

10212024 London UK AI Tour Keynote Satya Nadella

CEO EMEA Tour

AI Tour Keynote

Satya Nadella, CEO & Chairman , Microsoft

Jared Spataro, CVP Modern Works & Biz Apps

London, UK

Monday, October 21, 2024

(Applause.)

SATYA NADELLA: Good morning. It's fantastic to be back in London and in the UK. Especially at a time like this, when there's a new tech platform being born, it's exciting to be able to talk about it and see its impact in the UK, and talk about that vibrancy that comes with these platform shifts.

In fact, I was just recounting this morning, Microsoft has been in the UK for four decades. I myself have been coming to the United Kingdom for the last three decades across, I think, first – in fact, I was thinking about it. The first keynote, obviously, I didn't do the keynote, I did a breakout session back in the early '90s was, I think, Excel with VB demo was what I did.

It's fascinating. I came here when the PC, client-server was being born, and then went through the web and the internet, and then cloud and mobile, and now here we are at the very beginning of what's, at least, this new tech shift to AI.

I think it's always helpful, and I found it very helpful to ground myself in what's the core driver. I remember even joining Microsoft back in '92, the implicit understanding I had, at least of the tech arc, was all Moore's Law. In fact, I remember going to PDC in '91 and saying, wow, it's so clear what's going to happen next, which is X86 and the PC architecture will not just win the PC, but it's going to basically win the server. It just was the case by even the end of the '90s.

The same thing is happening in some sense in terms of a new law that I think we can all track, which is the scaling laws, as people describe them, in AI. It's an empirical law. Like Moore's Law was also not a physical law; it was just an empirical observation that we talked about it as if it was a law, but it held true. Similarly, the scaling laws basically are showing that every six months or so, we now have doubling of capacity.

In fact, one of the things that I think a lot about is performance. You can even say tokens per dollar per watt. That's the new currency, and you measure it by that. Every six months, we have doubling. Some of it comes because of compute power. A lot of it actually comes from just better techniques around how to use data and algorithms.

This is something that's inflected. You could say this began at the beginning of the DNN era in the early 2010s, but 2018, 2019 with the LLMs and the transformers, they inflected, and it continues.

It manifests in three fundamental ways. The first thing is, I think computing interface as we know of it is fundamentally changing. Once you have a natural language which is multimodal even, there is image, speech, text, video, in and out, that means every computer interface is going to change.

You have increasing reasoning capabilities. If you look at even what just came out with o1, it just gives you, whether it's planning, reasoning capabilities. Seventy years of computing history has been about digitizing people, places and things and making sense of it. We now have a new reasoning engine to make sense of it. Then lastly, you can feed it more context, more memory.

You put all these three things together, you're building out a very rich AI or agentic world, in which you are going to have these AIs or agents. There will be some AIs and agents that are personal agents. There will be things that will work in the context of a team, in the context of an organization or a business process, or even cross organization, this rich tapestry of AI agents that augment everything else that we built. That's the other part, the entire digital infrastructure and tools that we today have get augmented in this agentic world with all these AI agents that we build using the scaling laws as the underlying force.

Now, of course, it's great, all this technology is going to be there. The question is, what do we do? How do we, most importantly, parlay this, translate this into the most important mission that we have as a company, as individuals and organizations, is to be able to empower people to do things that they couldn't do previously with any other previous technology era?

To me, that's the ultimate test, which is, can we use – Clare started by talking about how right here in this country, some of the most seminal technologies of the industrial era and the Industrial Revolution were all created. The question is, can we now go back to that and really have that type of profound impact in human life and condition by really, whether it's scientific discovery or whether it is productivity, can we see that flourishing of innovation because of technology?

That's our goal, that's our mission, to empower every person and every organization right here in the UK to achieve more, whether it's small businesses becoming more productive, large multinationals right here becoming much more competitive globally, public sector becoming more efficient, health outcomes, education outcomes. That's really what this is all about.

Now, to us, to achieve that, we are building three platforms. The first is Copilot. To us, you should think of Copilot as the UI for AI. That's the simplest way I think about it. We then have the Copilot and AI stack. To be able for you to build your own AIs and AI

agents and Copilots, we have a full stack. And then lastly, this new set of devices, which are these Copilot devices. I want to talk about each of these platforms, starting with Copilot.

Now, as I said, if you start with this idea that this rich agentic world ultimately does need to meet us, and we need to meet it, that means you need a UI interface. Just like the PC or the phone was the user interface or the apps on a phone or a PC were the interface to essentially digital technology, these Copilots and Copilot is the UI for all of this AI.

Even in a world where there are a lot of agents that are working autonomously, they do need to raise exceptions get permissions from us, and the question is, how does that happen? It happens through this new organizing layer for how, in particular, work gets done. In fact, work, work artifact and workflow is going to change.

A great example of this is just a month ago, we launched something called Pages. Just like, say back in the day, in the '90s, we launched Excel or Word, which were basically editors to create new artifacts, Pages is the first, I would say, user experience to create new AI-first artifacts. I can search the web or my work for retrieving information, and then I can put it into pages, and it's a document that I can then share across the organization, and I can work with AI and humans.

In fact, I say, the metaphor I use is, I think with AI and work with my colleagues at work. That's the new workflow. What was the previous workflow? I thought on my own, I created artifacts, and I shared it across the organization and collaborated. But now I not only have a cognitive amplifier effectively with AI, where I do my work, and then I create artifacts and I collaborate with my colleagues in order to get things done.

That's really the beginning of this Copilot era where it's just not about a chat interface, but it shows how chat is just one modality of being able to retrieve information, but it does lead to more sophisticated workflows and collaboration.

Now, you extend. The other thing is that this is not just about any particular artifact editor or workflow we created, but you can extend Copilot with any agent you build. In fact, Copilot Studio is a low code/no code way for you to be able to build agents.

These agents are really grounded in a rich set of data sources, starting with, in fact, the most important database in most organizations is the database that contains all your office information, who works for whom, who are my colleagues on this project, what documents are there related to a particular team or a project, what is the relationships between all of these documents and people and projects, all of that. All the emails you've had, Teams conversations you've had, that's all, in fact, in a first-class database called the graph or the substrate that is now exposed through a graph in M365. You combine that with even Dataverse, which is all the business process data, data that you may have collected into something like Fabric, all that's available for these agents.

In fact, a great simple example of this is, let's say you wanted to build a field service agent that you want to interface into Copilot. All you've got to do is give it a system prompt, tell it that, hey, I want you to be a field service agent, point it to a SharePoint site where there's a bunch of documents related to field service, add to it even additional data sources, in this case, Dynamics as the system of record for field service, and you have an output which is essentially a field service agent, which now you can talk to and have a conversation with, just like you would with any other regular Copilot conversation.

That simplicity, it's kind of like back in the day, we just created an Excel spreadsheet. It's no more mystical than that. Just like how you could create an Excel spreadsheet that was a forecast, you can now create AI agents using a low code/no code tool like Copilot Studio, put it into Copilot. You can even think of these as the new form of applications that any one of us can create. It's not even that you have to wait for someone else to create this application for me able to use it. That's the Copilot and Copilot plus agent ecosystem that we're building out.

The impact of this has been tremendous. In fact, right at Microsoft, if I look at it, if you take functions like, let's say sales and marketing, we now have quantitative results which are double-digit increases in throughput effectively. This is top line impact.

You take something like customer service, help desk, IT help desk, HR help desk, that's where, again, our employee engagement, our employee satisfaction, our customer service agent satisfaction is going up, and the waste and costs are coming down. Same thing with legal, same thing with finance.

We now have at-scale evidence of how these tools are fundamentally changing, I would say increasing value and reducing waste. What industrial companies in the past have done with things like lean is happening, finally, with cognitive work at scale.

It's also happening right here in the UK. In fact, I had a chance to see customers who are already using this across the board, and it's fantastic to see the local Copilot example. I had a chance to meet with Clifford Chance, which is a legal firm, and they explain the workflow that happens around M&A transactions. It turns out that using things like Copilot and Copilot Studio, you can create these agents that streamline that entire process.

I also had a chance to meet with Unilever. I had not realized this, but for an organization like Unilever that has a lot of topline marketing spend, because, after all, they reach three-plus billion users around the world with a multiplicity of products, one of the biggest things they do is these creative briefs. The creative brief, the fidelity of it is so critical for them to be able to then have the marketing impact.

But the point is, the drudgery involved in creating the creative brief, or the efficiency or the accuracy of it, that's where they've gone to work and created this fantastic tool that has just changed both the prep of it and the quality of these marketing briefs. These are

just two examples of customers of ours who are already deploying these AI agents and Copilots to drive real high impact productivity inside of their organization.

Of course, today, the exciting thing for us is that we are making the announcement with the government of the United Kingdom to really ensure that this diffusion of tech happens not just in the private sector, but happens across the public sector. Because at the end of the day, I think perhaps one of the biggest pieces of impact we'll have is in the services that governments offer, whether it's in health, whether it's in education, whether it's in energy, or any sector of the government that will be transformed using this technology, and it starts by actually putting this in the hands of civil servants all over and for them to be able to use it in their own work. Because after all, when I talk about reducing the drudgery, improving productivity, which then impacts the services our citizens get, this is the place to have impact. We're very, very excited about this announcement.

In fact, talking about the skilling, the UK, the skills around AI, this is LinkedIn data. When I saw this last night, it was fantastic to see how we now have a monotonic increase in the total number of people who are taking credit for all the credentials they've gotten around AI skills. When you see something that's grown by 88% since 2019, it speaks a little bit to what Clare said in the opening remarks about the structural advantage of this economy, because the human capital and the rest of the infrastructure that's already getting built I think can catapult the United Kingdom in terms of what it can do in the AI era, and it's fantastic to see.

Of course, we're not stopping here. What we are really excited about is what's coming next. I'm very, very pleased to announce today for the first time, that we are taking the next big step in the AI platform with the announcement around autonomous agents that you can create.

In fact, today you will see us both make announcements around tools that are going to help create these autonomous agents, but more importantly, the agents themselves that we are building for products like Dynamics 365, which ultimately can dock to something like Copilot when it needs a UI for human connection, but also can stand alone.

To show you all of this, I wanted to introduce up on stage my colleague, Jared Spataro. Jared? (Applause.)

JARED SPATARO: Thank you, Satya.

I'm sure many of us are familiar with McKinsey & Company, one of the most successful management consulting firms on the planet. Now, for McKinsey, client experience is literally everything, and they're always working hard to ensure that they're streamlining and improving every touch point with their clients.

Imagine our delight when they agreed to partner with us to use Copilot Studio to create an autonomous agent to streamline a portion of that client experience. Let's take a look.

It all starts with an incoming email from a prospective client, much like you see on the screen right here.

Now, previously, they had had people on the back end essentially receiving these emails, parsing through them and figuring out what to do next. Who should it be routed to? What expertise did they have in the firm?

But this is where the autonomous agent comes in. Now an email comes in, and the agent springs into action. What you see here is that it will begin to parse out the email, moving through the ambiguity of human language to, for instance, find out what the engagement is about, to check the engagement history, to also map it to their industry standard terms, and then finally, to try and find the right person to take the next step within the firm. With all of this information in hand, the agent then goes about writing an email that takes all of this information and summarizes it for the receiving partner.

What you see on the screen is exactly that. In comes a whole bunch of human written email. The agent processes it, summarizes it, and sends it to the right partner in the firm to take that very next step.

Now, it's worth pausing for just a moment here to reflect on what you're seeing. It happens so fast, you might miss it. But essentially, this agent has been given a loose set of instructions, kind of like you would to a human, and it deals with all of the messiness of human communication, figuring out what the right next touchpoint is for the customer.

Now, this is magic, but it's only half of the magic, because now we're going to go behind the scenes to see how easy it is to actually create an agent just like this. For this we will move over into Copilot Studio.

Here you see that we have programmed up with McKinsey the agent, but not using a sophisticated programming language, instead using natural language, the same way that you would tell a colleague to get ready to do this task.

You also see that what makes this agent autonomous is that we can set what's called a trigger. In this case, the trigger is set to watch an email address and to react immediately when an email comes in. But in fact, you can set it to look for events across a whole wide range of systems, sitting there working for you 24/7 waiting for an event to come that gets it going.

You also, just like a regular human colleague, add knowledge. Here we see a Word documents, a SharePoint site and a database about engagements. But of course, you can add additional knowledge sources. That includes line of business systems like SAP or ServiceNow, or even databases.

And finally, to finish up what you give this agent to do its work, you give it a set of actions, and we saw those in the flow. These are actions that include things like pulling

out the relevant information or summarizing what a human has written. All of this together makes the agent powerful because it can deal, again, with all of that ambiguity that a human throws at it.

Now, what we saw was one email coming in about one new client engagement. But the exciting thing here is that this scales. How does it scale? Well, to see that, we'll go over to the activity pane, where we can look at the long list of engagements that it's working on.

Zooming in up top, for instance, we can see that it's worked on over 1,300 engagements, and there are 33 in progress. If we want more details, we can go into the analytics tab. What this means is that this agent is always working on behalf of the firm, and that's very exciting for us.

Now, from here, we also see that, although the agent's amazing, it does sometimes need some human help. We're going to jump into a case, the second from the top here, where we will see that it gets a little bit stuck. As you look, it's gone through those steps that we saw previously, but it's stuck here at that one where it's looking for the partner. If we zoom in, we can see why.

Here, for instance, we see that it's picked the right partner, but that partner has now left the firm. It has an instruction that says, if that's true, it needs to escalate to a human manager to give it someone else to go to.

Now to see what that looks like, we're going to switch over to Copilot and see that interface with that human manager. Here at the bottom right, you will see that a notification pops up in Copilot, then the manager gets all the information he or she needs and can provide the right person to route the email to. Back at the ranch, going back to our agent, we can see that it takes that information and fills out what it needs to do.

Now, we're excited about this because of the business value it can drive. McKinsey, in its trials, has shown that it can reduce lead time by 90%, reducing administrative overhead by 30%. And as you look at this list, what we envision is an orchestration layer, just a bunch of agents that can be out there helping individuals, teams and entire functions to streamline and automate their processes, no matter what industry they're in.

And they're so easy to make, anyone can do it. You design and set these Copilots out to work in Copilot Studio. You interact with them in Copilot.

We're very excited about the technology, and to give you a sense for how it's already being used, why don't we roll the film and see a couple of our customers.

(Begin video segment.)

VO: At Unilever? We believe there is a huge opportunity in the end to end marketing model to leverage Copilot, transforming a largely manual part of our marketing process, the Brief.

VO: At the Brief kickoff meeting without even prompting, Copilot could begin to gather all of the relevant content across data sources. It could create a first draft and when that's ready, share it with relevant stakeholders for their review. We call this smart briefing. That's quite powerful.

VO: This is not just about embracing change. It's about strategically positioning (inaudible) for the future while maintaining our core mission of delivering superior legal service. The Copilot agent, working alongside the lawyer, can autonomously identify risks, confirm the inclusion of required documents and notify stakeholders of next steps.

VO: That just allows me and my team to be free and available to deal with much more complex issues. So the concept that you could have your own personal compliance agent is so exciting.

VO: The beautiful thing about Copilot agents is it can connect into our existing ecosystem. When a potential fraud scenario is identified, the agent will autonomously go and gather the relevant information from that source system, perform an analysis and provide that analysis back to the fraud colleague. That always ensures that humans in the loop. We actually built a working prototype in the space of a week.

VO: It enables us to operate at scale and at speed, efficiently, accurately. And it allows us to take control over areas that we need to be focusing on to improve our customer experience.

(End video segment.)

(Applause.)

SATYA NADELLA: All right. Hopefully, that gives you a bit of a feel for how Copilot is evolving to becoming this organizing layer for work, workflow and work artifacts. And Copilot Studio Plus agents is really this orchestration layer that docks with the Copilot in helping all these agents work together in the context of whether it's individual's work or organizational work or business process. So, that's really how all this comes together.

Now, I want to move to the next platform, which is the Copilot Plus agent stack, or the AI platform.

Now, ultimately, you want to be able to have every layer of the technology stack underneath what we have done with Copilot and Copilot Studio, and the agents that you just saw, all available to every software developer to be able to build any AI system of their own. And it starts with having the widest, broadest footprint of raw infrastructure.

When we think about Azure, we think about Azure as the world's computer. And we're building this out in 60+ regions. Right here in the United Kingdom, we have UK West, UK South. In fact, last year, we were excited about the announcements and expanding this with \$2.5 billion of investment. And this is going to be ongoing. We're going to bring the best infrastructure to the UK, both traditional compute and AI compute, so that you have the base infrastructure available.

And we're doing a lot to make sure that this infrastructure is optimized for AI workloads, starting right from silicon, whether it's the work we're doing with Nvidia. In fact, I was just looking at even the GB200s that are coming online – I mean, in pilots with all the liquid cooling that goes with it. In fact, we borrowed some of the liquid cooling work, because we were doing it for our own silicon with Maia. And that's what's available now, even across different silicon. We're working with AMD.

And so, we have fantastic partnerships on the silicon layer. You then parlay that by really building the best optimizations, whether it's for training or for inference, so that you can then build your Copilot. So, a lot of work going on in core infrastructure.

Now, the other very important consideration for any developer building AI applications is data, because ultimately, whether to train or to inference and do things like retrieval augmented generation, you need to really have your data estate in order. And in order to do that, that means you want to be able to bring all your data to the cloud and rendezvous with AI. In fact, there's gravity. Wherever there is AI compute is where data will go.

And so, one of the things that we are doing is whether it's your Oracle estate, whether it's your Snowflake estate, whether it is anything else, you should be able to bring that to the cloud.

And then on top of it, we are, ourselves, building a very first class, I'll call it, cloud-native infrastructure, data infrastructure for everything from OTP, whether it's Cosmos DB, whether it's SQL, or whether it's Postgres or Fabric, which is a first-class analytics database built for the AI era.

The data estate is in a place where perhaps some of the best work is happening. And in fact, thinking about AI plus data, there's no such thing as an AI application that's not state full. AI APIs are stateless, but once they meet a real application and a real workload, they become pretty state full.

In fact, ChatGPT, by the way, is one of the biggest customers, users of something like Cosmos DB, for example, or Azure Search as another example. It just shows you, once you build an application like Copilot or ChatGPT, you need the data estate to be robust.

Now, we are also building an app server. If you say you have infra, you have data, then you have to have an application server. In fact, I remember even coming in previous eras here and talking about .NET. I think Scott is going to be here later, talking about AI. He's

the guy who did a lot of .NET work as well. And guess what? We're back again with another new app server era.

And in this case, by the way, the app server we built for cloud-native applications, whether it's containers, whether it's app services, all of those, AKS and functions are all still needed. In fact, again, when I look back at, say, the architecture underneath ChatGPT, what's the base compute provisioning that you do? In fact, for every GPU they use, they have a ratio of how much regular compute with AKS provision they have.

So, it starts with app servers. Then there is a new AI app server. And the AI app server, in fact, starts with the broadest model selection. We are very excited about our OpenAI partners with all of the innovation, even o1, to GPT-4o, so the all of the latest frontier models. But it's also everything in the open source, whether it's from LLaMa and Mistral or other closed source providers like Cohere. You have the broadest selection of models.

And once you have the broader selection of models, the next service in AI app services is for you to be able to do SFT or you supervise fine tuning of it. You have Finetuning as a Service, even, on top of all these models, so that you can build this into your application.

You have all the tools for things like retrieval augmented generation. Things like Azure Search help you ground your data, or rather, ground your applications and your use of LLM with data. On top of that, you even have services to be able to make sure that you have guard rails around your application.

We are very excited about all of the innovation happening across the length and breadth of the app server.

You have the app server, you have the AI app server, and then you have the best tool chain. In fact, the thing that I'm most proud of is Microsoft, since 1975 to today, gets most excited when talking about the tools we build for software developers. That's what is core to us. With VS Code, and GitHub and GitHub Copilot, there is an absolute new frontier for what a software developer can do with this tool chain.

In fact, one of the coolest things I have seen recently is with o1 coming to GPT, or rather, to GitHub Copilot, you can use AI to do the next level of optimization. I think what you have on the slides behind me is the auto encoder we use for GitHub Copilot is being optimized by o1. Think about the recursiveness of it, which is we're using AI to build AI tools to build better AI.

It's just a new frontier, the productivity. And in fact, and in fact, I think next week, or the week after, is GitHub Universe. I can't wait for it in terms of what you will see us do with even GitHub Copilot Workspace.

In fact, right here in the UK, we now have 3.7 million GitHub developers. It's fast growing. It's, I think, the fifth largest in the world. We have, I think 22% growth. It's fantastic to see, and I think there's going to be a really inflection, because the barrier to

entry, I always joke now, every weekend, I can go back and code again, which is just fantastic, just because, thanks to GitHub Copilot, any one of us can go ghetto to repo, clone it, and get to work on it, and actually finish a project in finite time. (Laughter.)

It's very, very exciting stuff, and then you can see it right here. In fact, this morning, I had a chance to meet with many partners of ours, in fact, starting with British Heart Foundation. What they've been doing using machine learning and AI for a long time, it's been fantastic, but even things like using Azure Speech services to do simulation of these emergency calls. I had not realized how important it is to actually help people get comfortable making even the emergency calls. That's just a great example of how you empower people at crucial times.

HSBC, there are many, many customer journeys, but one of the journeys they showed is how a relationship manager doing credit approvals, that entire process has been transformed using some of the AI tools they built, effectively their AI agents themselves, that intercept with their mobile apps, as well as the applications that relationship managers used.

I had a chance to meet the developers of Mandara, who are doing things, which are pretty unbelievable. They building an entire digital twin for the food safety and food supply chain for every retailer in the UK and beyond, in order to be able to improve the sustainability of it. The real-world impact that this product is going to have is fantastic to see.

And one of the startups out of here that I've been excited about for a long time is Wayve. And Wayve has taken an approach which is an AI-first approach to ADAs, how to build, really, from first principles, a very new AI driven way to think about automation. So let's roll the video to give you a flavor for it.

(Begin video segment.)

VO: Well, it's clear to me that artificial intelligence is completely changing the world. And we've seen so many advancements in the world of internet and software, but at Wayves, we're really passionate about bringing AI into the physical world, and we started with autonomous driving. That's exactly what we're building at Wayve. We're building an artificial intelligence system that can learn how to drive based on what it sees.

Our vehicles collect terabytes a minute from the cameras that we have around the vehicle or the radar, but the video data puts a lot of stress on the bandwidth and the scale that we have to work with. That's where our partnership with Microsoft has been fantastic, because we've been working on supercomputing technologies that really make this possible.

Ultimately, we're talking about billions of parameters and transform models that need to be trained on petabytes of video data. And this is a scale that autonomous driving hasn't seen yet, and we want to push the boundaries.

Even Azure ML has been brilliant for us, because it's meant we haven't had to build teams and people working on this technology, that we've had a managed platform that's enabled our systems to train effectively and efficiently on the Azure's cloud.

I'm excited to see our partnership grow with Microsoft as we invest more in the compute infrastructure that makes training at the scale possible.

I think these tremendous forces that, when brought together, will be what creates an autonomous driving solution. It can scale to everybody through it.

(End video segment.)

SATYA NADELLA: It's fantastic to see, again, that rate of diffusion, where it's no longer about, hey, this is going to come, but you can see how already, some very sophisticated applications and use of these AI platforms right here.

I want to then move to the last platform, which is the Copilot devices. Copilot+ PCs, we just launched a little over half a year ago. We're very excited about what this means. This ushers in a complete new era, where you have the CPU, the GPU and an NPU all available at the edge, because if you think about it, up to now, the scaling laws have worked and worked super well in the cloud.

But I think going forward, what's going to happen is this AI era is going to be defined not just by what's happening in the cloud, but what's happening in the edge. And we will have, in fact, fundamental breakthroughs even in some of the model architectures, which allow for this hybrid use of the fabric that gets laid out. This is not about thinking of this as old-world client and server. That won't work. You've got to think of this as one continuous, distributed fabric, and that's our architectural approach.

We don't think of the Copilot PC standing on its own. Of course, you can use it for privacy, but more importantly, you can use it in conjunction with everything that you're doing in the cloud. And that's what these devices will usher in.

Just to give you a flavor for what's happening with these devices, let's roll a video.

(Product video segment.)

SATYA NADELLA: It's one of the exciting things for those of you who are gamers, is going to be when you have your GPU in full tilt, and then you have all your NPU left for all the vector operations. That's the type of application development that is coming to life, and it's just going to be a new platform. Just like how people build novel, new applications for the PC with the GPU, you're going to start seeing now people build for essentially, what is going to be an NPU plus GPU plus CPU world. And it's going to be a pretty interesting world.

I just want to close out by talking about what perhaps underlies all these three platforms, which I think is very critical, which is trustworthy AI, and trustworthy privacy and trustworthy security, because trust in the technology ultimately, is going to be core to all of the diffusion I talked about, because if you don't trust it, you're not going to use it. And that's not going to be any good for anyone.

What we are doing is pretty straight forward. One is, first, have a set of core principles. Whether it's in security, whether it's in privacy or in AI safety, have concrete set of principles. But more than just the commitments we make, the more important thing is, what are the actual capabilities we are building to make progress on the commitments we made.

And in fact, even today, when you think about it, let's take security. You're deploying a new AI model. The first thing you want to really do with security is to test the adversarial attack. It's not just finding bugs, but it's more about something like prompt injection. What does it do to this model? And so, to be able to simulate the adversarial attack, that's an important consideration; or things that we're doing with confidential computing around privacy, that's another thing that one has to do in conjunction with the latest and greatest models.

Same thing around, for example, we know LLMs hallucinate. When you talk about AI safety, one of the things is to, in fact, use AI to measure groundedness of any output. These are things which are concrete capabilities that we are building in to our platform as we make progress, so that software developers have trust, and people who use products that software developers build using AI platforms have trust in the products that they're using. That's, I think, core to us.

And all of this leads back to where I started, which is ultimately, technology has to translate into real-world impact, one person and one organization at a time, right here in the United Kingdom, for us to be able to drive that economic surplus and growth that truly then improves the outcomes in this economy.

I couldn't be more excited about what I think all of you and your partners will do with all this innovation. And I look forward to coming back here and seeing some of that innovation in the years to come. Thank you all very, very much. Thank you.

(Applause.)

END