hosted agents and live enterprise knowledge with Work IQ, Fabric IQ and Foundry IQ. Across any store, any region and at any scale. And it works with any model, any framework, and any tool. And with that, back to you, Asha.

ASHA SHARMA: Thanks so much, Elijah. With every platform shift, it redefines how we create. Computing made it personal and connected, and AI makes it ubiquitous. And that is changing how every company needs to operate. What you saw today expands what work can be.

Every person can now shape intelligence to fit their job, their team, their ambitions, and they can stop scaling by effort and start scaling by capability. And they can do that with the tools that they already know across M365, Copilot Studio, GitHub, Fabric and Foundry. Judson, back to you.

JUDSON ALTHOFF: Thanks. Awesome. Well, I'm excited to talk about ubiquitous innovation in real life and to do it with one of my favorite brands, Mercedes Benz. And I'm pleased to be joined by Daniela, the CIO of Mercedes Benz North America. Daniela, welcome to the Ignite stage.

DANIELA DIMITROVA: Thank you, Judson. I'm really happy to be on this stage with you together.

JUDSON ALTHOFF: So, Mercedes-Benz has truly been one of our best partners on this AI journey, and we chose to highlight our partnership here because you really have represented ubiquitous innovation across the firm. Maybe tell us a little bit about your journey with Copilot, GitHub and the other assets that you've deployed across Mercedes-Benz.

DANIELA DIMITROVA: OK. You know, next year we will celebrate 140 years since the brilliant engineer Carl Benz invented the automobile. But it was his wife who took it on a long ride and made it really popular. So innovation is deeply rooted in our DNA. And fast-forward to today, when we talk about AI, AI for us is our best digital colleague. Talking about Copilot. We started with 10,000 licenses. We quickly increased to 15,000. Now we are heading towards all employees have their own Copilot.

I extended my workspace with another monitor, so I have my Copilot working for me 24/7 with Copilot Studio. We have 52 areas in our company that are currently automating workflow and creating their own agents. My favorite example is the finance team, that without IT, without any coding, they automated the invoicing workflow.

You mentioned GitHub. GitHub is the place where the best minds behind the million lines of code in our company are meeting together. And since the introduction of the GitHub Copilot, we see 70% increase of the engagement. So the developers nowadays just focus on creative tasks, on more logic related topics. And the AI does the routine coding. So, with the Microsoft AI tools, we actually give a lot of power in the hands of our employees and they can do this really faster.

JUDSON ALTHOFF: That's awesome. So you've also done a lot of work with AI on the frontlines in the factories. Maybe you want to share a little bit about that innovation and the unlock that's provided for your employees there?

DANIELA DIMITROVA: Yeah, we are bringing AI also on the factory floor, and we have M0360 data platform that is connecting 30 plants, passenger plants, across the world. So we have real-time data, historic data, so we can analyze this data and improve the efficiency across the factory. With the digital factory chatbot ecosystem, now we have an intelligent multi-agent AI system, so we can get all the actionable insights in just an instant. Maybe we can have a look at a demo that we provided.

JUDSON ALTHOFF: Yeah, let's take a look.

COLLIN VANDAMENT: Thanks, Judson. As a quality assurance engineer, I'm responsible for ensuring we meet Mercedes' stringent production standards. And I've just been alerted to a decline in efficiency on the CLA 250 line, here in the Sindelfingen plant. In the past, diagnosing this issue would have taken days for a team of experts to investigate. But now, I can leverage M0360 Unified data and ask Digital Factory Chat to investigate the cause. The system engages specialized agents across disciplines that analyze production data and generate a consolidated report.

But I want to verify this information and understand the underlying analysis, so I can click show all messages to see the underlying data and the code outputs used to generate the final summary. This allows me to quickly and confidently validate the findings. Powered by Azure OpenAI models, Digital Factory Chat saves days of manual data investigation and empowers employees with a team of AI agents to solve complex challenges.

Mercedes is also leveraging AI to make their production processes more sustainable. Teams at Mercedes Rastatt plant use M0360 to identify opportunities to reduce energy consumption, which is exactly what they did in the paint shop. The topcoat process takes place in a paint booth where a full air exchange happens several times per minute. The new air supply for the paint booth must match strict temperature and humidity requirements for high paint quality. To balance the paint performance and energy efficiency, Mercedes created a simulation in Azure and trained in AI models to analyze temperature and humidity values.

The model uses real time sensor data to adjust operating parameters, ensuring optimal paint quality while minimizing energy consumption. By deploying this AI-powered solution, Mercedes reduced energy consumption by more than 20% and are scaling this across Rastatt.

Moving from operations to infrastructure, Mercedes is leveraging Nvidia Omniverse on Azure to build factory scale digital twins for six of their production plants, including this site at Kecskemet. Before making costly adjustments to align, engineers can design, test and optimize every stage of a vehicle's journey, unlocking unprecedented agility in planning for future market demand. It's incredible to see all the innovation happening here at Mercedes. Back to you, Judson.

JUDSON ALTHOFF: Thanks. So awesome things happening across the spectrum, across the employee base, across operations and innovation. But I think, perhaps, my favorite part is how you're enriching that engagement with customers. Share a little bit more about Hey, Mercedes. Yeah.

DANIELA DIMITROVA: Hey, Mercedes is our voice assistance on the wheels. And with Hey, Mercedes, we actually invent another workspace. We call it the third workspace. We have our office, we have our car. And now you can take our productivity also with you when you're on the drive. Hey, Mercedes is powered by Azure, OpenAI and Bing.

And now it feels really natural. You can check the weather in San Francisco, or if you're working parents like me, you can take your calls, you can check your calendar, you can take your notes. And all this magic happens just with the power of your voice. But I would suggest that we have a look how Hey, Mercedes works here, on our car.

JUDSON ALTHOFF: Super. Let's check it out.

MARTIN DUREJA: I'm excited to show you how Mercedes Benz is innovating with AI to improve our in-car experience. Hey Mercedes, I'm driving on the Golden Gate Bridge. Can you tell me more about this bridge?

HEY MERCEDES DEMO, AI VOICE: The Golden Gate Bridge is a stunning suspension bridge connecting San Francisco with Marin County, known for its iconic international orange color and engineering marvel. If you have a specific question about a topic, feel free to ask, and I'll assist.

MARTIN DUREJA: OK, who's the chief engineer of the bridge?

HEY MERCEDES DEMO, AI VOICE: Joseph Strauss was the chief engineer of the Golden Gate Bridge project. He was instrumental in its design and construction, although the final design was a collaborative effort.

MARTIN DUREJA: And can you tell me about the conference happening in the Moscone Center in San Francisco this week?

HEY MERCEDES DEMO, AI VOICE: Microsoft Ignite, a prominent technology conference, is happening at the Moscone Center in San Francisco from November 18th to November 21st, 2025. This event brings together IT professionals and industry leaders to explore advancements in technology and innovation.

MARTIN DUREJA: Sounds like an incredible event. Back to you, Judson and Daniela.

JUDSON ALTHOFF: Well, thank you so much, Daniella, for showing off all of the amazing things we're doing together between Microsoft and Mercedes Benz. It's been a pleasure.

DANIELA DIMITROVA: Thank you very much for having me on this amazing event.

JUDSON ALTHOFF: Thank you.

DANIELA DIMITROVA: Thank you. Bye-bye.

(Applause.)

JUDSON ALTHOFF: We love the partnership we have with Mercedes Benz and how it's unlocking ubiquitous innovation across everything that we do. We're inspired by how we can invoke new levels of creativity and innovation at every room in the house of the Frontier Firm,

from how we empower citizen developers, professional developers, enable rich collaboration with IT, and then provide the fundamental building blocks of context and intelligence through Work IQ, Foundry IQ and Fabric IQ to empower the success of the Frontier Firm.

But you may be wondering, how do we get all of this done inside of our company, Judson? How many of these assets from Microsoft do we have to have to pull off all of the magic that you've shown in this room? Well, when I said we need to democratize intelligence this morning, I meant it. And so, today we are announcing Microsoft Agent Factory. It's one way to harness all of the power of Work IQ, Foundry IQ and Fabric IQ through a single meter. What does that mean?

Well, what it means is that whether you're trying to automate claims processing or freight forwarding or supply chain management, you can meter the entire experience, measure the ROI for your outcome so that you know the ROI and AI upfront. You can leverage our forward-deployed engineers, our partners and our skilling assets to deliver the solutions that you've seen in the room today for your business. We're super excited about Microsoft Agent Factory, and we hope you take advantage of it.

So you might be feeling a little bit of tension in the room. And to be fair, the constructive tension between innovation and governance is real across many firms around the world, across all kinds of industries. And as you start to see all of the innovation happening around the room, you might be wondering, how do I manage all of this? How do I govern, how do I secure, how do I optimize? Well, it all comes down to observability. And to talk more about that, I'd like to welcome Charles Lamanna to the stage.

CHARLES LAMANNA: Good morning, everybody. It's great to be back at Ignite. And as you saw earlier from the demos with Asha and team, agents are showing up everywhere, embedded in every workflow. And these agents will be built by everyone in your organization will be key to getting AI value moving forward. And to put that scale into perspective, IDC predicts that by 2028, there will be 1.3 billion agents deployed. And one of the big challenges we'll all be facing as part of this is how to track, manage and govern all of these agents.

With agents becoming part of every process and team, we need to make sure they are secure and observable by default, so we're going to need new tools to safely unlock the full potential of agents. And we need these tools across IT teams, dev teams and security teams. Every part of the stack needs a single place to manage and secure agents. That's why we're so excited to announce Agent 365 today, now available in Frontier.

(Applause.)

Agent 365 is your agent control plane. It helps you safely scale your agents across the whole company, no matter where or how they were built. And these include agents from Microsoft, the agents you create, and even agents from partner clouds and companies. And to do all of this, Agent 365 unlocks five key things for you. First, it provides visibility into every agent in your enterprise, so you have a single trusted agent registry. Second, it will help enforce access controls to data and systems. That way, as information and processes move between people, agents and apps, you stay compliant and protected.

Third, it provides an easy way to see and understand your agent ecosystem, its usage and its risks so you can manage all of your agent estate in just one place. And fourth, it provides a safe and secure way for agents to interact with your data, your apps and your comms. We, of course, have first class support for all the Microsoft ecosystem, our apps, our data and even Work IQ, but it also interacts with agents from other partner clouds and platforms.

And last but not least, it will secure all your agents. It proactively detects and responds to threats. It also prevents data loss. And ultimately, just as you protect your apps, your phones or PCs, you want to protect your agents. And the best infrastructure for managing agents is the one you already use to protect devices and people, so all these capabilities integrate into the existing products and tools you already use for Microsoft today. Let's see all of this in action with a great demo from Jared. Over to you.

JARED SPATARO: Thanks, Charles. Good morning, everyone. As an IT leader here at Zava, I'm busy helping the business use agents to improve the way that we serve our customers. But as all those agents come marching into work, I need visibility and control. Enter Agent 365, the control plane for all my agents, and it starts right here in the M365 Admin Center. This is the same place that I already go to manage my users, my devices and even Copilot. And now I can come here to manage all of my agents.

The overview gives me a summary of all of my agents. I can see as an example, the agent registry here tells me how many agents I have across the company. I can also look at active users. Down below I can look at agent analytics. I can see the agent publishers, those created in my organization, those created by external partners. I can also look at the platforms that were used to create those agents, including Copilot Studio, Foundry, Workday, Google, ServiceNow and others.

Zooming out, I can look at active users over time and see that increasingly, people are truly using agents to get their work done, and I can even get a glimpse into those agents that are most popular in my organization. Up top here we see a claims processing agent, but we also see agents by Workday, ServiceNow and Manus. Now, all of this is built on top of a foundation that starts with what we call the agent registry. The registry, as I click into it here, gives me a complete list of all of my agents. You can see them listed out here, each by name, platform, security risks, even data that's important as I'm trying to understand exactly how these agents are used and how they're providing business value.

Now, the best part about this view is it allows me to focus my time and my energy, and I do that by coming into agent risks here. Here I'm able to see the agents that need my attention now. For instance, this comms agent looks like it has a couple of security risks. I can see the platform it was built on, the owner, and sure enough, there we go, those two security risks. And by clicking in, I can see the details.

It looks like the agent has an abnormal sign-in frequency, perhaps even being used by a risky user. But that's not all, this agent is protected comprehensively. Here we see Microsoft Purview protections that include monitoring the agent activity, sensitive data, and evaluating compliance

gaps. We can even see the activities monitored over here and get a sense for some of those sensitive interactions that have been prevented. I even have the ability at just one click to block the agent at any time. I am in complete control.

Now, when there's a new request for a new agent, that request gets routed to me up top here. And here I see a list of agents for my review. Clicking into the staffing agent, I can see that same set of details, including, for instance, data and tools that the agent has access to and even permissions here. I can look at those permissions and if everything checks out well, I can publish it for immediate use. Perhaps my favorite part of Agent 365 is the ability not just to see one agent, but to zoom in and actually get a sense for how all those agents fit together. This is what we call the agent map.

Here I can see agents by platform, Foundry, custom, Google, Amazon, and as I come down and click on a particular agent, I can actually see the other agents that it works with to get the job done. I'm also able to see a little bit about hotspots. I'll click into this one, zoom in, and sure enough, look at that. There's a high exception rate on this particular agent. I'll zoom out and see if I can do something about that. Let's see.

There we go. There's some detail I'm able to see. Sure enough, there are some errors in that magenta exception, but I have a hunch about this one. Let me go to the overview. It was built on Foundry. Oh, there it is. I'll zoom in. My good buddy, Seth Juarez. I love this guy, but he's pretty busy and not always attentive to details. Might be something that we're going to have to look at just a little bit later.

Now, there's so much to show you of this product and so little time, so I'd encourage you to learn more throughout the week. But before I go, if you remember just one thing, remember this: No matter where your agents come from, no matter what they do or how they do it, Agent 365 has got you covered. It gives you the visibility and control you need. Back to you, Charles. And, hey, make sure you take a closer look at that agent, the one that Seth did.

(Applause.)

CHARLES LAMANNA: Thank you so much, Jared. Awesome to see Agent 365 in action. And I'm super fortunate to be joined by three of our launch partners for agent M365: Tao from Manus, Dean from Workday and Paul from ServiceNow. Thank you so much for joining us. I thought we could maybe have a conversation about some of the things you're seeing from users of agents and companies today, so let's get started. Tao, I look and I see what are common challenges and opportunities that companies are facing as they integrate AI into their workflows and their agents right now. OK.

TAO ZHANG: Right now, we are still like a consumer-focused product because a lot of people are using Manus for their work, so definitely, we got a lot of questions and inquiries from enterprise too. And from the early feedback we saw that from the enterprise perspective, the first thing they care about is about how can they integrate Manus into their existing workflows. Because right now, even though Manus is great, because every company already has a lot of IT

assets, they want Manus to have some way to integrate into the existing workflows, like how can they access the files? How can they operate their internal systems? That's the first part.

The second part is always about like compliance and security. Because the enterprises just want to know that because we are generating, actually, we can do a lot of different things. So the idea that we just want to make sure that the agent is accessing the right level of permissions about the files, about all the data, that's the second thing. And the third thing is about they want to evaluate the benefits, like what is the agent doing, what's the outcome? So these are the three things they are asking for.

CHARLES LAMANNA: It makes a ton of sense, and Manus is an awesome product. Dean, one of the things that we also hear from a lot of customers is this idea of shadow IT coming back in the era of AI and agents. What are you seeing for customers who deal with that?

DEAN ARNOLD: Yeah. I think shadow IT is really interesting. With the rise of AI and agents, we're moving from shadow IT to shadow AI, and this happens when in the businesses, we don't provide our employees the adequate tooling they need to do their work day to day. They want to be effective, they want to get job done, so what do they do? They go and pull out their personal AI. They bring out their Copilot, they bring out all the different tools they have available to them, and they use this at work.

So this can create certain challenges around for the enterprise, like data loss or governance and compliance problems of like GDPR and HIPAA, and that's something we need to solve. Now, how can we do that? As enterprises with IT and business technology teams, we should be going from guardians to being enablers in the workplace, so the best thing to do is to reach out to employees, understand what they're trying to do and provide the tools with them.

And if the tools have deep integration with their productivity suite like Microsoft 365, then that will become the default path. That's their fast path, and that's the path they want to use, and they'll adopt that rather than our own tools. And that's great because that allows us to combat shadow AI and shadow IT. I would also say that as agents become more prevalent, they become more capable, and they take on larger tasks.

Adopting tools like Agent 365 and at Workday, we have our agent system of record, and you can use those tools to be able to register agents and add kind of governance layers, and enabling that with solutions for security, like a Entra with agent ID as well. That allows us to give identity to agency, know what they're doing, who's interacting with them, and derive value. And I think together this creates a holistic solve for how we deal with shadow IT and shadow AI.

CHARLES LAMANNA: Makes a ton of sense. Make the safe and secure path the easiest path, and you get more usage. Now, Paul, let's finish it off with a question in terms of how you're seeing companies balance the risk from a compliance or protection point of view to the value of agility and quick innovation with AI. How are you seeing people think through that?

PAUL FIPPS: Yeah. First of all, Charles, I just want to say a huge thank you for having us up here at your fantastic user conference. Our relationship with ServiceNow and Microsoft has just

exploded in the past couple of years, so great to see what we're doing and what we announced today. From an observability standpoint, what I'm hearing from every executive out there is, how do I move fast but stay in control? Almost like an F1 driver, right?

And the reason is because these agents are rolling out faster than anybody anticipated. And if you think about it, I move too fast and I break things, I move too slow and I get left behind. I'll give you a great example. AstraZeneca, a great customer of ours, is using ServiceNow AI Control Tower to manage workflows across lab work and operations, but they're using a Control Tower to get visibility into those agentic workflows and make sure they can trust them. That's driving real results. Actually, they're saving 90,000 hours they give back to researchers to build those lifesaving drugs that we all benefit from.

Now, the reason they do that is because they can see AI so they can trust AI and then they can scale AI. And what we see happening in the importance of scale is that we think there'll be more AI agents in the very near future than human agents, and if you don't have something like this, it's really hard to manage AI. And so, you can see unchecked work, you can see security issues, all the things that will happen without this. We think about observability as a difference maker between accelerating your business with AI or slowly and in some cases not so slowly eroding it.

CHARLES LAMANNA: Really incredible. I think kind of all the stories that you were sharing today, based on the partnership we're doing with agent M365, so much more is possible. Thank you again for joining us. We really appreciate you being here today.

PAUL FIPPS: Thank you.

CHARLES LAMANNA: Now, it's not just Manus, ServiceNow and Workday that are integrating with Agent 365. We're committed to an open and broad ecosystem. Today, we already have dozens of other large companies and small AI-forward companies integrating with Agent 365. That includes classic app vendors like SAP and Adobe, as well as AI Frontier Labs like OpenAI and Anthropic, who integrate with Agent 365 to make it easier to get your agents deployed.

Now, Jared's demo earlier only scratched the surface of what you can do in Agent 365, but it's not just going to be IT and security teams that need to manage their agents. Developers will also need to manage the agents that they build. And today, developers are already in Foundry building agents and training models, but they need to go beyond that and start to monitor and manage the agents that they're building, too.

That's why we're so excited to share the brand-new Foundry Control Plane capabilities that will allow you to observe, control and govern all the agents that you're building. Dev teams will be able to scale more agent projects without losing control. Now, earlier in his demo, Jared showed us the procurement agent generating a whole bunch of exceptions, and that agent was built in Foundry. Let's see what additional observability is possible by having Seth take a look. Over to you, Seth.

SETH JUAREZ: Hey, thanks. I'm developer lead at Zava. I don't know if there's a whole bunch of exceptions, OK? I'm the developer leader at Zava. Thank you. It's always the devs fault, right? But I'm constantly building and managing innovative agents, which can make it tough to prioritize my time. That's why I rely on the Foundry Control Plane, which you can see here.

The Foundry Control Plane lets me keep an eye on everything happening across all of my agents. Prioritized alerts help me stay focused on quality as well as other security threats, which is pretty cool. Earlier Jared said there was something going on with one of our agents. I'm pretty sure it's a user error, but I think we should check it out just in case. Notice I do have an evaluation pass rate alert for the procurement agent, so I'm going to go to view and build.

And what that's going to do is that's going to take me to the procurement agent dashboard. We've been using it a lot. Looks like -- oops. Looks like there was a little bit of an error rate. And if I scroll down, we can also measure things like groundedness, task adherence, task completion. And there are certain metrics that measure from one to five, so we're doing good here. Task adherence looks like drops, and that is a little concerning, so let's go to the playground and see if we can't figure out what's going on.

I have a suspicion about what's actually happening, so I'm going to ask you to generate a procurement plan. Let's cross our fingers and hopefully we don't have -- oh-oh. Oh, man. I think I know what happened. This is really embarrassing. As I'm looking at the actual MCP server, I think I have the dev one there, and now my AI agent is putting me on blast, too. This is not good.

Thankfully, we have Foundry Guardrails that allow us to block these kinds of things. Guardrails allow us to mitigate things as they actually happen. And task adherence is a super important measure of how closely an agent's actions match the user's goals, so I can actually fix this by going down to guardrails, changing the guardrail. I'm going to assign a new guardrail to task adherence guardrail, and notice that as I scroll down here, I can actually just block this. This agent is really good at a lot of stuff. I just need to switch it to the right MCP server. Don't tell Jared.

Second thing, I've been asked to manage costs a little bit. We want to make sure we get this under control as soon as we can. And the good thing is that we can actually, maybe switch models without sacrificing fidelity. Evaluations let me confidently compare different models across each other so that I can actually see, hey, can I actually switch models and the quality side by side?

For example, this is using GPT-4.1 and this is GPT-4o mini. Notice that there's even an improvement here. With cost, I can actually save 13x by switching models while still knowing that this is something that I can do, which is pretty cool.

Last but not least, thank goodness we have here a jailbreak attempt was attempted on our actual agent, but the good news is that Foundry already blocked it. The bad news is that hopefully this isn't a pattern. What we're going to do is we're going to notify our Zava Stock Analyst, because this notifies Microsoft Defender if this is something that needs to be taken a look at.

Why don't I leave it to the adults in the room to take care of that? Back to you, Charles.

CHARLES LAMANNA: Awesome. Thank you so much, Seth.

(Applause.)

Seth just showed us how Foundry Control Plane blocks a jailbreak attempt. Bad actors target employees every day, and now they're going to be targeting agents, too. You need tools that understand these new threats.

At Microsoft, we're extending our existing security capabilities that we have for identities, supporting over a billion users a day, as well as a million and a half customers, to make them support agents, too. This includes threat protection from Defender, identity from Entra, and data security from Purview. And it's all fueled by threat intelligence from over 100 trillion signals every single day. And all these shared security primitives have been woven into Agent 365 to seamlessly protect both people and agents.

Now, let's go show you how a security professional would navigate this new threat landscape. Irina, take it away.

IRINA NECHAEVA: Thanks, Charles. I'm in Zava security operations, and my job is to protect users' devices up and now all Zava agents from bad actors and insider risk. What I love about Microsoft Security in Agent 365, I can do it all in my flow of work, which is this Defender portal.

Remember that Foundry alert about the jailbreak attack? Well, it's already here in my incident queue. That's because Defender and Foundry share signals in real time and join forces in protection agents.

Now that I know that this agent was targeted, I want to understand what happened. And as I expected, I see other alerts related to this incident. Well, that's because most attackers try multiple ways to break in. We'll investigate further.

This incident graph shows me the full scope of the attack, where it originated from, how it spreads from my environment and how suspicious entities connect to the procurement agent. To get the full picture, we're going to look deeper into each of those alerts, starting with the jailbreak attempt from Foundry.

Here, Defender reminds me that a jailbreak is a type of attack that tries to bypass AI safeguards and trick an agent into doing something it shouldn't. Defender confirms that Foundry Guardrail blocks this attack.

Next, I'm going to loop deeper into this high priority phishing alert. And here it looks like somebody tried to send a procurement agent a malicious web link, but Zava policy blocked this attack.

And with all the sensitive data in procurement, I'm a bit concerned about this data security incident, but I know that Purview can identify and protect sensitive data. And I love that my investigation goes much deeper, because all Purview data security alerts are available to me right here in Defender.

In the investigation view, I see that someone instructed our agents to share a secret file about Project Orion, but Purview blocks this action and prevents the data loss. My investigation is complete. I confirm that a malicious act had targeted procurement agents, but Zava defenses blocked this attack. I can now resolve the incident and let procurement agents go back to work.

Now, back to Charles.

CHARLES LAMANNA: Awesome.

(Applause.)

Really incredible to see. Thank you, Irina.

Microsoft Security helps make your agents enterprise ready with minimal effort. The idea is that you could become a Frontier Firm with Copilot and agents directed by human ambition to go transform every part of your organization. And all of this needs to be secured and observable at every layer of the stack. And that's where Agent 365, Foundry Control Plane and more come in.

Thank you everyone, and right back over to you, Judson.

(Applause.)

JUDSON ALTHOFF: Thanks, Charles. Microsoft has no more important partner than Adobe, and I'm delighted to be here with Adobe CEO, Shantanu Narayen.

Shantanu, a warm welcome to Microsoft Ignite.

(Applause.)

SHANTANU NARAYEN: Thanks for having me, Judson. This is a pretty amazing show. I've been a huge Warriors fan, but being on the floor is truly a special experience.

JUDSON ALTHOFF: Right. Well, there's high expectations now. We'll see how you can do on your layup.

We've talked a lot about this notion of the Frontier Firm. We've talked about ubiquitous innovation. We've talked about AI and the flow of work. We've talked about observability, and I

know that Adobe shares many of the same views. Do you want to talk a little bit about how Adobe brings to life AI through your strategy and your customer orientation, and how Copilot and agents come to life inside of Adobe?

SHANTANU NARAYEN: Sure. First, congratulations on all the incredible announcements. When I read and hear about everything that you're doing with M365, with Copilot, with Azure AI Foundry, I think across all of those dimensions we're partnering.

As we think about Adobe, we really have three customer constituencies where we are focused on as we think about empowering everybody to create. First, on the creative side, I think everybody would acknowledge people use our products to do any creative endeavor that they want. On the business productivity side, what we are doing around Acrobat and PDF, there was a lot of conversation about unstructured information in your keynote. That's the other area.

Perhaps, the area that a lot of your customers may not know as much is we're also one of the largest providers of marketing technology in the world, something that we've partnered on from day one. As we think about what we are doing around AI, I think the themes that you outlined is exactly what we are doing, which is, first, how do you bring AI into the application? How do you think about ubiquitous access to AI within the applications?

And we do that across the data layer, the models layer, the agentic orchestration layer, as well as what we are doing with the core infrastructure. It's really great to see the multiple areas in which we can partner.

JUDSON ALTHOFF: That's awesome. And you just had Adobe MAX a couple months ago and had your own slew of announcements that were really magical. I think, suffice to say, the pace of innovation is really frenetic right now. Talk a little bit about our partnership and how that comes together in light of all of the new things you're releasing to market.

SHANTANU NARAYEN: Well, as I was sitting through this keynote and just seeing the amount of announcements that you made, it's mind boggling. And we're trying to do the same, which is, how do we innovate at the pace of AI for all of our customers?

A couple of the key announcements that we have made: First, as you talked about AI Foundry, within Photoshop right now, you can actually support any third-party model. As you talk about Model Router, what we have done is within Photoshop, you can pick whether it's Adobe Firefly as a model, whether it's MAI that you're going to work on, OpenAI and other partners. The models and supporting all of these models, because they'll all have slightly different characteristics, is an area where we made a significant amount of innovation.

On the agentic side, you know, making sure that the task that you want to perform so that the model can actually reason on your behalf, and you can do workflow, whether it's in Photoshop, whether it's in Acrobat or whether it's in our marketing products, we've also released a huge amount of agentic stuff. And I'm really excited about what you talked about with the ability to do governance with Agent 365. And so, a significant amount of agentic innovation that we have delivered.

And then finally, on the marketing side, I think the area where we have partnered from day one, as you're working on sales, as you're working on service, and we think about for every knowledge worker in the enterprise, how they need to be both productive and creative. I think that the work that we're doing together to make sure that we can embed creativity and M365 as well as integrate our marketing stack with what you're doing on the service and sales side, I think those are just a few of the areas that we're working on together.

JUDSON ALTHOFF: Indeed. And we've been using Adobe prolifically throughout our marketing organization at Microsoft, also doing co-innovation and engineering work, even with Firefly. And it's really important how Adobe takes on this notion of enterprise grade. It matters to Microsoft. It matters to our joint customers. At the core of that is trust. Maybe your view on how AI has to come through and deliver on the promise of trust in the enterprise.

SHANTANU NARAYEN: I think that's so true, especially for the people in the room today, because if you're developers, you're IT professionals, and you're thinking about enterprise grade software and how you release this for customer engagement, which is an area that we're really interested in, for knowledge, productivity, scaling this so that it's a truly scalable solution across all enterprises is super critical.

We're thrilled that you are deploying our solutions. The reverse is true, which is, we use Copilot every single day in order to make our developers, or GitHub Copilot. But I think the biggest issue that we're working on together, which I'm really proud of, is what is the data governance? How do people get assured of the fact that all the data that they are entrusting, whether it's to the Microsoft or the Adobe solutions, the governance around that is absolutely iron clad.

We were thrilled that we announced this thing called the Content Authenticity Initiative, and that's all about saying, we respect the intellectual property that the creators have. And so, the ability to make sure that their partners, both LinkedIn and Microsoft, are partners. We now have 5,000 partners, so dealing with this issue of content provenance and trust, both within an enterprise as well as outside the enterprise, I think that's an area that we can continue to innovate pretty dramatically.

JUDSON ALTHOFF: It's awesome. Well, look, it's a fantastic partnership, one deeply ingrained in engineering and joint work together, a partnership with a purpose and a partnership that, above all, stands for customers.

Thank you so much, Shantanu, for being here with us today. I really appreciate it.

SHANTANU NARAYEN: Thank you for having me and congratulations on all your innovation.

JUDSON ALTHOFF: Thank you. Appreciate it.

SHANTANU NARAYEN: Thank you.

(Applause.)

JUDSON ALTHOFF: We've talked a lot this morning about the three traits of the successful Frontier Firm, AI in the flow of human ambition, ubiquitous innovation, observability at every layer of the stack. But at the end of the day, you have to bring the magic to life. You need to have the core cloud platform and all of the services that enable AI to reach millions of customers around the world, unlock that human ambition and really help deliver on the promises of all of the innovation that we talk about and your ambition.

You need the cloud and AI platform that lights that up, and there is no one better than the man in the red shirt himself. Please join me in welcoming the one and only Scott Guthrie to the stage.

(Applause.)

SCOTT GUTHRIE: Thank you, Judson. Well, thanks, everyone. It's awesome to be here.

Now, everything you've seen so far in this keynote runs on Microsoft Azure. It's the infrastructure that is powering all of the Microsoft Cloud and all of our AI services. And Azure provides both the infrastructure and the cloud platform that you can use to supercharge your organizations and make them Frontier.

Microsoft has over 70 Azure regions worldwide today, and offers data residency in 33 countries, more than any other cloud provider. And we're continuing to dramatically expand our data center capacity so you can put your apps and your services closer to your customers and employees everywhere around the world. And we're providing a clear sovereignty commitment giving you control over where your data and your AI lives and how it's managed. And we're optimizing Azure to be the world's AI supercomputer, and we're delivering AI innovation across the entire system.

Last week, we unveiled our latest Azure AI data center in Atlanta. This Azure data center is the size of 20 football fields and contains hundreds of thousands of the latest Nvidia Grace Blackwell GPUs. It has the highest density of GPUs of any data center in the world. Let's watch a video of it.

(Video segment.) (Applause.)

Now, this one Azure data center delivers 10 times the performance of the world's fastest supercomputer. And this step change is what's accelerating the next wave of AI model improvement and AI capability breakthroughs. And AI infrastructure like this is also enabling exponential reduction of AI cost. You look at where we were even two years ago compared to today, the price of a model like GPT-4 has plummeted by 93%. And continued infrastructure and model innovation is going to continue to drive cost reductions and in turn, enable you to leverage AI for even more use cases.

Now, we recognize, with all this power and scale, also comes responsibility. Microsoft is now one of the largest buyers of renewable energy in the world, and we're on track to meet our goal to have our data centers powered by 100% renewable energy by the end of the year.

(Applause.)

And this includes the data center you just saw in that video, as well as all the others that we're building around the world.

We're not just using renewable energy, we're also using innovative sustainability approaches to ensure zero water waste. The data center in that video is a liquid cooled facility, meaning we cool the servers using water, as opposed to traditional air coolers.

And in this picture here, you can see the liquid chilling system that feeds cool water into the data center. And we use what's called a closed loop system to ensure that the water is continuously reused with no evaporation or water loss. The initial fill of the water into the system equals only about 20 homes' worth of water consumption. And then once the initial fill is done, we'll reuse that water for six years without ever having to add any more.

(Applause.)

With Azure, we're setting a new benchmark for sustainable operations.

Now, GB300s are the latest NVIDIA GPUs, and they're the most advanced GPUs in the world. GB300s deliver 11 times the performance of Nvidia's previous generation H100s, and delivers 12 times the performance of AWS' Trainium 2. Microsoft Azure is the first cloud provider to deploy GB300s in the world, and Azure has more GB300s in production than any other cloud provider today.

Customers are already running production workloads on GB300s in Azure. Here, you can see a cluster with tens of thousands of GPUs in it. Being first to run the latest generation of Nvidia hardware means that you can train and inference AI models faster, more efficiently and at a dramatically lower cost.

And we're investing in differentiated silicon across both GPU and CPU workloads. Microsoft Azure's Cobalt silicon delivers industry leading ARM 64 price performance compute in the cloud. Cobalt based VMs provide up to two times the performance improvement with .NET apps. And developers from Snowflake, Databricks, Elastic, Adobe and our own Microsoft Teams are already taking advantage of Azure Cobalt in production, seeing up to 45% better performance, which translates into 35% fewer compute cores and VMs needed, which ultimately saves a lot of money.

And today, we're excited to announce our new Azure Cobalt 200 offer. Azure Cobalt 200 is a major leap forward, with more compute cores, more cache, faster memory and higher performance, all built on the latest TSMC process technology. Cobalt 200 will deliver up to 50% better performance and will be Azure's best price performance VM in the market.

Now, all of this Azure infrastructure innovation, AI model improvement and cost reduction is powering a new era of apps and agents. And it's going to enable all of you in the audience to integrate AI into every workflow and build transformative solutions that weren't possible before.

AI leaders and more than 95% of the Fortune 500 are already building differentiated solutions on Azure. This slide here just includes a few of them. Now, one of the companies is OpenAI, who builds ChatGPT. ChatGPT is built on Azure and uses the exact same Azure services that you can use as well. Services like Azure GPU VMs, Cosmos DB, Azure Kubernetes Service, Azure PostgreSQL and Azure Storage. ChatGPT needs to be able to scale their application tier across tens of millions of compute cores around the world, and they do it with just about a dozen engineers, which is pretty incredible.

Now, how do they do that? Well, they do it with AKS, which provides a highly scalable Kubernetes service for cloud-native applications. AKS is a fully managed Kubernetes service available in every Azure region, and it streamlines operations at any scale. AKS offers automated deployments, auto-healing, automatic patching and built-in security safeguards. And these enable applications like ChatGPT to scale without significant operational resources.

And with our new AKS Automatic capability, we're making it even easier for teams to get started with Kubernetes. AKS Automatic is built in best practices for security, reliability and governance. And it allows you to go from code to fully deployed Kubernetes clusters in minutes. I'm excited to announce that AKS Automatic is now generally available as of today.

Now, data is the fuel that powers AI. Great AI solutions are built on highly-capable data platforms. If you want to build a solution like Microsoft Copilot, or ChatGPT, you need a database capable of storing petabytes of data, handling trillions of transactions and supporting limitless growth. And Azure Cosmos DB provides that.

Azure Cosmos DB is a globally distributed, multi-modal database service with guaranteed millisecond latency and uptime. With Cosmos DB, we built a database service that can automatically replicate your data to any Azure region around the world to give your users lightning-fast performance, regardless of wherever they're accessing your applications.

And with both Copilot and ChatGPT as examples, as users interact with the apps, conversations, prompts and metadata are stored using Cosmos DB. And this enables these apps to maintain context across sessions for hundreds of millions of users, delivering a natural user experience with low latency and high reliability at truly global scale. Cosmos DB also allows you to elastically scale your storage and performance throughput with zero application downtime.

You can start small and scale to exabytes of data and trillions of transactions per day. And you can start with processing, say, just 100 operations per second, and then scale to millions of operations per second if you need to. And best of all, with Cosmos DB, you only pay for the storage and performance throughput that you actually use.

Azure Cosmos DB delivers incredibly fast response times in five nines of availability, meeting the most demanding performance and uptime needs for any application. That's why innovative leaders across industries and around the world are using Cosmos DB to power their most critical apps. This slide includes just some of the logos of companies using Cosmos DB today.

Now, we also know that customers want to be able to scale without limits with the relational databases as well. And today, I'm really excited to announce that we're introducing a new cloud-native database solution to our Azure PostgreSQL offering family. Azure Horizon DB enables you to power mission-critical apps with performance and speed at any scale. Horizon DB is fully compatible with Postgres, and it supports scaling out a Postgres database to over 3000 cores and 128TB of storage.

It supports sub-millisecond, Multi-zone, commit latency. And horizon DB also allows you to build smarter apps using database AI, vector indexing and semantic search all integrated with Azure AI for seamless innovation. Now, data is everywhere. It's stored in NoSQL databases, relational databases, and unstructured sources. It's in every app, every system and every user interaction. And while that data holds incredible value today, it's too often fragmented, making it hard to manage and even harder to use effectively.

One of the biggest things that we hear from all of you is the need to bring all of that data together into a single, cohesive foundation. You can unify your data. You can unlock better analytics and provide the fuel that powers your AI solutions. And that's where Microsoft Fabric comes into play.

Microsoft Fabric is an end-to-end data platform that takes you from raw data to AI and by value, all in one unified experience. It brings together purpose built workloads for engineers, analysts and BI professionals. Earlier in this keynote, Asha announced Fabric IQ, the semantic layer that gives your data meaning.

With Fabric IQ, Power BI and Copilot agents can now share the same underlying understanding of your business. And Microsoft Fabric is built on OneLake. Our unified AI-powered data lake that brings all your enterprise data together in a single, trusted foundation. OneLake provides unified AI-powered data lake for all enterprise data, structured, unstructured and everything in between. I'd like to do now is turn it over to Patrick to show a demo of how Fabric and one lake work in action.

PATRICK LEBLANC: Thanks, Scott. I am so excited to be here to talk about OneLake, and with all the IQ that we've talked about, I feel like my IQ has gone up just a little bit. All right. I'm Patrick LeBlanc, the data architect at Zava. And earlier in the Fabric IQ demo, we showed you how we could uncover product sales inventory insights across multiple platforms.

The reason that worked is because all the data that Foundry was working with was already unified in the OneLake. OneLake unifies all the sources together in a single place. So let me take you on a little tour behind the scenes to show you how we get this. So, I'm going to head over to my Fabric workspace, and you'll see that you can take a look at this task flow. And I've mirrored

data from platforms outside of Fabric. And I also have a Fabric SQL database as part of the solution. All the data landing in the OneLake.

When I click on the unify model, I'm going to jump directly into my lake hub. And with Microsoft Fabric, all your data, no matter where it lives, comes together in one place. I've created shortcuts to BigQuery, to Databricks, Oracle, Snowflake and even SharePoint. Instantly, they all appear side by side in the OneLake. No copies, no movement, just single seamless Access.

Now, if I walked in here and I had a magician's outfit on, you probably like, what's going on with Patrick? This is not a costume party. But what I'm about to show you is truly magical. I'm going to switch over to my SQL analytics endpoint. And when I go in this endpoint, what you're going to see is every table that I created will automatically be available to me. And then what I can do, I could write a simple query with some simple selects and joins. But I decided to write something more complicated, something to really push this engine.

I created a CTE with some aggregations and some groupings and some rankings. And when I run this query, what you're going to see is cross-cloud data unified into one lake with the simplicity of SQL all in just a few seconds. What we're giving you is, instead of it depends, we're giving you it's done. That's Fabric. That's OneLake. Back to you, Scott.

SCOTT GUTHRIE: Great. Thank you. Patrick. So the services that we've talked about today in this keynote are incredible for building new AI apps. We also know that most of you have existing apps that you want to infuse with AI as well. And to do this, you want to be able to modernize these apps quickly.

One of the ways we're enabling this is with some of the new capabilities that we're bringing to Azure Copilot. Today, we're announcing six new agents in Azure Copilot that help you migrate and modernize existing workloads even faster. Our new migration agent in Azure Copilot helps you create landing zones and gives you the flexibility with both IaaS as well as platform as a service modernization options, so you can move existing on-premise apps to, say, VMs. Or you can alternatively use managed platform services like AKS, Azure App Service and all of our Azure database services. Bruno, show us a demo of Azure Copilot in action. Take it away, Bruno.

BRUNO BORGES: Well, thank you, Scott. As part of our push at Zava to become a Frontier company, we are focused on using the most modern tech to power our systems. But recently we made an acquisition. And now we need to bring some legacy systems to the cloud. As a cloud architect, I'm using AI to make that happen as fast as possible. I will start with the new Azure Copilot migration agent to modernize this workload and move them to Azure.

First, I will ask the agent to show me the discovered inventory. As you can see, Azure Copilot created a full inventory of everything running on my VMware data center. It has servers, databases, .NET applications, and Java applications. Now, I am going to ask the agent about one specific application I am already familiar with. It is the Zava Inventory Asset Manager. This is a Java application currently running with PostgreSQL.

The immigration agent recommends a path approach for this workload. The database should move to Azure Database for PostgreSQL, and the web application should be containerized and deployed to Azure Kubernetes Service. Now, I'm going to go and click on Go to assessment, which opens a readiness report for this Java application. Here we can see the options. We can lift and shift to the virtual machines. Or we can use paths. We'll stick with the recommended paths approach for now.

Next, I can connect this application with GitHub and have GitHub Copilot assess the code and add insights for modernization. I'm going to request this by creating a GitHub issue to generate those code insights in the background. Now, once on Visual Studio Code, I can see those insights and use GitHub Copilot app modernization to run the containerization task.

GitHub Copilot drafts a plan and then executes on it by generating all the artifacts such as Dockerfiles, documentation and the container images. At this point, the application is built, containerized and tested. GitHub Copilot can even deploy the application for us. So here it is. The application is now up and running on AKS. Thank you. Back to you.

SCOTT GUTHRIE: Thanks, Bruno. Microsoft Azure is a global platform, the world's AI supercomputer, and an integrated cloud foundation. It provides you AI-ready infrastructure that lowers your costs as you embark on your journey to becoming Frontier. I'm looking forward to seeing the great solutions you build with it. Back to you, Judson.

JUDSON ALTHOFF: Thanks, Scott. Well, we have one last great customer to share with you here live. I'm joined by Paul McEwen through the magic of Teams from UBS. Paul, thanks so much for joining here and a warm Microsoft Ignite welcome to you.

PAUL MCEWEN: Excellent. Thank you so much for having me.

JUDSON ALTHOFF: So, Paul. Look, UBS is one of the best adopters of cloud and AI in the financial services industry. And we've had this amazing partnership that we've shared over the years. It's no secret that you're ahead in AI, because you were also ahead in the adoption of cloud capabilities. And you've worked this strategy from end to end. Can you share with us a little bit your approach?

PAUL MCEWEN: Sure. So, you know, I'm happy to say that DBS is one of the largest global wealth managers. We have over 50% of our workloads in the cloud. We have created over 100 Azure services. For a firm of our size, this represents a very big footprint. And as you say, today we are one of the largest users of Azure in the financial services industry. And while we've come a long way, we still have a long way to go, and we've been at this for a long time.

Our adoption journey started back in 2014, long before our first apps were running in the cloud. And this is because we were focused on a massive infrastructure modernization program. This really did pave the way for where we are today, and where we are going in the age of AI. For us, the cloud was never about lift and shift. It's about reimagining how we build our applications to

better serve our clients, manage the security of their data and the service they receive. This is central to everything that we do today.

Just as when we started our cloud adoption journey, the biggest challenges are often not technical, but humans. Getting people to think about their applications as cloud native first, which is not just where the app is hosted, but how you run the app as well. Now, we're in an era where people expect to have computing on demand. They expect to have things like elastic compute, automated provisioning and all the great features that Azure offers to us. There are things which of course we offer, and now people take for granted.

But I think one of the key things that this journey has allowed us to do was as part of our Equity Suite Integration, we were able to move a greater pace. The services offered from Azure allowed us to move our platforms across the world and our users across the world, and create great efficiencies for us.

JUDSON ALTHOFF: Well, Paul, I want to thank you for being here with us today. And when I sort of ask you this last question, we'll turn it over to demo. But you've talked about the importance of data as the foundation for everything that lights up AI. You've built this amazing stat solution that helps your wealth managers connect with customers. Maybe you share a little bit about the demo. And thanks again for joining us here at Ignite.

PAUL MCEWEN: Yeah, absolutely. I think that the first thing for us is we need to get our employees to think about data and how important data is in the world that we're going to now in the era of AI. The key thing for us is that we are a highly regulated organization, so we take the security of our data and the regulations we have to follow around our client information very, very seriously. Part of ensuring that data foundation is we work through programs like our enterprise database.

But we're also focused on a lot of data governance, which unfortunately, in the regulated world we have to focus on governance, and as you'd expect, that becomes part of our foundation as we start to build these things. And what I recommend always to everybody is that when you build these foundations, whether you're a company of 100 or a company of 100,000, as part of the AI-first mindset is how do you structure and how do you know what the data is that you want to use?

I believe that our AI journey today is built on the back of the cloud journey that we started back in 2014, and our investment in the cloud and our partnership with Microsoft has allowed us to jump to the next paradigm, which is AI. One great example is an app my colleague Jim will take you through, and just one example of how we're serving our clients through AI-enabled applications with Azure services as the foundation.

JUDSON ALTHOFF: Awesome. Thanks, Paul. Over to you, Jim.

JIM KENNEDY: Thanks, Judson. I'm excited to show you some of the AI capabilities that our smart technologies and advanced analytics team or STAAT have developed working hand in hand with our advisors. So imagine I'm a financial advisor. I start my day, like most of us do,

heads down checking my email. Statuses appears as a side panel within Outlook to help me complete client requests efficiently and to uncover new opportunities right in the flow of my work.

I've received an email from my client, Aaron. This is a common request: He needs to transfer some funds. Traditionally, I would have to manually search for information relating to Aaron, but STAAT Assist automatically classifies the email, extracts key information, retrieves data like account details, withdrawal limits, available cash and transfer guidelines without me doing anything. The agent then auto populates the transfer request form for me, meaning that all I have to do is review and then approve with a single click.

Now, it looks like an email has come in from my client, Sam, looking for advice on the current market. STAAT Assist detects that this email contains a client inquiry and automatically offers me relevant advice. So I can ask STAAT Assist questions about investment research, the potential impact of import taxes, or insights on certain funds to help me prepare for this client engagement. STAAT efficiently locates and summarizes the research and insights, equipping me with real-time information right at my fingertips.

Now that I'm prepared with the right information for my meeting, I can include a link to the STAAT meeting scheduling agent in my communication back to the client, helping streamline the back and forth of scheduling logistics and making it easier for my busy clients to find a convenient time to meet. AI at UBS is helping financial advisors and their teams spend less time on administrative tasks and more time doing what they do best, serving and advising their clients to help them succeed. Back to you, Judson.

JUDSON ALTHOFF: AI can help democratize intelligence, bringing information to the edge and making it accessible to everyone, everywhere. We're working with Land O'Lakes, an agricultural cooperative to put data-driven intelligence with AI into the hands of agronomists like Tyler, to help farmers grow their businesses while feeding their communities.

LAND O'LAKES: An agronomist helps a farmer make decisions so they can get the most return on their investment. This time of year, we're out harvesting. The crop protection guide is an 850-page guide, I have to know where all that information is and be able to pull it out. With AI, I can add another layer of insights to make quicker decisions so that we get higher ends of the yields to get the best results for farmers.

JUDSON ALTHOFF: AI can help us accelerate innovation, empowering builders with the tools to tackle big challenges through agile experimentation. We've partnered with industrial automation leader Siemens to infuse AI into industrial automation processes right on the factory floor for engineers like Svenja to build what's next.

SIEMENS: Innovation isn't just an idea on a slide, here it comes to life. You can see it, you can touch it, and you'll be part of it. Coding is the foundation for everything, but incredibly time consuming, and I would rather spend the time creating new things. For me, it's about letting AI do the tedious, repetitive work so I can focus on what really matters. My title might be engineer, but I am a maker at heart.

JUDSON ALTHOFF: AI can help us unlock creativity, removing barriers to innovative thinking and enabling makers to spend more time bringing their bold ideas to life, and we're doing just that with Nestlé. One of the world's most beloved brands and companies in the food industry, they are weaving AI into the creative process and inspire marketers like Felipe to imagine and design what's next.

NESTLE: My passion is finding new ways to connect with people. We're creating, adapting and localizing digital twins for virtual versions of our thousands of products, from KitKat to Nespresso. This opens up new opportunities to create personalized and scalable content. I'm proud to deliver high-quality, meaningful moments, delighting our consumers. This is what gets me most excited.

JUDSON ALTHOFF: AI can help us enable opportunity, applying tools and technologies to foster more equitable experiences. We're working alongside the State Secretary of Education in Sao Paulo, the governing body for education across all of Brazil, to build AI-powered solutions that help students learn with confidence and help educators like Katia rediscover the joy of teaching.

SAO PAULO SEC. OF ED: (In Foreign Language.)

JUDSON ALTHOFF: AI can help us obsolesce the mundane, reducing tedious and time-consuming tasks that get in the way of meaningful work. We've teamed up with telecommunications provider Telstra to deliver intelligent AI solutions for retail associates like James, so they can spend more time doing what they love: Helping people.

TELSTRA: I've always been a people person. I love welcoming our customers into the store. We have so many different products and services, and our customers trust us to be the experts. Having a tool like Ask Telstra provides answers at my fingertips and ensures I can give our customers the quality of care they expect. I can just ask a question and get the right answer in real time. If I can even be a small part of making someone's day better or easier, that's a win.

JUDSON ALTHOFF: AI can help us fundamentally change the future, pushing the boundaries of what's possible to drive societal progress at scale. We're collaborating with Insilico Medicine, a clinical biotech company, to apply the most advanced AI models to accelerate drug discovery and help scientists like Dr. Heng Zhau discover breakthroughs and find cures.

INSILICO MEDICINE: Developing new drugs can take decades. You never know when you might find that one breakthrough that could change someone's life. Our team uses AI to identify potential biological targets faster and then run simulations in our platform to create new molecules in a matter of days. We've already developed 30 drug candidates using AI, and we're just getting started. I'm more optimistic than I've ever been.

JUDSON ALTHOFF: AI and the flow of human ambition, ubiquitous innovation, observability at every layer of the stack, all on the fundamental building blocks of intelligence and trust, empowered by Work IQ, Foundry IQ and Fabric IQ because it's about your IQ, your bold

ambitions, the human ambition that lives inside of your company, all brought together so that your organization can achieve its highest aspirations and growth potential.

At Microsoft, our mission is to empower every person and every organization on the planet to achieve more, and that is what we're going to do by empowering all of you to become Frontier Firms. Thank you and have an amazing Ignite.

END