11172026 Ignite Keynote Judson Althoff

Microsoft Ignite
Keynote
Judson Althoff – CEO, Microsoft Commercial Business
Moscone Center, San Francisco, California
Monday, November 16, 2026

JUDSON ALTHOFF: Thank you. Good morning, everyone. Good morning, and welcome to Ignite. It's awesome to be with you all here in the arena and everybody watching online. The rest of our keynote this morning is going to be anchored around these three durable assets of frontier transformation and the journey that lies ahead for all of us: human ambition, intelligence, and trust. And we're going to start with human ambition.

Microsoft Copilot is a platform for human ambition. It enables amazing, everyday AI experiences for all of your people, and on top of that, high-valued experiences using the best state of the art models. Simply stated, anything you can do in Claude or ChatGPT, you can do better with Copilot and for less.

Now, my ambition since the beginning of my career has always been about helping people and organizations solve problems using technology. I thought a great way to start off this keynote would be to show you how I use Copilot to support my ambition. To do that, let me take you back a couple of Sundays ago, very specifically, Sunday, November 8th. I've been super busy preparing for this event, meeting with customers. And let's just say I've been a little bit behind on my email and the Teams messages I've been getting outside of the context of Ignite. I'm here in Microsoft today, which is an awesome new feature inside of Copilot. It appears both in the Outlook canvas as well as in the Copilot app.

And the cool thing about today is that it scans my entire email inbox, all of my Teams messages, and shows me what's most important. Now, you'll see a lot of things redacted and grayed out here. That's because this is my real environment and not a demo, and so we have to hide a few things. But here you see that today has surfaced up this email from my Chief of Staff, Caroline Sanderson. You'll see this email is calling to attention that in January I have lead keynote at NRF, which is Retail's biggest show, and I really haven't done anything to prepare for it yet. I'm going to go into Copilot and actually get started on some preparation without bothering my team because they're super busy working on this event.

I went into Copilot, and I created a project. Projects are super cool because they allow you to chat with your business and use Cowork for advanced tasks. I can use all the models that I need and only the models I need to get real work done. You'll see that I kick things off with a simple prompt. If you look here, I've just said basically tell me what's going on at NRF and what the hot topics are going to be. If you scroll down, you can see that Copilot has done some really cool things here.

First, it has scanned all of my work IQ, all of the external information, and it's used hybrid intelligence here on my Surface Laptop Ultra, which means it's pulled in content that's local to my device. My speaker notes that I downloaded before finding out that I was speaking at the event. And so what it's done here is it's pulled together a great synopsis. It says there are seven big topics ranging from agentic commerce through to the return of the physical store experience and also supply chain and sustainability.

I chat back with Copilot and further prompt: Hey, Copilot. Given that Ignite is just two months prior to this event, wouldn't it be great if we took the seven stanzas of NRF and merged them with the three here at Ignite: human ambition, intelligence, and trust? Copilot goes off and creates an awesome outline for me. The cool part about this is Office is embedded in Copilot. This is not a sidecar. This is the full-blown Office. You see all of the sensitivity labeling carrying through my ribbons and even the ability to track changes. And you can see that it's got some news around Dynamics 365 supply chain.

That's really cool. I can close that out and then actually open up a richer body of work and use Copilot Cowork to take this a bit of a step further. You can see the prompt I've asked here, I've added here is a bit more robust: Tell me more about what's happening, why people love retailer, retail experiences, what the retail craft is all about. And build me a storyboard using Word, a presentation I can share with my team once we get done with Ignite so we can get started on the keynote prep as well as a tracking document that shows me the top customers attending NRF and what they're doing with Microsoft, because we want to stay situationally fluent as we deliver the content.

Here's what it built for me. It multitasked. It used the models that it felt were best for each task producing great work in the most efficient way. You can see the PowerPoint presentation that it built for me. Again, this is PowerPoint. I can go in and edit this document, make it to my liking. You can see the rich storyboard that it built and that it picked up on my notion to really pull out the human ambition here: why people love retailer. What is the merchant's craft in this whole retail experience? And it's giving me a minute-by-minute outline of how to run the keynote that I can share with my team.

Now, that context for them is so powerful because normally we would have to spend meeting after meeting so that we could get those ideas out of my head and into the context of a show, so this is super beneficial. Lastly, it's pulled together this customer tracker in Excel. Again, there's things redacted here because this is from Work IQ and Fabric IQ pulling together what we do with these customers and how important they are to us as partners, whether they're speaking at NRF.

Fantastic. There's one last thing I've asked Copilot Cowork to go and do, and that is to take this awesome storyboard and turn it into a 3D model using code. And I'll come back to that and show you what it did. But meanwhile, I want to breathe life into this whole thing. I want to keep the project going in an ongoing way. To do that, I bring in my Autopilot. Now, I've named my Autopilot "Amelia" because I'm a little bit of a flying junkie, and Amelia is going to go do some really cool things for me. She's going to do two major bodies of work on an ongoing basis. She's going to keep that spreadsheet I just showed you up to date with everything that's happening in

the market, what's most important to my customers. And she's also going to attend all of my product meetings. Big part of my job is taste making our products and working through product roadmaps. Amelia is going to keep track of everything we're working on and help me understand what's best to present at NRF once we get to January. She's also going to inform the rest of my team about all of this work.

With that simple prompt, Amelia has gone off and created routines for herself. She'll do these things on an ongoing basis. Amelia doesn't just work for me. She can work here in Teams. Here you see the NRF planning channel that I've created. She's told the team the three things she's going to go do. And what's really cool is in Copilot, I can set up a workspace that maps to this NRF planning channel and Teams. And the really cool thing here is I've got Autopilot and my people working together. My Autopilot, Amelia, my teammates' Autopilot. And they can work on curating this content. Again, this is all on a Sunday. Not bothering anybody. Me working on my own to help prepare the team once we get done with Ignite.

If we step into code, I'm going to take that markdown file that Copilot Cowork built for me. And you can see here I've just given it a simple prompt: Build me a 3D demo based on the markdown file. I can actually give it some cues about the colors and the fonts and the background and tastes make it to my delight. And you can see the demo that it's built here. Really, pretty cool. Here's a story about how human ambition comes to life in the retail experience. Here's the intelligence behind all of it, the data sources, the things that belong to the retailer, their unique IQ, the shopping agents, merchandising agents, logistics agents, and the trust platform that enables the retailer to manage, govern, secure, and get business value out of all of these assets and all three layers coming together.

Now, the really cool thing about this demo is even though I didn't write a line of code, there is real code behind it. I can hand this JavaScript file to my professional developers. They can dump it into GitHub Copilot and plumb the back end with databases and hit the ground running.

Now, again, the awesome thing with all of this is I did all of it without bothering my team on a Sunday. I spent a couple hours doing it. I can track my usage. If you look in here, you can see that on that Sunday of November 8th, I spent about 10,000 credits or about 100 bucks. Now, two hours of my time, $100 worth of tokens, Copilot using only the models it needed to use optimizing the experience, and my team hits the ground running for NRF without me bothering them.

That's how Copilot empowers my ambition to hear more about how Copilot can empower your ambition. Please join me in welcoming to the stage the Executive Vice President of Copilot, Jacob Andreou.

JACOB ANDREOU: Thank you. Thanks, Judson.

Now, what's so cool about what you just saw is that wasn't a special build. This is how we're getting our work done in Copilot right now. And you'll notice what he didn't have to do. He didn't jump into a developer tool to write code. He didn't jump between five different applications. He did everything right in Copilot.

Now, a ton of what you just saw is already in your people's hands, and some of it's coming soon. But let's start with Home. Home is where Chat and Cowork come together in a single composer. This means that whether you have a quick question or you're kicking off an entire project, Home is the perfect starting point. And you also saw Office. We've all seen Anthropic and OpenAI embrace documents, slides, and spreadsheets, but their simplified viewers and editors force compromise on power and collaboration, and real work needs both.

Knowledge work already lives in Word, Excel, and PowerPoint. It's where your team is already collaborating, and that's why we brought the full functionality of Word, Excel, and PowerPoint directly into Copilot. This means that you can collaborate with AI and your team at the exact same time. This includes track changes in Word. It includes real advanced formulas in Excel, and of course, the ability to @ mention a teammate to bring them into the doc that you're working in. This means that your whole team doesn't have to blow up their workflow, doesn't have to switch which apps they're using, and all the files stay at rest in your tenants with security labels and compliance policies fully intact.

The best part is all of Office and Copilot is included in your Copilot subscription. That means that Copilot is the only way to draft a PowerPoint presentation without spinning a meter. You also saw Judson use Code. Code allows anyone to describe what they want to build in their own words. Copilot picks the best approach, whether it's a tracker, an app, or a dashboard, and it builds a full working app. Code is fully sandboxed locally or hosted securely in your tenant by default. It makes hosting and sharing as easy as sharing a word document.

Now, none of this innovation matters unless it really holds up at scale. And few are at the scale of Chevron. And so I'd like to welcome Ryder to show us how the new Copilot is transforming work. At Chevron.

RYDER BOOTH: Thanks, Jacob. At Chevron, powering the world to enable human progress is our purpose. Society depends on us to provide reliable, affordable energy and to do this while protecting people and the environment.

One of the ways we achieve this is through an engineered basis of design. It's the blueprint our engineers use to ensure we execute and operate safely and reliably. However, building this can take months. To improve this, we turn to Copilot, where I typically start and end my day. Let me walk you through the demo.

We started in Copilot Chat and asked it to write a Cowork prompt to build a basis of design. You can see here's the prompt. It's 300 lines of technical instructions to Cowork, and what enabled this is Work IQ in the background, sourcing relevant files, and with one click, I'll send this over to Cowork, and work begins. Cowork's already gone through hundreds of engineering files. It's highlighted to the engineer right here where they need to focus their time to build out the basis of design. And here at the bottom, our first two initial drafts of basis of design document and an Excel workbook. However, these two documents were based on a relatively small data set compared to our complete enterprise data.

Next, scroll up here. We asked Cowork IQ to enrich both deliverables with our engineering standards and with all of our relevant data across the company. And what enables this is Fabric IQ. You can see here in the bottom right corner in our plugins, Fabric IQ. And without changing any of our infrastructure, our data architecture, it's compared our basis of design to thousands and thousands of company-operated designs and operational reports. I'll scroll down, takes just a second here, to our deliverables. You'll see our two updated deliverables here. And I'll click on our basis of design, over 80 pages of technical information on a basis of design based on our company template, informed by our company standards and pulling in from across the company data and our institutional knowledge, located in multiple systems across the company, including a system from a company we recently acquired, all enabled by Fabric IQ. I'll close out of this and get back into our session.

And now that we've completed that task, we asked Cowork to turn all of this into a reusable skill, and I'll go ahead and just scroll down to the output. You can see it working here. And here it is. Bottom of the screen. Assets. Basis of design. Skill. Now, anyone on the engineering team can use this skill. In fact, any engineer in the company can use this skill to create a basis of design. We're almost finished. We go back into Cowork. We give it a prompt: To summarize the work we've done, post the two deliverable files into our engineering Teams channel for further review. And with one click, we're done, and I never left Copilot.

JUDSON ALTHOFF: That was really awesome, Ryder. I bet the question that everybody's asking in the room is: How much did all of that work cost you?

RYDER BOOTH: Well, Judson, costs always matter, but in Cowork, I can hit "/cost." Let's see. Oh, in this session, 4,500 credits, around $45, which is significantly less than if I'd stood up a coding team to build this as an agentic solution. And our engineers' time went from months to minutes.

JUDSON ALTHOFF: Fantastic. Thanks so much, Ryder.

You saw how Ryder visualized his cost with just one simple command, but it's really cool to be able to see this from the entire perspective of the IT organization. I'm here with Steph Waters from Chevron's enterprise data and AI team. Steph, show us what this looks like across all of Chevron.

STEPHANIE WATERS: Will do. Agent 365 gives IT visibility into the AI spend across the organization, so we can see where investment is going and manage it at scale. At a glance, I can see what we've used, what remains, and what's driving spend across our services, agents, policies, and models. And I can drill down into our users. And here I can see Ryder's, and I can see the models he's using, as well as the services in the use cases and the value that he's getting for using Copilot.

Now, importantly, IT needs to set the spend limit before anyone can use Cowork. And once that limit is reached, a user must request an increase to continue to use Copilot that prevents runaway stems and alerts, let IT step in before any issues occur. I can also set which models are available for use as well as create ones from templates. Ultimately, I'm not just managing cost. I can see

the value that we're creating and make better decisions about where to optimize and where to invest next.

JUDSON ALTHOFF: Awesome job, Steph. Thank you. Jacob, back to you.

JACOB ANDREOU: Those guardrails help keep costs in control, but Copilot works to maximize how far each of those dollars go. We do that in three ways. First, we do it with the Copilot subscription. The Copilot subscription is the best way to give everyday AI to every person at your company. It includes Chat, all of Copilot and Office, which means that you can do drafting, summarizing, and analysis all without spinning a meter or generating additional charges.

Everything you saw is also powered by auto-routing. Auto-routing selects the right mode, the right model, and the right effort, all from a single composer. That means that you get the best possible response. And if you're using Cowork, we help manage the costs of each request under the hood.

It also benefits from Work IQ. Work IQ grounds everything that you do in your company context and how you work. The most impactful way to keep costs under control is to make sure that you have the context to do the job right the first time. Copilot knows your work. It's able to pick the best model for every task, and all of that is included in a high-value, fixed subscription. Whether it's everyday AI in everyone's hands or advanced AI for your very hardest tasks, Copilot saves you money. This is AI built for work.

So everything you've seen so far is available now. And I want to take a second to pivot to some of the really cool things rolling out over the next couple of weeks. This is when Copilot goes from helping you with your work to actually moving your work forward.

Judson started this off by showing you Today. Today is our new proactive, personalized experience in Copilot. It's a command center that spans across mail, calendar, Teams, meetings, tasks. It shows you what you missed, what needs your attention. And Today doesn't just surface the work; it actually helps move things forward. The draft is already written. The proposed schedule change is already made. Everything is waiting always for your approval.

You also saw Autopilot. Autopilot is your digital teammate. It's the first-ever enterprise-grade, long-running agent. It lives in your tenant. It has its own identity, its own memory, its own email, and its own computer. And that means that Autopilot is not just limited to M365 apps; it can also use all of the legacy applications that your business runs on. Autopilot watches channels, it can follow up on threads, and it can push recurring work forward even while you sleep.

We're also really excited to announce Workspaces, which you saw in Satya's keynote at the top. The future of collaboration is bringing people and AI into one place, and Autopilots come to life in Workspaces. Until now, AI has been mostly single player, and Workspaces makes it truly multiplayer. In a Workspace, people and agents come together around a common goal with shared context.

The Workspace automatically pulls in the files and the people needed to get something done so that conversations and decisions all live in one place. Everyone has full context, they know where things stand, and they know what to do next.

We think all of these things can make Copilot an amazing way to serve every person at your business, but what sets your company apart likely lives in the functions. It lives in legal, in sales, in finance, in HR. Each of these functions in your company has its own systems, its own data, its own way of doing things.

Copilot is a platform for human ambition, and plugins is how it supports each function in your business. It brings together the skills, the connectors, and the tools that each of your teams relies on. Plugins is how Copilot stops being a generalist and actually becomes an expert in the job to be done.

It approves plugins once and it works everywhere Copilot does, across Chat, Cowork, Code, and even Autopilot. Please welcome Andrew from BlackRock to show you how they're putting the new Copilot to work. Andrew.

ANDREW GIDNEY: Thank you. Thank you, Jacob. Sales is a relationship business. It's about knowing your clients, understanding what matters to them, and bringing the best solution to them at the right time. Today, the challenge isn't having the information to do that. It's connecting the information across conversations and systems, and that's where AI gets really powerful. It brings those signals together and gives our relationship managers a much better starting point, so they spend more time with clients and less time searching for information.

At BlackRock, the opportunity isn't simply adding AI to the way sales works today; it's reimagining how we manage client relationships. Now, we're working with Microsoft to shape the Copilot and Dynamics experience around the needs of our teams and our clients to enable our sellers to mine insights, take action, and accelerate commercial outcomes. So let me pass it to Nirja to show you what we're working towards.

NIRJA TRIVEDI: Thanks, Andrew. Hi everyone. I'm Nirja Trivedi from Microsoft. Let me show you what we're building with BlackRock. So let's say I'm a global sales lead. My day starts in Copilot, specifically in Today, the command center for my work. Here, information from Dynamics 365 sales is pulled in along with information across my work systems as well as my business systems like financial health and pipeline performance.

So here, for example, I see from Dynamics 365 sales that my team is at 73% of their goal. In order to hit our goal, we need about 3x pipeline coverage, but we're only at about 2.2x, so we have some work to do. Copilot spots this gap and recommends a cross-sell play around our new private markets fund. So I'm going to dive into this a bit deeper with Cowork.

Cowork will pull in information across Dynamics 365 sales and services plugins, like my service history and client information to identify the top clients that are most likely to commit. These clients represent about $110 million of potential commitments. More than that, the top clients are

in decisions and could commit by December. So I want to turn this insight into action, so I'm going to jump back over to Cowork.

I'm going to ask Cowork to help me build a package for each client with the standard information from BlackRock's knowledge base, a personalized email, and an invitation to meet with our private fund specialists in New York. I also want to account for some of the travel time in between. Lastly, I want to send each package to the client's relationship manager to review.

Cowork will pull this information from Dynamics 365 to route it to the specific account – to a specific client relationship manager. More than that, nothing goes out without that relationship manager's explicit approval. And if I want to, I can not only read from the CRM, I can write back with my approval as well. To see this in action, I'm going to hand it over to Efe.

EFE ABUGO: Thank you, Nirja. My name is Efe Abugo, and I am a relationship manager on Nirja's team. In my Today, Copilot has pulled Nirja's request right to the top of my view with those three personalized emails to send out to my clients, and below I can also see other recommended actions I could take today.

So I'm going to hit start task, and Cowork is going to begin building out those three emails for me to send, and Work IQ pulls in meeting and email contacts from Microsoft 365. And with Dynamics 365 client activity and history, I'm able to actually see personalized emails based upon conversations we've actually had.

Now, two of my clients are actually going to be in New York next week, where our private fund specialist will be, so Cowork has proposed meeting times based upon his actual availability. Now, all three emails look good to me, so I'm going to hit send on all of them, and they're all going to go off all at one time.

Now, this is a lot of work to stay on top of, so to help me keep track of everything, I am going to put Autopilot to work. Autopilot will work on my behalf, keeping me updated on key items, as well as helping me move critical work projects forward. So here I'm going to ask Autopilot to help me track my fundraising outreach opportunities, client news, and service cases. And if a client hasn't responded within a week, I wanted to draft a follow up for me to review and approve. And additionally, I want to get a weekly view of relative fund performance and firm level news.

So here we see Autopilot has created the two routines. First, we have that daily check where Autopilot will look across Dynamics, Outlook, and Teams, and we have that Monday morning intel brief. And the guardrail is clear: Nothing goes out to my clients without my approval. Autopilot is operating within the supported controls IT has already put into place.

And if we come out here and navigate to the right, I can see all of the routines I currently have running, including the two that we just created together. So at any point I have visibility into the actions Autopilot is performing on my behalf. Now, let's turn it back to Nirja to see how things are going.

NIRJA TRIVEDI: This is great. Thanks, Efe. All right. So it's been about two weeks, so let's check in on the progress of my team. I can see in just two weeks, we're up from 73% all the way up to 81%. We also have $59 million of new pipeline generated from this play. So you can see here that Copilot really becomes the front door to my business processes. Client knowledge that once only lived in seller's heads is now intelligence my whole team can act on. And with Dynamics 365 plugins, sales and service intelligence is now directly in the flow of my work so I can keep the momentum to have a strong close to the quarter. Back to you, Judson.

JUDSON ALTHOFF: Thanks, Nirja and Efe. Really appreciate it. Well done. So Microsoft Copilot, again, is a platform for human ambition, and plugins enable you to wire in these experiences for your people and how they work, whether they're sales professionals or other walks of life.

Now, you would have expected us to have plumbed that demo with our own assets to see Copilot working well with Dynamics, but what you might not have expected is that we're doing this same thing with Salesforce. We're announcing three major new product innovations in conjunction with Salesforce today. First, Copilotforce, which is the new Salesforce plugin for Copilot, the Salesforce Cloud Sales Cloud in Teams, and the integration of Data 360 and Fabric with Zero Copy. To share more about that, please welcome to the stage President of engineering, Rohan Kumar. Rohan, welcome back to the Ignite Stage.

ROHAN KUMAR: Well, thank you, Judson, and good morning, everyone. After nearly 28 amazing years at Microsoft, it's wonderful to be back at Microsoft Ignite, now representing Salesforce as we bring together our power to shape the future of enterprise AI. It's clear that the foundations of computing are changing and people actually had to adapt to software, navigating applications, moving through systems, and actually piecing together their work.

And what we're seeing essentially is the ability for AI to become the interface of the future, applications becoming capabilities, and enterprise users basically spending their time focused on their business outcomes that they're trying to seek instead of navigating software. This is exactly the vision we landed at Dreamforce earlier this year when we introduced AIforce, bringing the power of Salesforce closer to wherever you choose to work.

And for millions of business users across the world, that place is Microsoft 365. Salesforce has your complete customer context, business understanding, and workflows. Microsoft 365 is where your people collaborate and get their work done. Why should they not work together? And as Judson mentioned, we are so excited to introduce Copilotforce.

We are bringing the best of Salesforce directly into Microsoft 365 Copilot and Microsoft Teams. With this ability, it's not just about your CRM data. It's about your complete trusted customer context, your business understanding, and the ability to take action, all with Salesforce security and permissions. So instead of asking AI, what's the next best action with the customers, you can actually get it done, get it executed using Copilotforce.

This is the power of the agentic enterprise: intelligent, open, and interoperable. And we are thrilled to partner with Accenture, both Salesforce and Microsoft as we build Copilotforce to

bring this part together, and we are thrilled that they're deploying this across their sales teams. Back to you, Judson.

(Applause.)

JUDSON ALTHOFF: Thanks, Rohan. So we're – that's right. You can clap for that. So we're really excited to have Accenture share what this is like in their environment. Please welcome Tony Leraris, the chief information officer of Accenture.

TONY LERARIS: Thank you Judson. At Accenture, we all work in Teams. It's where more than 800,000 of our people meet, collaborate, and increasingly use Copilot to deliver impact every day. Now, with Copilotforce, our sellers will have customer contacts built right into the tools they already work in. Accenture is a design partner, and we're customer zero. We help build it, and we will prove it on ourselves first. Let me show you what that looks like.

Here's Priya, one of our sellers. She relies on Cowork to track her sales priorities for the day. Copilot reasons over her entire Salesforce pipeline, from open opportunities to the latest customer activity on each account. And within moments, briefs her on the top deals and where to focus.

First up, a new RFP from (Coastal Foods?) that needs her attention. With a quick tap, she uses Cowork to set up the deal room to tackle the whole pursuit. The opportunity in Salesforce, a Teams channel, a closed plan, and a weekly deal sync and everybody's calendar. Day one, and we're already organized. Now, let's jump to the deal room in Teams.

Here, the team is grounded on the customer opportunity with the RFP and the closed plan ready to go. Let's hop over to the Teams chat, where Priya has shared with the team that they've made it to the short list. Without needing to switch apps and directly in the chat, she asked Atlas, her Autopilot, to update Salesforce. Using the Salesforce plugin and grounded on Work IQ, it can pull insights across CRM, emails, and chats and meetings, and proposes five updates to send to Salesforce. One click, and Salesforce is current.

Now, this isn't just for sellers. A week later, Priya's manager opens Copilot to see how her team's updates have moved to forecast. Using the pipeline review skill, which is one of more than 38 Salesforce skills available on Copilot, it combines CRM data with context from email and chat to give an instant snapshot of a region. You can see it flags the two biggest deals at risk with a recommended next step for each.

With a simple confirmation, she sends the account manager a nudge in Teams tailored to the opportunity in Salesforce. That's sales leadership and the flow of work. Copilot, Teams, and Salesforce working as one. With Copilotforce, our sellers will spend less time on process, more time with clients, and by proving the value on ourselves first, we can help every company do the same. Judson, back to you.

(Cheers, applause.)

JUDSON ALTHOFF: Thanks, Tony. Intelligence is the second durable asset that we all need to build on the frontier transformation journey. Not artificial, yours. We all need to build, retain, and continuously improve the assets inside of our organization so that intelligence compounds from within. Last year, we launched Work IQ, Foundry IQ, and Fabric IQ to release to you the fundamental building blocks for this intelligence platform.

Now, over the last year, we've been thrilled to see so many customers adopt this same concept. You'll see a Lilly IQ, a Toyota IQ, and a Levi IQ. We really want to bring this whole concept to life for you by showing you Unilever IQ. Please join me in welcoming to the stage, Deepalok Roy Chowdhury from Unilever to talk to you about Unilever IQ.

(Applause.)

DEEPALOK ROY CHOWDHURY: Thank you. Thank you, Judson. Every day we have 3.7 billion consumers use a Unilever product, and that helps us generate 20 million orders and 5 million claims a year. At that scale, we need intelligence at the core, and that's Unilever IQ. How our people operate, how our business runs, and how we generate value for our customers. Sharoon, you want to showcase the trade claims agent we have been building.

SHAROON SRVIVASTAVA: Absolutely. Thank you, Deep. Here is the trade claims agent. It works where our team works. You can see it has already processed 14 claims on its own, and now it's asking me for help for the Northwind claim. Now, let's take a closer look. Looking at the first line, you can see that the customer is billing us. However, we notice a discrepancy with the contracted rate. More than $2,000 over.

DEEPALOK ROY CHOWDHURY: $2,000 on one line? I don't think a human would have picked up out of the millions of lines that we look at. And so, I am sure the agent did some magic to get this going spontaneously. How did it do that, Sharoon?

SHAROON SRVIVASTAVA: Great question. The agent in this case ran a three-way match. You can see that it looked at the deal sheet from Foundry IQ and matched it with the claim file and point-of-sales data from Fabric IQ. Now, if I'm working on this claim, I'm trying to understand, well, has this happened before? So now you can see that the agent goes through all my past claims and finds four similar claims with the exact same pattern.

But even more interesting, it looked at Work IQ to find the people behind those decisions. And that is Unilever IQ in action. Our contracts, our data, and our people all in one place. So, Deep, should Unilever pay this high rate?

DEEPALOK ROY CHOWDHURY: Hey, hold on. I don't want an agent to make a judgment. Let's bring in some people in and have the agent and people work collaboratively. So let's have Daniela call up Jordan, who is the account lead for the Northwind account, and see what she has to talk about it.

SHAROON SRVIVASTAVA: Exactly. So it seems like Jordan has already responded here, and he's recommending to settle this as a one-time adjustment. Now, I can apply the override here,

but here is the best part. Jordan's judgment doesn't just stay in a chat thread, it's saved back to Unilever IQ so that the next time this comes up, the agent knows how to handle it and also helps with auditing.

Now, finally, for the four items that are still pending, the agent drafts a letter leveraging all of this context from Unilever IQ and clearly articulates whatever is needed to close this claim out. That's nearly $12,000 protected in value on this claim alone.

DEEPALOK ROY CHOWDHURY: And $12,000 on one claim. Think about the millions we save on the other claims if we find such discrepancies. Great job.

SHAROON SRVIVASTAVA: That's right. Now all of this comes back to Unilever IQ. Every decision our team makes is added here. Our best analyst, tacit knowledge, is now part of it, and every claim makes it smarter. Back to you, Judson.

JUDSON ALTHOFF: Thank you. Well done.

(Applause.)

JUDSON ALTHOFF: We're taking this a step further today by launching Microsoft IQ. It's your enterprise context management system that enables you to create these durable business plugins that you can use with Copilot and across your agentic estate. Now, of course, it works with GitHub, Foundry, Fabric, and Azure, but it is also open, model diverse, and heterogeneous at every layer of the stack.

Here to tell you more about it, please welcome to the stage EVP of Azure, Arun Ulag.

(Applause.)

ARUN ULAGARATCHAGAN: Thank you. All right, thank you, Judson.

What you deliver both, and what every company here can build are agents that understand exactly how your business works. Today, we are excited to announce Microsoft IQ, the enterprise context management platform that allows you to bring together all of the context across your entire organization to a single place, so that it can be used consistently across all your agents. This includes all your collaboration data, both humans-to-humans and humans-to-agents. It includes all your business data, your unstructured data, data from the Web, as well as context from pretty much any enterprise system that you already use.

It's really easy to get started because you already have the building blocks. For every Microsoft IQ customer, Work IQ makes your collaboration data available. This includes your Teams chats, your email, your documents, your spreadsheets, your SharePoint sites, and so much more. Work IQ enforces the permissions in Microsoft 365, so every agent can only see what it's supposed to see.

For every Power BI customer, Fabric IQ makes all of your metrics and KPIs that you've already defined in your Power BI semantic models available to your agents, again, with all of the security and permissions intact. Power BI already connects to all your data, not just from Azure, but from AWS, GCP, Databricks, Snowflake, and so much more. It's ready to go without having to rebuild your semantics in yet another platform.

With Foundry IQ, you can build knowledge bases that bring together everything your organization knows, from contracts and precedents to technical manuals, case studies, and beyond, making that knowledge accessible to agents. And when you need to selectively bring in content from the internet, for example, regulatory guidelines or the news and specific topics, Web IQ can take care of that for you.

Microsoft IQ goes beyond the Microsoft ecosystem. With 350+ connectors to third parties, including ERP and CRM systems, you can bring your entire business context into one unified platform. And it works wherever work happens in Copilot for your people, in Foundry for your agents, and because it's completely open, in assistants like ChatGPT and Claude as well. This context is your intellectual property. It's what makes your organization unique; it is your IQ.

Our goal at Microsoft is to ensure that your IQ stays yours, and it compounds in a way where the benefits accrue to you, not to Microsoft, not to the model companies, and definitely not to your competitors.

Now, I'm going to invite Eren Orbey to show you a demo, so you can see Microsoft IQ in action, unifying context for the business process, packaging up the context into a business plugin and using it seamlessly in Copilot.

Eren, over to you.

EREN ORBEY: Thanks, Arun.

You just saw what a trade claims agent can do with the right context, and Microsoft IQ is the platform that helps enterprises manage that context.

Our company's knowledge used to be scattered across different sites, handbooks, systems, and imaginations, and now it's consolidated and codified here in a brand new interface as our IQ. These are business processes that we've minted with building blocks from Microsoft plus context from our own systems.

Let's take a look at that trade claims process. It sits under legal and compliance standards that apply company-wide. The Legal team has vetted certain skills and policy documents that point agents to canonical sources like these. And then within Legal, we have guidelines for contracts, and then specifically for commercial contracts.

And down here in trade claims, we automatically inherit from that hierarchy. everything above flows through skills, documents, data, and official Web sources. And then our claims team can add bespoke agent skills for checking a charge or escalating exceptions to the right approver.

These skills sync up with the responsibilities that we've defined for people in our organization, and we've pulled in the live Teams thread where we triage escalations day by day. Foundry IQ comes in to help us index files like contracts and deal sheets for agentic retrieval. And Fabric IQ brings in operational and analytical data with OneLake from across clouds, in open formats, plus semantic models, like this one, which define metrics and measures atop that data.

Every source here can be customized with instructions telling agents when and how to use it, so they load only what they need and no more, saving tokens along the way. And as a subject matter expert, I can easily evolve this process by bringing in assets from our enterprise and beyond. Of course, there's a Copilot experience to help me evolve the process further, and all of this context can be permissioned so that only select stakeholders on our team can view and edit it.

At the top level, this context can also be certified by IT as a source of ground truth for agents, as we see here, and put it all together. And our trade claims context can be published as a durable, interoperable business plugin. It's governed in Agent 365 and accessible in Copilot, in Foundry, and in other applications that we use.

Here's what that means at the most essential level for a claims analyst. When I ask Copilot why one retailer's claim looks high, our IQ steers it straight to trusted sources like these. The response that we get charts the charges that don't match up according to our data and our contract. It lays out next steps based on our process, and for anything that needs human judgment, Copilot drafts an email to the approver specified by, you guessed it, our context.

With Microsoft IQ, every agent and every analyst can follow the same process that we defined, getting answers grounded in in-house expertise. That's the power of curated enterprise context.

Happy Ignite, folks. And Jay, over to you.

JAY PARIKH: Thanks, Eren.

(Applause.)

Now, if you ask any CIO about agents, you're going to hear the same thing. Building them is easy, but running them in production is not. A production agent needs things like identity authentication, sandboxing, access to enterprise data, monitoring, red teaming, and the list goes on. And that's where most agents stall. The question isn't how to build an agent; it's how to build and run thousands of them, trust them, and make them better every day.

We built a foundation to help you do that, a comprehensive end-to-end platform. GitHub is where you build, Fabric is where your data is organized, and Foundry is where your agents run in production and keep getting better.

Today, I'm going to cover a couple of key features in Foundry. Let's start with models and model choice.

Foundry has more than 11,000 models, including OpenAI, Anthropic, Meta, Mistral, open weight models, and our own MAI models, but the number isn't the important point. It's having one place to evaluate them and switch when a better one chips without having to rebuild your agent.

Now, with our model router, you don't have to pick just one. Simple requests go to small, fast models. Complex ones go to the frontier. Our data shows that this is up to 65% cheaper than sending everything to a frontier model. This allows you to strike the right balance between the quality of the output and the cost of getting to that output.

The next big question is, once I have my agents up and running, how do I improve them over time? After all, your business isn't static. You're working hard to improve it all the time, and your agents should be doing the same.

Foundry helps you do that in three ways. First, agents on Foundry get better the more you use them. Foundry allows you to feed agent traces and user input back into the system, creating a continuous learning loop.

Second, Foundry looks across every run for ways to make your agents better. The system finds and tests improvements, shows you how each one trades off across quality, speed, and cost, and then you get to decide what shifts.

Finally, Foundry provides a range of other capabilities that allow you to improve progressively, from lightweight skill improvements to agent configurations, all the way to training your own model using your own data. And all of this learned intelligence stays with you and only you.

Alaska Airlines runs this kind of loop, and they are here with us today to show us what they've built. Three things to watch: First, how they define better; second, how far optimization takes them; third, what training adds.

Please welcome Charu Jain from Alaska Airlines and Seth Juarez from Microsoft.

(Applause.)

CHARU JAIN: Thanks, Jay. Airlines are in the business of time. Being reliable and running on schedule is how we earn your trust and deliver a remarkable travel experience. Our approach to innovation starts with the guest. We look for pain points across your journey and work to solve them. We're working with our Seattle neighbors at Microsoft to put powerful tools in the hands of the people who are closest to our operation.

What you're about to see starts with routine maintenance. Every maintenance report gets an ATA code that helps us track the work across our fleet. The more precise and consistent that coding is, the easier it is to see patterns and identify issues we may be able to address proactively, rather than having them show up close to departure.

This is very great work, and I'm excited for you all to see it. Let's go see.

SETH JUAREZ: How are you doing, Charu?

CHARU JAIN: I'm good, Seth. What have you been up to?

SETH JUAREZ: I am so excited. When I heard about this scenario, I generated a whole bunch of synthetic data that matches exactly what you're doing.

Now, you talked about ATA code. ATA stands for Air Transport Association, and they have chapter and verse of every single code that maps to every single. You can see it up there. There's a lot of them, but it's also hierarchical, so it's fairly complex.

Do you want to see the system we built to help out with that?

CHARU JAIN: Let's see it.

SETH JUAREZ: Okay, so this is the maintenance reliability software. And what you can see is that during a general inspection, identified number three break is near its wear limit.

CHARU JAIN: And that's routine maintenance.

SETH JUAREZ: Absolutely, this is routine. It's the system working, but here's the thing. The engineer coded it as a 3200 in the logbook, which is technically correct. It's part of a landing gear, but 3240 turns out to be a little bit more precise.

CHARU JAIN: Yeah. See, these precise codes will help us make better decisions and solve the right problems.

SETH JUAREZ: A hundred percent. And so, since it's text coming in, we thought maybe we could make some agentic software to look at the text and see if it could figure out the ATA code. But again, the important bit is that the engineer is always in charge, so they can look at this and then confirm whether something is correct or not. Isn't that cool?

CHARU JAIN: That's really cool.

SETH JUAREZ: All right, so let me show you how we built this. This is Microsoft Foundry, your place to build, optimize, and train your particular agents inside of one service.

Let me go to the agents pane here, and what you're going to see is the Alaska ATA agent is right there. There you go. And what it is, is it's an agent that we built, but we need to know how it's working. There is something called traces that enables us to see inside of these agents.

And so, I already have a trace already open up, and you can see it right here. Notice that you can see the entire run of this particular agent, including getting your Alaska IQ, which includes these ATA codes plus some other stuff that we're going to talk about later. It's really cool.

CHARU JAIN: Seth, how hard was it for the model to understand airline maintenance codes?

SETH JUAREZ: Yeah, see, this is the thing that's most important. Just because we made an agent doesn't mean that we know it's working. What we actually did is we created a rubric that grades this agent. Foundry did this for us, but we can validate. And let me show you three things that it checks. This is like a rubric, like any school grading rubric.

And notice that the first one is correct ATA code is super important, but you also have as part of your Alaska IQ your internal house codes, which are amazingly, super important as well, the way you do stuff, which is good.

We ran this. I made a (naïve?) agent the first time. Do you want to see the results?

CHARU JAIN: I do.

SETH JUAREZ: Okay. It turns out that I did not – this is not as good as I thought it could be.

CHARU JAIN: You can do better.

SETH JUAREZ: I think I can do better. Okay, there's two things you can do whenever you're making an agent better or more accurate. You can correct the assets of the agent, like your instructions, your tool descriptions, or you can fix the model. Let's do that. Let's do the first thing first. Let's see if we can fix the assets.

Inside of Foundry, we have something called the Foundry Optimizer, and what it does is it looks at your agent, and it iteratively goes through fixing the system prompt and the tool descriptions to give us actually a better one. And notice that it actually changed a lot of stuff. It changed a lot of stuff.

And do you want to see how much better it did?

Okay, so here it is. Notice that it's getting way better accuracy. We were using a frontier model, GPT 6 Astra, and I have to come clean. We did a little analysis of this thing, and buckle your seatbelt.

CHARU JAIN: Oh, hold on, Seth. That is way too –

SETH JUAREZ: Yeah, I know, but here's the thing. Here's my explanation. It turns out that because we're using so much more tokens preloading for the instructions, because we're getting more instructions, the input and output token cost went up a little bit. I think we can do better.

CHARU JAIN: You can do better.

SETH JUAREZ: Let's do it. Okay, the next thing we can do is we can actually fix or train a model to know more. One of the things that I do when I'm looking at this, a hint is if you have a lot of stuff in your instructions, that could be an indication that you could train a model to have those instructions in its weight instead of the instructions.

That's what we did. If you go over here, notice that we actually did a training job on MAI-Code-2-Flash, so that it would actually do a little bit better. And notice, if you look at the monitor, as it goes up, you're going to see a chart that actually sees the things increasing.

And are you ready for what we were able to do with this model?

Ninety-seven percent on exact ATA code match and 93 on your coding room, but here, it gets better. It gets better! Everyone, buckle your seatbelts one more time.

CHARU JAIN: Nice. That's fantastic!

SETH JUAREZ: Yeah, we should clap for this.

(Applause.)

This is MAI Code 2 trained, but more, the more that we do this, the more efficiency we get.

How are you feeling about this?

CHARU JAIN: I'm feeling really good. This is fantastic.

SETH JUAREZ: This is amazing. And so, we actually made it work. It's pretty cool.

CHARU JAIN: Yeah. Seth, what I love about this is that it started with people close to the operation. And they saw a problem they could solve, and then they now have the tools to build a solution themselves. People always remain at the center. AI makes a recommendation. Our maintenance engineers make the decision.

And that's how we're using technology to empower our people and deliver the reliability our guests depend on. And that's the innovation we're unlocking at Alaska Airlines.

SETH JUAREZ: Hey, do you know people at Alaska Airlines? Is there a way to get me from Atmos Gold to Titanium?

CHARU JAIN: I think you're going to have to fly more. How about one of our new long hauls from Seattle to Japan or London?

SETH JUAREZ: Let's go! Let's buckle up one more time.

Over to you, Mustafa.

(Applause.)

MUSTAFA SULEYMAN: Thanks, Seth. Thanks, Charu. Great demo.

Our mission at MAI is to build you the very best models in the world, and to keep every developer and organization operating right at the absolute frontier of capability, but what I hear time and again from every CIO and CFO is that right now, token spend is going through the roof. And at frontier prices, the math just isn't sustainable. Our commitment is to give you the very best price performance out there, so that you can really scale your agents.

And we also know that you want models that feel truly yours, where you get to control the weight. We don't just hill climb on frontier quality alone. We co-optimize quality and price together, and we keep doing it over and over. We ship; we learn; we improve; we repeat.

Take coding, for example. Averaged over four widely used benchmarks that target the everyday enterprise coding tasks that we see in GitHub Copilot, MAI Code 2 gives you 98% the performance of GPT 6.1 SOL at just 40% of the cost. And against Opus 5.5, it's an incredible 5% of that cost.

Same goes for Code Flash. It matches 6 Luna and Haiku 5.5 on these everyday tasks at just 65% and 20% of the price, respectively.

But agents are increasingly needed to make decisions, especially where speed and cost are critical. Decision 1.1 provides model routing integrated right into your applications. And on the AI Decision Index, it beats Jev, Google, and OpenAI.

Now, it's super important to say that we've built these models to prioritize the coding work that our customers do every single day, so they're explicitly not yet tuned for the harder problems like fundamental scientific research and very long horizon agentic work. And we won't chase those more niche capabilities at your expense, because after all, when the same model is built to also push the frontier of science, you pay for that capability on every task, even for a routine bug fix.

Our commitment is to keep pushing that price-performance frontier on what matters most to you. And a great example of that is AT&T's custom version of Code 2 Flash. It outperforms 6.1 SOL at standard network analysis for 5G, while being a whopping 50 times cheaper.

Now, it's the same story in cybersecurity. The last few months have really been, I think, a wake-up call and a warning shot for all of us. We're facing a completely new reality, where these attacks and exploits are cheaper, and easier, and more ubiquitous than ever before.

And that means that agent security, alignment, and containment are going to be the most important things that we need to get right in 2027 and beyond. MAI-cyber-2-flash is our response. It's ahead of GPT-6 Astra on CyberGym end to end at less than 2% of the cost per task. Truly remarkable. That's the power of Microsoft's decades of proprietary data and world-class security expertise, enabling us to bring you a model that not only detects the vulnerabilities, but recommends the right fixes at an incredible price.

We're also moving just as fast in multimodal. We're now the number-two lab on image generation, and MAI Image 2.7 just this morning is the world's number one image editing model

at a fifth of the cost of the competition. Customers like Adobe, Sports Illustrated, and Manus say that it gives them professional quality with super-precise and responsive edits. And our transcription model is now the fastest and most accurate in the world. And, of course, medicine is probably the most important, most complex, and high stakes of any of the AI applications. And so MAI Transcribe has now been deployed to 10 million Copilot Dragon patient encounters a month whilst halving the error rate.

Judson said it best, I think at the top: Your core assets in this new platform shift are your IP and your workflows. The years of commercial data that you've assembled, essentially everything that makes your company what it is, your IQ. And I've heard it again and again from leaders. No one wants to pour their IQ into somebody else's model. And if you build on open weights, can you really trust where they've come from or how they'll behave when you use one of them as the foundation of your most valuable IP and data?

And so we heard you. We started with custom models, as we just saw with Alaska and AT&T, but also Docusign, Rakuten, Woven by Toyota faster and better at the tasks that meet your specific needs. But today, we're taking a big step forward with MAI licensed weights, models under your control. You take a frontier MAI model, customize it using your proprietary data if you want to, and then you run it securely in your own Azure environment. You control who has access to it. Your adaptations and your data stay entirely yours. And we'll work with you through Frontier Company to deliver the custom model and embed it in your own products and agents.

That's MAI, a family of world class models with exceptional price performance. The world's best image editing model. The world's best transcription model. Frontier cyber defense at 2% of the cost and world-class everyday coding models at a price that you can truly scale. And crucially, all of it with your IQ, your agents, your weights, all under your control. This is frontier intelligence that you can trust from the company that you can trust. Judson, over to you.

JUDSON ALTHOFF: Thanks, Mustafa. About the best way to round out our intelligence segment is to take the view from the top. The very top. Please join me in welcoming my dear friend Arthur Sadoun, the CEO of Publicis. Arthur, welcome.

MUSTAFA SULEYMAN: Thank you. Welcome.

JUDSON ALTHOFF: Good to have you here in my office at Ignite. So, Arthur, I'm going to be blunt, since we're friends. You're not supposed to be here. If you sort of read the press three years ago and the dawn of generative AI, the very first companies that were supposed to have died off were media companies. But you're still here. You're actually doing better than you've done before the dawn of generative AI and an all grounded on this notion of having intelligence at the core because of the platform you've built. Tell us more about that.
ARTHUR SADOUN: I mean, yes, I should be dead, and I'm here on the Steph Curry thing. I mean, for a French guy, I was sending pictures to my family. This is a great pleasure. Now, more seriously, yeah. A couple of years ago, every AI guru will tell you advertising is dead. Everyone is going to be off job. I remember Sam Altman from OpenAI saying 90% of the people that work in advertising will disappear in the next two years. Four years later, we move Publicis from

90,000 people to 120,000 people. And we made at the time three strategic bets that were tough bets that we actually made with you.

The first is we invested massively into identity and technology, buying Epsilon and Sapiens to really have the power that we need for our clients. The second, which is something that I think is important for everyone in this room, is that we said it's time to break down the silos because technology and AI doesn't work if you don't break the silos. And we have seen great examples today, and we have this philosophy that was "no silos, no solos, no bozos." And we could get rid of all of our bozos and make an exceptional organization.

And last but not least, it was ten years ago now. We said advertising will be led by AI. No one in our industry thought it was the case, and one day I received a mail from Satya saying I'm interested about what you say. Maybe we should build together, together, something together. And this is why we started Marcel, and we moved from not being very well-placed to the number-one position in our industry today.

JUDSON ALTHOFF: Tell us more about how the work actually changed. You talked about the evolution of your organization. What's different about the ambition of people inside of Publicis today?

ARTHUR SADOUN: I mean, I think what we should not forget, and I'm sure it's the case of a lot of people here, is we are a service business. There is a lot of investors that came to me in the last year saying, Arthur, you have to move to really become an AI company. So we are all about the people. We are a service business. Yes, we have the best tech, we have the best data, we have great AI partners, but we are a service business.

And when I heard everything, I heard this morning about human ambition and how do we make sure that we bring people to a better future? I think this is exactly where we are. And I don't know if you're seeing the demo here. Marcel is our AI platform where we bring Publicis IQ with Microsoft IQ. And this gives us a collective intelligence that is the best of our industry, and in all of our people, not only to learn more and to share more, but actually to do more. And that's what makes a big difference for us.

JUDSON ALTHOFF: Awesome. You and I spent a lot of time with customers. Tell us about the partnership, what it means to you and what customers are saying about our work together.

ARTHUR SADOUN: The thing we can do with AI and data today that we couldn't do yesterday are just amazing. We are able to know the customer of our client better. We can find billions of prospects. We can connect all the different media, particularly creators, with influencers and commerce, and then drag it onto the LLMs. We are able to deliver the right kind of content, and we have shown great example here, intelligent content that we can measure. And we are starting to co-create data assets that make the agent of our clients unique.

Now, the things that we are confronted with every day is that it's not enough to have the capabilities, the level of change management you need to bring into the organization is pretty amazing. And again, I don't think anyone in this room is going to be replaced one day by an

agent or a robot, but I do think we can actually all get fired if we don't bring the right level of AI adoption into our organization. And basically, where I think we did a great job together is that you guys are really the trusted partner that get us there. And if I want to be a bit provocative, I will say that in terms of subject matters, we have done a great job with us. You have great engineers, but where you have made a big difference is in the change management, is in this ability to make sure that we continuously improve.

And I have to admit, in an industry that is a bit creative, you have been the adult in the room, and when you look at what we do with Rodrigo now through the frontier, it's exactly what we are trying to replicate for our clients.

JUDSON ALTHOFF: That's super. Well, look, I can't thank you enough for the partnership, and everything we have done together is deeply meaningful for us and all you're doing to help Microsoft.

ARTHUR SADOUN: And it is. It is really for me. And maybe I close with that. And this is why it was so important to me. I didn't know it would be in this place, which I like a lot. But I've been CEO for ten years now. It has been a tough journey. I had to fill big shoes when I took my job. I felt the distress of the financial market after it went down like this before going up, I went through Covid, I went through cancer. We went through the war in Europe. We went through the distress on AI. But I think that what is making the difference, and this is what I feel in this room, is that when you have the trust of your people, when you have great clients, and when you have partners like you, everything becomes possible. Thank you very much.

JUDSON ALTHOFF: Thank you, Arthur. Really appreciate it. Our final pillar of this morning's keynote is arguably the most important, and that's the pillar of trust. In order to enable trust inside of your organization, you have to have observability at every layer of the stack. You have to be able to manage, govern, secure, and get business value out of all of the AI assets that you deploy. Here to talk more about that, please welcome Charles Lamanna.

CHARLES LAMANNA: Thanks, Judson, and good morning, everybody, or I guess this afternoon now. It's great to be back here at Ignite with all of you. Work now runs on AI. Every company is deploying AI, and work is increasingly being done by agents and with AI assistants. But further adoption is constrained by security and governance risks. Without observability and management of your entire AI estate, you really can't start to reach scale. You have to be able to manage the cost, the security, and the agent sprawl. And that's why one year ago here at Ignite, we announced Agent 365.

Agent 365 is your observability platform for AI. You can manage your agents centrally in one place. You can govern their deployment and their enablement, and you can track the cost and value of each agent that you build. And you're also able to keep your employees and agents secure over time, so you can take advantage of the capabilities without compromising on risk and compliance.

An Agent 365 does all of this across Microsoft and non-Microsoft platforms. Agent 365 supports the key partners across the AI ecosystem. That's platforms like AWS, GCP, OpenAI, and

Anthropic, and apps and data backends like Databricks, Snowflake, Salesforce, and more. Agent 365 is that one place to manage all of your AI investments? And Agent 365 does that by maintaining a registry of all your AI assets, things like your agents, your plugins, your skills, models, and more. And you're able to join signals and create policies that span all of them. And that registry is just the starting point.

Agent 365 applies policies around cost. You can do things like set limits and alerts at the agent, departmental, or individual level, and you can even empower IT or business leaders to manage their own budgets and usage at their level. Companies around the world are already operating AI at scale with Agent 365. They're managing over 50 million agents, and we have tens of thousands of companies like Lumen, AT&T, KPMG, EY, and others across every industry and every region. To show that in action, let's welcome Jeremy Legg and Vandana Sharma to the stage together to show how Agent 365 is used at AT&T.

JEREMY LEGG: All right. Thank you, Charles. At AT&T, we connect hundreds of millions of people every day, so when we implement AI, we have to do it at scale. We're operating over 60,000 agents as well as a multitude of skills across our enterprise, and it cuts across all of the divisions of our company as well as our millions of customers. And while the technology is incredibly important, the transparency, the governance, and the visibility is even more essential. And that's why we're excited to use Agent 365 as part of our enterprise. Let's see a little bit about it.

VANDANA SHARMA: AT&T is already running agents at an incredible scale. Jeremy, let's look at what's new in Agent 365 for admins to go even further. This is Agent 365, the observability platform for admins to see, manage, and govern all of their agents across Microsoft and Non-Microsoft services. And with the expanded capabilities in connected platforms, admins can get deeper insights into services like Databricks that you see here. And right here in the registry, let's look at the wireless sales agent. This helps sellers recommend the right plan to customers. This ID gives it a unique identity in Entra so admins can govern it like any user. And everything that's connected to this agent is visible via this beautiful Agent Map: tools, resources, people, AI models, all here. And this red line shows admins any exceptions or errors so they know where to look for to troubleshoot.

JEREMY LEGG: This ontology is incredibly important because when you're operating agents at scale, you need to know both what they are doing as well as what they shouldn't be doing, matching the intent of the agent against the actual results of what is performing so that you can take the appropriate actions across the enterprise.

VANDANA SHARMA: Exactly. And sometimes, it's really hard to know where to start. And that's where Action Center comes in. It brings everything that needs my attention right here. Groups by category, ranked by priority. Even gives me the next steps. Let's look at this alert for agents that haven't been used in the last 90 days, a threshold admins can adjust. And instead of chasing each one down, admins can now set automation for cleanup at scale so it notifies owners, blocks idle agents, removes what's no longer necessary. Less clutter. Less risk.

JEREMY LEGG: Forgotten agents can become rogue agents. You can think of them essentially as digital employees. And how do you monitor what they're doing, knowing that they maintain levels of access into your systems, your privileges, your access layers, your authentications? And so you have to have the observability associated with those and be able to take automated actions at scale so that you can both manage the cybersecurity side of it as well as your costs.

VANDANA SHARMA: Exactly. And speaking of costs, Action Center keeps track of that as well. Here we can see agents that recently had a cost spike. And spikes aren't always bad, but admins do need a control over that. And two of these agents recently switched to a different model. And this graph right here shows spending forecasts for three different options. One, if they did nothing with it; two, if they reverted back to the old model; and three, it shows the price right before the spike happened. And to mitigate this risk, admins can now set spending policies. And right here, you can see that once these two agents hit a certain threshold, the right person is notified. And with Microsoft agents, it goes a step further. Admins can actually set hard limits, so you should just stop as soon as it hits the cap. And with spending groups, business users can now set budgets on a day-to-day while admins can set guardrails in place.

JEREMY LEGG: So for all of you CTOs and CIOs out there, this feature is the one that you use to please your CFOs. And how many times have you been asked about token costs? And so you need to have visibility back to the business so that there's transparency not only about what the costs are, but also the benefits and the revenue that you're generating for the company. You don't want to guess at those numbers.

VANDANA SHARMA: Right. So that's what's new in Agent 365 for admins, so they can have the visibility and control to manage AI at scale across the entire state so they can move faster with confidence. Thank you so much, Jeremy. Back to you, Charles.

CHARLES LAMANNA: Thank you so much, Jeremy and Vandana. AT&T's story really shows that one of the most important parts of Agent 365 are its cost controls. Cost controls tell you where the money is going and keeps your spending on track, but calculating your return on investment also requires understanding the value that's being generated. That means connecting investment to business outcomes across your agents and projects. And the new Agent 365 Business analytics experience brings AI spend, adoption, and value together in one integrated experience. It measures the work that gets done, not just the tokens that are spent.

It covers the tasks that your AI completes, how much time it saves and how much value it creates. It gives leaders a directional view of how AI is helping their team and their business move forward. And you can change things like your budget, your model mix, and platforms to balance cost and performance over time. The goal is not simply to spend less on AI, is to get more value from every dollar that you invest. Now, let's see how this works with another demo. I'm excited to have Stephanie join us and show us the experience for business leaders.

STEPHANIE CRONIN: Thanks, Charles. We just saw how IT admins manage agents and their costs in Agent 365. Now, let's switch seats to a business leader. Here, we can see a research and development operations leader has just received next month's AI budget. They're working in the new business analytics app with a view that scoped to R&D. This overview page shows spend,

use, and value across Copilot Cowork, GitHub Copilot, and agents. But every budget needs a plan, and so they start in cost and forecast.

This chart follows every dollar from each of those services to the models doing the work. And they can see that Copilot Cowork is currently 39% of R&D usage today. And they want to see what it's doing. The use cases and summary map shows exactly that. It shows task inferencing so they can see what R&D is using Cowork for like hypothesis testing. And they can see the cost right next to the value. And leaders determine how this value is calculated using their own task times and costs so it reflects their business. And it's just a conservative estimate of the work's value, net cash savings.

Back on the previous chart, Copilot Cowork usage is spread across several different models. This is auto-model at work. Selecting the right model for the task at hand. Optimizing value and cost, and maximizing the value of Cowork. Now, that is the power of multimodal. But they also want to take a look at GitHub Copilot. And so if we scroll up to the top, they can see that the cost drivers for GitHub Copilot are these groups and these job functions, and these are Teams and engineering. Looking closer at the costs and credits over time, they can see that monthly usage has been steadily increasing, and they need to preserve room for other services, so they run a custom forecast. They use a model mix that leverages lower-cost cloud models, and immediately, they can see that it has a substantial impact on the budget.

But they also want to understand the impact on the work itself. And so before taking any action, they send this forecast to the engineering department head, who can then run the pilots and necessary evals to ensure things like output quality and business continuity. Now, there's one more piece to the plan. The agents themselves. Here in the Agent Registry, they can get a single view of both Microsoft and Non-Microsoft agents. Let's take a look at this first agent here. It's built in Copilot Studio, and they can see the cost of the agent. They can also see things like the use case summary that shows the work it supports, and they can see model insights and group comparisons.

Now, since the cost here is being pulled directly from Copilot Studio, they can set up things like a spending alert or a spending limit that stops spend when it's reached. Now, Agent 365 is open by design. And so if we go back to the agent registry and take a look at this trial recruitment agent, we can see that this agent that's built on Amazon Bedrock shows information that's been pulled from the connected platform and being shared to Agent 365 so we can see things like the estimated cost. And with that, they can set up spending alerts so that when that cost reaches a certain threshold, they notify the right people.

Now, the more a platform shares, the more controls Agent 365 can offer. Now, if we take a look here at spending policies with all of these AI insights, they're ready to take some action. The leader has already created a spending policy for Engineering's Copilot Cowork usage. And it has a monthly limit. And it's being charged back to the engineering team via their Azure subscription that's tied to their cost center. That is Agent 365 Business Analytics and the budget for next month plans. See it, forecast it, budget it, and charge it back. Everything you just saw is available today, starting in Frontier Public Preview. And we have delegated spending policies and non-Microsoft agent cost visibility coming in December.

Thank you. And back to you, Charles.

CHARLES LAMANNA: Thank you, Stephanie. And that's how we can move from just controlling AI spend to really starting to understand and manage the value it creates. Agent 365 helps you manage your AI across your entire company, across Microsoft and other AI platforms like OpenAI, Anthropic, AWS, and more. But operating AI at enterprise scale also means securing it. I want to invite Yina up on stage to show you how Agent 365 helps you do exactly that.

YINA ARENAS: Thank you, Charles. As agent fleets grow across the enterprise, so does the need to manage their identities, access, and actions. Let's take a look at the new security capabilities in Agent 365 that give you greater visibility and control. The agent registry is we can see and manage all the agents across your organization. Here in shadow AI, I can discover agents running on employee devices. And we're expanding the local agent types we can discover for a total of 65, including OpenClaw, OpenAI, Perplexity, and many more. Let's take a look at OpenClaw.

Here we can see which devices the agent type is deployed and immediately see where my attention is needed. Microsoft is introducing a Dynamic Agent Risk for every managed agent, including local agents. This Dynamic Agent Risk brings together runtime, exposure, and potential blast radius so you get a total risk picture.

The related incidents show what contributed to the risk using the same incidents that the security team is already working on. That way, IT and security teams work from the same prioritized list and risk assessment, which follows the agent and adapts as risk posture changes. Here, IT can also govern agents on devices. With Intune integration, it can block the agent, flag devices that are not compliant, and scope the agents access with MXC from Windows. Let's apply it.

And that's the unique value for IT and security teams with Agent 365. They operate over the same prioritized inventory, complete risk understanding and policies. Back to you, Charles.

(Applause.)

CHARLES LAMANNA: Thank you, Yina. As these agents access data, use applications, and start to take action, security has to become a critical part of your agentic estate. To show you how we can secure that broader estate, please welcome Aleš to the stage.

ALEŠ HOLEČEK: Thank you, Charles, and hello, friends. It's really great to be here with you, especially talking about security. As we all know, advanced AI models and increasingly more autonomous agents are changing every industry, and cyber is no exception. Over the past 12 months, frontier and open-weight models have become remarkably good at discovering and exploiting vulnerabilities.

As a result, we see three significant trends. Number one is increased volume of vulnerabilities. Just this month's Patch Tuesday, we addressed more than 1,000 vulnerabilities. Compare this to

the fewer than 100 in March. That is 10x increase in less than a year. And it's not just Microsoft, the same pattern holds true for all major platforms, browsers, and open-source packages.

The second trend we see is reduced cost. Bounty programs pay about $300,000 for high-quality vulnerabilities, with the help of AI, researchers, and also the bad guys are now finding them for less than $5. And then the last trend is time. The time window between discovery and exploitation is approaching zero.

This is driven by AI being able to immediately exploit vulnerabilities using toolkits like ARTEX, but also because the process of responsible and coordinated disclosure is collapsing. We are on a mission to responsibly put the same technology attackers are using into the hands of defenders. That's why we built completely new cyber stack designed for agentic security.

At the foundation are the products 1.6 million customers already trust: Defender, Entra, Purview and Sentinel. They act as sensors and raise alerts when something is going wrong, but detection is not enough. When under attack, you need to act at machine speed, and that's what the actuators in our security portfolio do. They let you contain a device on the network, revoke a compromised identity, or create a block at the authorization or data layer.

We understand your security stack is heterogeneous and ours is too. That's why we have built more than 500 security connectors. Our products, combined with these connectors, generate more than 165 trillion signals every day. And we take these signals and we enrich them with threat intelligence from Microsoft and also threat Intel from you and also your business intelligence. We turn all of this information into a real-time, token-efficient security context that agents can reason over and act on.

Every day we evaluate this context for precision and recall so agents can act on what matters. At the top of the stack are the agents, the harness, and the models. Specialized agents work together to deliver protection and security outcomes. All of this comes together in a Microsoft Perception, our new agent security system generally available today. Please welcome David Weston, our own frontier security engineer and the architect of Perception, to show you the product.

CHARLES LAMANNA: All right.

DAVID WESTON: Awesome. So let's told you how our attackers that are moving at machine speed, but we built Perception, our new agentic security system, so defenders can move just as fast. And let me show you how you can do that. So we're hearing Perception where you can do things like ask questions, but you can also delegate work to singular agents or teams of agents. You can set them up to automatically trigger, and you can manage them as teams, just like you would a human set of security folks.

So let's jump into the agents and get you a view of what these look like. So our agents are created to mimic the way human security teams operate. So the first thing we have is a team of agents we call our blue team. So they can hunt, they can detect, they can investigate just like expert humans can, and you can delegate work to them.

Second, we've introduced this concept of red team agents. They can proactively analyze your environment, highlight risk just like highly-skilled pen testers can. They can even operate with threat intelligence to mimic sophisticated attacks. And then finally, this is what really matters, the green agents. They're the fixers. They can come in and repair your posture, fix code, help with detections, all sorts of very useful things.

And where this really gets interesting is the concept of playbook. Think about being able to create your own team to do highly sophisticated tasks like hunting, analysis. Anything you can imagine, you can create in a custom playbook. So let's take a look at my favorite playbook.

You can see here right away that you can actually transparently see what it's going to cost and what the time will be. You can also set up a trigger of when it should operate. And really cool, you can see the threat intelligence agents, the recon agent are involved here, even a web app pen test agent detecting authoring and posture. So this all comes together to do sophisticated work in a really simple way.

And then finally, this is the place I call the frontier engineer's inbox. This is a Kanban board where you can manage agents. You can see which ones are requiring human in the loop or human input. You can delegate tasks and you can understand. And then finally, just as important, we have our FinOps dashboard that explains what agents ran, what they cost, and most importantly, what the outcomes are. And this is all propagated in M365. So this all comes together to be a frontier agentic security system and an amazing tool for defenders. Back to Aleš.

ALEŠ HOLEČEK: Dave, that is some awesome sauce. Thank you. Perception was built to deliver measurable security outcomes, not just raise alerts. We evaluate every agent and playbook against rigorous benchmarks grounded in real security work. Our multi-modal design roots each and every task to the most qualified model, and that includes MAI-Cyber-2, closely co-developed with Mustafa and announced earlier by Satya.

We trained the MAI cyber models using decades of real-world security data and hill climb in learning environments designed by industry leading security experts, reflecting the most up-to-date threat landscape. The result is effective protection better than GPT-6 Astra on CyberGym end to end, and a 2% of cost per task making agentic security affordable at scale. Let's hand it back to David to see how Perception helps security practitioners like him reduce the exposure and defend with agents. David.

DAVID WESTON: Okay, so I'm going to transform into a security engineer at Caldova. And as a frontier security engineer, I have two big jobs. The first is I need to protect the estate. I need to analyze configurations, exposures, apply patches. This keeps the bad folks out and creates a fortress around my company. The second job I have is detecting and responding. If someone gets through that fortress; it's all about interdicting that attack, detecting, and responding quickly, and healing the enterprise.

So let me show you how Perception makes this extremely easy for a frontier security engineer. So the first thing you can see here is that Perception has popped for me something I need to be concerned about, which is there's been a Patch Tuesday push to all of my company. And within

that Patch Tuesday, there is a vulnerability that's being actively exploited. This is something I cared about deeply.

So I'll pop open a Perception session, and the first thing it's doing for me is a bunch of work to analyze in my enterprise which machines have actually been patched or not, and which ones would be vulnerable. So I get a great report here on my client devices, my servers. You can see here about 52,000 devices and a few thousand servers that still need to be looked at.

This Patch Tuesday is a big one because of the agentic push. There's about 1,500 distinct CVEs in there. I want to push out patches that I need to immediately, based on exposure and criticality and exploitation with threat intelligence, so I get all of that in a single system. Normally, this would have required me to write custom PowerShell, do analysis, dig through CVEs. All of that information is compacted into a singular pane here with the agent, and it does all that work for me.

So I want to move quick, and I want to expedite, so I push this out. Right away, the critical patches are starting to be applied on client systems. But on server systems I want to be a bit more conservative, so I use our new endpoint shielding technology to protect these servers, to give me a little more time to analyze the impact. Again, that's something that would have been really complicated without these agents. Here, it's done super simply for me.

So now I'm going to switch to my other job, which is analyzing alerts in my enterprise and acting on them. This is a place where I spend a lot of my time in the detection role. You can see here there's a bunch of cases that need to be triaged that are open. This is where Perception does the hard work for me.

So each one of these might represent hours or even days. Here I can just jump in, ask my agent team to do the work for me. So I want this to be triggered at any time there's a new alert, and each of the triage attack investigation and the detection authoring agent are going to work on each of those alerts. And I can see that it's going to take between 10 minutes and an hour, depending on the sophistication of the investigation.

So right away, my agent goes to work, and you can see all these different investigations. The first one I want to look at actually demonstrates how Perception can use third-party data. This is an alert that was detected based on the usage of Okta data and AWS cloud data. So we'll jump in here. We can see that Perceptions has done a lot of the heavy lifting. It's already triaging. It's generating a report and findings, and it's actually going to create some custom detections if needed.

So let's drop in on our agent doing a bunch of the work. We can see here that the attack investigation agent has already pulled up some findings for us, and the detection authoring agent is starting to enrich. So we'll let this run for a second, and the first thing we see here is there's a new detection report. So we'll jump into that.

And what it's giving us an idea of is as the investigation happened, we did a good job of detecting, but we can actually leave it better than where we found it by writing a few more

detections to augment our chain. Traditionally, this would take an expert security engineer to do this and someone who can also write custom code. But with a combination of APIs and actuators and the coding model in use here, we can actually write custom detections here to cover the entire attack chain.

So I'm human in the loop here because this is a sensitive thing. I want to make sure I'm in control. I can review all those rules, publish them with a single click. These are also being checked for false positives, things like that. And in the end, I can see that three new detections were deployed, my scope of detection got better, and overall my enterprise is better protected with Perception and agents. With that, I'll hand it back to Aleš.

ALEŠ HOLEČEK: Yeah. Thank you, David. Thank you for walking us through how Perception takes on threats on the perimeter. But AI also introduces new risk inside of our organization: Agents. While not inherently malicious, they are frequently overprivileged, have access to a broad set of tools and high-value data, and are increasingly more autonomous executing business processes around the clock.

Perception and Agent 365 share the same context and work together to secure agents across your entire enterprise at machine speed. AI has changed the physics of cyber, and that's why we built Perception, an agentic security system that helps you secure your foundations, defend against threats, and protect your agent insiders. Judson, back to you.

JUDSON ALTHOFF: Thanks, Aleš. So I'm going to close this out this morning, and I'm going to do it by sharing the perspective of arguably what has become one of the hardest jobs in the entire industry, the Chief information officer. I'm here with Leigh-Ann Russell, the CIO of BNY. Leigh-Ann, welcome.

LEIGH-ANN RUSSELL: Thank you. Thanks for having me here.

JUDSON ALTHOFF: So we have this amazing relationship. You have done so much leading technology and engineering at BNY. Tell us a little bit about the work that you're doing.

LEIGH-ANN RUSSELL: So as someone who leads technology at a highly regulated industry, governance and security is something I'm super passionate about. At BNY, we put that at the heart of everything that we do, including in our AI journey and our build of a platform, Eliza. And I firmly believe that if you build in governance and security at the heart of a project, that's how you scale fast. If you do it later, it becomes a break and that slows things down.

Now, our teams have been putting those principles to the test and something very, very cool. Under our agentic negotiation protocol, we have brought together a Microsoft agent and a BNY agent, and together they have executed a Treasury trade negotiation, which is a market first.

JUDSON ALTHOFF: Awesome.

LEIGH-ANN RUSSELL: And now they're going to learn how to scale that.

JUDSON ALTHOFF: Super cool. So what gave you the confidence that you could manage, govern, and secure these agents operating in your environment, and how did you get after it?

LEIGH-ANN RUSSELL: Well, our teams working together put those principles to the test. First of all, they really sweated the detail on what they were going to build. And this is really important in agents at scale, in regulated industries. First of all, they said, what are the roles of the human? What is the role of the agent? What are the permissions that are going to govern this trade, and what's the accountability? Where does the accountability lie?

And then they tested, tested, tested before we went to a live market. And even when we went live, there was a human over the loop who looked at the whole process and had the ability to pause or even stop the trade. Now, when it came to the final trade execution, a human made the accountable decision, and that's how the trade took place. And those principles allow us to have that confidence.

JUDSON ALTHOFF: Awesome. Let's pivot to the cyber side of things. You've been involved in our private preview of Project Perception and now Perception's GA. We're very excited about that. But share a little bit about what it's been like to go from a reactive mode to a proactive mode in the world of cyber.

LEIGH-ANN RUSSELL: Absolutely. When the bad guys have AI at their fingertips and go in AI time, our teams have to also have great AI tooling and also work in AI time, not human time. So we're super grateful that we were part of preview, and that allows our teams to have the best tools at their fingertips to fight cyber-crime.

But it's not just about having the tools. These tools have to give our cyber professionals great visibility about what's going on so that we still can have our cyber teams really in control of the defenses. Now, my word for the year was coherence. My word for next year is interlock. It comes from my background in physical engineering. There's four "I's" to interlock. Intent, how do we know what we're building is designed by intent? Integrity, governance, and security by design. Infrastructure, Satya talks about it. Do we have the right infrastructure supporting our AI build out. And then impact, everything that we do has to show up with an impact for our clients.

JUDSON ALTHOFF: I think you framed it super well, but the pressures of doing all of this in a highly regulated industry have to be super intense. But bring us all together on the color around that, what it's like.

LEIGH-ANN RUSSELL: Well, it's about having these controls. That's what gives us confidence to move forward. And you don't get to be 242 years old like BNY is without being at the forefront of innovation. And we want to innovate fast, but we want to do it within a framework of trust. Our clients depend on us to be resilient. Our clients and our regulators depend on us having discipline. And I think technology can solve both of those things, but it has to do it with that framework. And I think at the end of the day, trust is the product.

JUDSON ALTHOFF: Awesome. Thanks so much for being here with us and thanks for the partnership, Leigh-Ann. I really appreciate it.

LEIGH-ANN RUSSELL: Thank you.

JUDSON ALTHOFF: Thank you. This morning, you've seen how we're building the frontier ecosystem with our customers. You've also seen how the three most durable assets for this transformation are human ambition, intelligence, and trust. Every technology wave has been created to serve humanity, and this one will be no different. Infinite cognitive compute will never replace human purpose.

I had the opportunity to host an event recently for Microsoft employees, and we organized a party for the kids to free up the parents. At the front door, we set up these posters, and I asked the kids to write down their biggest ambition. I found these to be inspirational for how we think about the journey ahead.

Now, there's some fun answers here like, I want to go to Japan. I want to become a football player. I want to touch a King Cobra. And while it's true that AI might be able to help with some spelling, if you look beyond that, you see the hopes and dreams of the future generation. Nowhere on here do you see the words "prompt engineer," and it's pretty clear that our children will not want to live in a chat window.

Frontier transformation is intended to enable human ambition. People want to see AI lit up in the canvases that they use every day, in the flow of their work, so they can spend more time on what matters most to them. We've given AI models far too much credit in the intelligence category. Intelligence is not the aggregation of information or the inferencing over the same. Intelligence means something more: Your unique ideas, your organization's unique differentiation.

Throughout history, the gaps between the haves and the have nots have correlated to the access to knowledge. AI gives us the opportunity to democratize intelligence more than any technology has ever done before and solve real problems for society. The cure for the incurable. Advances in science and research. Nourishment and prosperity for all.

Intelligence must compound from within. The frontier ecosystem is not about the success of any one model or a model lab seeking the world's GDP as its total addressable market. It's about your IQ, your unique differentiation amplified in an open, model diverse, and heterogeneous way, achieving your highest aspirations while protecting the durability of the same.

In the end, it comes down to trust. Take this completely out of context for a minute. Consider the last time you invited someone into your home to do some work. There's this initial moment where you must decide to trust this person. They're coming into your most personal space. There's an actual contract for service and the associated costs.

You want them to perform the job at hand and only that job, and you want to see value, and you need to feel a sense of security. That's the kind of agreement we have to have with AI. The only way to gain trust is through observability. If you can't observe, manage, govern, secure, and measure the value of the AI solutions in your technology environment, you simply can't trust it.

We have over 50 years of building security into everything we do within the enterprise, and we are applying that to the safe application of AI. Everyone in this room has something that they're passionate about, and we ask the employees from the customers presenting with us here at Ignite about their ambitions because technology doesn't have ambitions, people do.

When technology has the potential to reshape every job, every workload or workflow, every company in every industry, you need a trusted partner. I can say confidently that Microsoft is the adult in the room relative to the transformation journey that lies ahead. And today, we've shown you a comprehensive AI platform for building the frontier ecosystem.

We're not here to empower every AI model or agent or every GPU. We're here to empower every person and every organization on the planet to achieve more. Let's all stay grounded on the principle that AI must do more for humanity and go build the frontier ecosystem together. Thank you.

END