Build 2017
Satya Nadella, Chief Executive Officer, Microsoft
May 10, 2017

SATYA NADELLA: (Cheers, applause.) Good morning. Good morning. How are you all doing? (Crowd responds.) All right.

Welcome to Seattle. Welcome to Build 2017. Finally, spring has arrived in Seattle. They tell me that it's a beautiful, bright, sunny day outside. And so we thought it would be an awesome idea to bring you all into this dark room and talk about cloud computing. (Laughter.)

You know, for me, these developer conferences have been the way I have, in some sense, marked the passage of time. The progress that we as an industry make, we as a company have made, we as a set of developers in the ecosystem have made.

Recently, someone posted a presentation I did perhaps in '93. It was one of those classic code-on-screen developer conferences. This was pre-web. And we were so excited about Windows NT coming out and the birth of client-server, and I think the slide was all about connecting through an SNA server to a host machine, getting data into Excel, using DB automation.

And I was sort of listening to this and I was thinking, "Wow, you know, it's so cool to get excited about technology." And, in fact, my daughter saw this picture and said, "Who is that?" (Laughter.)

And, you know, a lot has changed. In fact, one of the developers, as I was prepping for it, said, "Hey, there seems to be a real inverse relationship between the amount of hair you have and the amount of computing power that we have now grown." (Laughter.) And then later he realized that he made a big mistake because what about the future? I've reached the limits of physics now, I don't have more hair to lose. (Laughter.)

So we truly, truly are in that post-Moore's Law and maybe GPUs and FPGAs or what have you that will have to come and help us.

It's stunning, though, to see the progress across the depth and breadth of our society and economy and how digital technology is so pervasive.

There was this one stat, because we were talking about the past a bit, in 1992 when I joined the company, the total amount of Internet traffic, this is pre-web, was around 100 gigabytes per day.

And now, per second, we have 17.5 million times that. In fact, 90 percent of all the data has been generated in the last two years. An autonomous car is going to generate

something like 100 gigabytes per second. And by 2020, it's estimated that there are going to be 25 billion intelligent devices.

So whether it's precision medicine or precision agriculture, whether it's digital media or the industrial internet, the opportunity for us as developers to have broad, deep impact on all parts of society and all parts of the economy has never been greater.

Now, but with this enormous opportunity, I believe comes enormous responsibility.

We talked about this even last year. But to me, the choices we as developers make are going to have pretty profound implications.

I'm an unrepentant tech optimist, there's no question of that. But also I'm grounded, though, that there are unintended consequences of technology.

And it's not that we can just use more technology to solve for those problems. And, definitely, technologists by themselves cannot solve for these. But I do believe that it is up to us to ensure that some of the more dystopian scenarios don't come true.

I mean, if you think about it, what Orwell prophesied in *1984*, where technology was being used to monitor, control, dictate; or what Huxley imagined we may do by just distracting ourselves without any meaning or purpose. Neither of these futures is something that we want.

So the question is: What are we going to do? What are the practical ways we can make progress? And, to me, it comes down to a set of design principles and choices we make. I believe we can make these practical design choices that help enshrine our timeless values.

The first one is that we should empower people with technology. That's got to be the front-and-center consideration for everything that we do going forward. People have to be in the loop. Let's amplify their capability, amplify their ingenuity.

Second, let us use technology to bring more empowerment to more people. When you have these amazing advances in computer vision or speech or text understanding, let us use that to bring more people to be able to use technology and to participate economically in our society.

So this inclusive design, ultimately, can be an instrument for inclusive growth.

Lastly, I think it's up to us as the world becomes more technology driven, digital technology is part and parcel of every aspect, building trust in technology is crucial. And I think it starts with us taking accountability, taking accountability for the algorithms we create, the experiences that we create, and ensuring that there is more trust in technology with each day.

These principles, these are practical things. But, yet, they capture, I think, the essence of the timeless value. And this is what grounds us in our mission, when we say we want to empower every person and every organization on the planet to achieve more, and I think it especially speaks to empowering every developer.

After all, we're a company that started as a dev tools company. We want to think about people, but we also want to think about the institutions people build. We want to think about what we do across the globe, every part of the world. And perhaps most importantly, it's not just celebrating any platform or technology of ours, but it's about celebrating what you can do with that technology to have broad impact.

And that's what we've been realizing. That mission, that sense of purpose is what has been driving us in this mobile-first, cloud-first world.

And the opportunity that you all have around Microsoft platforms is pretty stunning and amazing. Think about the 500 million Windows 10 devices that you can now reach through the Windows Store. That reach is what's going to drive our ecosystem going forward. It's going to give each one of you more of an audience, more users, more engagement in a secure way. And we're going to talk a lot about the advances in Windows Store throughout this conference.

100 million MAU of Office 365. This is the most valuable set of users in the core of the enterprise using these tools every day. And we've done a lot to make all of your applications more discoverable, more usable, inside the user experience scaffolding of Office.

I mean, take Teams. It's so amazing to see how we built Teams with that third-party extension being front and center, bots, tabs, extensions, all of those being surfaced within the user experience scaffolding of Teams. I think huge opportunity. Outlook extensions. All through Office.

So you have now the ability to reach 100 million of these most valuable users with any application you build.

Cortana is getting used by 140 million users every month on Windows today. I mean, just to put this in perspective, these are devices, by the way, that all have near-field audio and screens.

And so when you start writing skills that is going to take audio input or text input and have responses that are both audio and screens, you have an amazing install base already that's rich that you can extend.

In terms of organizations, there are over 12 million organizational entities in Azure Active Directory. So this is both large businesses, small and medium-sized businesses, educational institutions -- 12 million of them are already inside of Azure Active Directory. That means they're using one of the cloud services we have.

And that means your applications can do single-sign-on across all of these 12 million organizations, reducing friction, making it possible for you to reach them.

Over 90 percent of the Fortune 500 are using our cloud today. So this is all a foundation for the opportunity that you have in front of you as you build your applications and services.

And throughout these three days, we're going to talk about how we're continuing to advance these platforms, continuing to advance this opportunity so that you get the best return on the investments you make on Microsoft platforms.

Now, one of the things that always happens is you as developers with the apps you build and the feedback cycle between those applications and the platform is what shapes the evolution of platforms.

Last year, we talked a lot about conversations as a platform, and this year, you will see in demo after demo, people using conversations as a platform. People building agents, personal assistants that span devices. Skills for those agents, bots, natural user interfaces for these multi-device applications, whether it be text, speech, or even vision and gestures with mixed reality.

IoT, big, big, big workload in the cloud. AI is everywhere. The ability to reason over large amounts of data, create intelligence, distribute it.

Even the micro services, workflows, advanced analytics that people are building in the cloud are all pointing to what I think is a fundamental change in the paradigm of the apps that we are building, a change in the world view that we have. We're moving from what is today's mobile-first, cloud-first world to a new world that is going to be made up of an intelligent cloud and an intelligent edge.

And this entire conference is about talking about this massive shift that's going to play out in the years to come.

There are three fundamental characteristics of this new application pattern and world view.

First, is the user experience is getting distributed across devices. It's no longer just mobile first, in other words, it's not about one device, and app model for one device. The user experience itself is going to span all of your devices.

Easiest to understand it when you think about your personal digital assistant. Your personal data assistant, by definition, is going to be available on all your devices. And, in fact, as you move between devices, it's going to be there, helping you get your tasks done.

That multi-device experience is what now needs platform capability, you need new abstractions, not to just a single piece of hardware, but all of the devices that help you build out your application.

And, by the way, this is a user experience that is not just about input, it has also got stake in it.

And so that's what leads to, in fact, the second big change that is happening, it's AI. We know the currency of this platform shift is all about data. We talked about the car, right? 100 gigabytes of data per second for an autonomous car or connected car.

When you have that type of data being generated at the edge, data has gravity, computational power will move to it.

And so that means the AI you create is going to be, by definition, more distributed. You will, in fact, already see these cases where people are doing a lot of the training in GPUs on the cloud, but are deploying the models on the edge.

But we're reaching a point where you can't rendezvous all the data for training in the cloud, you actually will have to train and do your entrances in the edge, so you absolutely need a new set of abstractions for even AI in terms of how you learn and you train and you infer that span both the edge and the cloud.

And so as you see the experience as well as data and AI getting distributed, there needs to be a fundamental change in the inner and the outer loop of what we as developers do. Continuous integration, continuous deployment is going to change. And that's where things like micro services, containers are playing a massive role in the outer loop.

But one of the things that I think is going to completely change how we think about logic is serverless.

In a world where things get more distributed, event driven, you want to be able to write logic that reacts to these events, is not static. That means you don't have a piece of logic just being bound and running on one virtual machine in one location, but it's something that can actually be mobile.

I mean, think about even these AI models. One of the big challenges is going to be model drift. So you want to be able to tackle that. So serverless computation is going to fundamentally not only change the economics of what is back-end computing, but it's going to be the core of the future of distributed computing.

So this application paradigm shift of intelligent cloud and intelligent edge is going to be pervasively changing everything we do inside of Windows, inside of Office 365, inside of Azure. And the rest of what you will hear about in the next three days is about unpacking this.

So I think the best way to perhaps get into this would be to actually start seeing this in action. The first demo we want to show you is from a company, Sandvik, which is a Swedish industrial automation company, 150 years old, 43,000 people in the company.

They build industrial automation, industrial metal-cutting machinery that does high-precision work. They built a cloud application to be able to take all of the data from their machines, do predictions, preventive maintenance. And they realized that, in fact, they started using the time series capability. So that means you were able to see the anomalies as they happened, but with time series, the cool thing is that you could, in fact, go back and see what is the causal effect or what is the causation of the anomaly.

And so once you start seeing that, then they said, "Well, we should be able to make these decisions on when to shut the machine before it does any damage without having latency and bandwidth issues and so on." So they said, "Well, we need to distribute the logic that they've written in the cloud to the edge itself." And that's where they wrote Azure Functions, that's where the logic was written, and then they containerized it in Docker and moved it to the edge so that they could take those inference decisions and make them so much more faster.

And to show you this demo, I wanted to invite up on stage Sam George from our IoT team to sort of illustrate how this intelligent cloud and intelligent edge is already working at Sandvik. Sam? (Applause.)

SAM GEORGE: Thank you, Satya. Good morning, everyone.

The machine that you see costs over $1 million. It has over 900 moving parts. Any of those parts that break down could cause an interruption to that machine.

And this machine doesn't exist in isolation. It exists in factories. And there's many of these machines in a factory. And so down time to one machine can cause cascading failures to the factory and delays that cost millions of dollars per hour.

Now, years ago, these machines were simply serviced at regular intervals. As Satya just said, these are now being connected to the cloud so that you can predict the maintenance needs of these and to even shut them down in the case of an impending failure.

So Sandvik Coromant has these connected to Azure. And they trust Azure and Azure IoT to keep these machines running and to keep them safe. Let's take a look at what they're doing.

So here you see both the factory floor and the edge devices as well as the cloud.

Sandvik has software that's running on these machines and the software is collecting telemetry from these devices and sending them to IoT Hub, to our Azure IoT Hub service in a secure, scalable, two-way communication.

Once the telemetry is through IoT Hub, they're using a combination of Stream Analytics and machine learning to determine the maintenance needs of these machines and keep them running, and also to trigger a shutdown in the case of a pending damage to the machine.

And so when that happens, they simply trigger an Azure Function which calls through Azure IoT Hub and then down to the devices themselves.

Now, Sandvik, like many of our industrial customers, has also implemented failsafe logic that runs out on the machine. So if there's a local network interruption, they can still monitor whether the machine needs to be shut down to protect it, and they had to develop that again. They already had that logic running in the cloud, and what they really, really want is to be able to take that logic in the cloud and move it out to the edge.

Today, I'm thrilled to announce a breakthrough new capability we call Azure IoT Edge. Azure IoT Edge is a cross-platform runtime that runs on both Windows and Linux. It runs on devices even smaller than a Raspberry Pi. And it enables cloud functionality to be exported, run and managed directly on IoT devices remotely.

So now the picture starts to look dramatically different. Now Sandvik can simply take that logic and export it and run down on these devices. And what's nice about that is you get a nice symmetry where logic can be tested and developed in the cloud, you can run pipelines in the cloud, and on successful deployments, you can run the exact same code out on the devices at very low latency.

So to take a look at that in action, let's take a look at an operational dashboard.

So what you see here is a set of factories across the globe. And I'm going to open one of them up.

So these are all the different machines on this factory floor. Now, each one of these machines has its own unique telemetry that it's sending to the cloud, its own unique RPMs and power consumption and vibration. And all of that telemetry was used to create those machine learning modules that are fine tuned to predict the maintenance needs for each one of these machines, and to shut them down in the case of a pending failure.

So if I go over to the Azure tab here, I can see all of the different logic that Sandvik has developed for these machines. And the first three are for the metal-cutting machine that you saw when I came out.

But this new capability that we've developed called Azure IoT Edge and this capability in the cloud, we can now simply take that logic and export it. And it gets packaged up into portable Docker containers. And now when I go back to the factory floor, I get an alert on the machine that I just exported that there's new software available for it.

So if I open this up, I see the operational dashboard for this machine. And it's got all sorts of interesting information -- temperature, power consumption -- but there's a really important one I want to draw your attention to, and that's up in the upper-right-hand corner, the emergency response time.

And what that is is how long it takes for telemetry coming from these devices to be sent to the cloud, evaluated, and in the case of an emergency shutdown, for a signal to be sent back down to the device.

Right now, that's a little over two seconds. And with Azure IoT Edge, we can do dramatically better.

So these packages that I exported are now available. I'm going to open them up and select them. Now, I've already coordinated ahead of time with the IT and OT team so that I know that it's safe to update these machines.

And when I click "deploy," what's going to happen is we're going to network these containers down to the machine, we're going to activate them, and they're going to start running that protection code right on the machine. And you're going to see the emergency response time drop dramatically.

The reason why is that that logic is now running locally, so there's no cloud loop involved in it anymore.

So I go ahead and click deploy. It's now downloading, installing and running. And now you see we just drop down to a little over 100 milliseconds, that's a 20X improvement. And now Sandvik Coromant doesn't have to duplicate their logic in the cloud and on a device to protect these million-dollar machines.

This is the power of Azure IoT Edge and it's the power of an intelligent cloud working in harmony with an intelligent edge.

Thank you very much.

(Applause.)

SATYA NADELLA: Thank you, Sam. I mean, that shows you the simple, but powerful illustration of how you can take logic, distribute it to, in this case, improve response time.

Now, let's take it to the next logical step. In fact, let's take it to the extreme.

You know, when you think about the last 20 years or so, one of the most profound changes that we've seen is what happened with the web and something like search where all the text that was ever created could get indexed and searched, reasoned over. I mean, we're crawling all the time, we're able to understand all of the world's text, and then to serve it up knowing what's in it.

Just imagine if we can do that with any physical place. Suppose we can create these digital twins of a hospital, of an industrial setting, a factory floor. And fundamentally, you could start reasoning about people, their relationship with other people, the things in the place, all towards creating safety for human beings.

When you can start setting policy on what is safe interactions, that can absolutely change lives and make sure that some of the -- you know, in fact, there was a survey that was done by the National Council of Occupational Safety and Health, and basically said that pretty much all of the accidents in the workplace can, in fact, be prevented if you were able to detect these anomalies before they happen.

So that's the kind of impact some of the technology that we have now can have.

And so we've put together a demo for you using pretty much the runtimes that are all available to us. The new Azure IoT Edge, when you have lots and lots of devices in a place and you're creating a mesh. In fact, an ad-hoc data center gets created pretty much nowadays on a factory floor, at home, or anywhere else.

And for that, we even have Azure Stack as the edge of the cloud so that you have the computation and storage capability, you have the ability to run these Cognitive Services, things like computer vision on a commodity camera.

Bringing all of this edge compute together and edge intelligence together with the cloud, you can turn any place into a safe, AI-driven place.

What we wanted to do is roll a video to show you some of the scenarios, and then Andrea will come out and give you a demo of this AI-driven workplace for safety. Let's roll the video first.

(Video segment.)

ANDREA CARL: Good morning, everybody. So let's take a look at how you can digitize real-world environments like what you just saw in the video.

As Satya said, today we're going to show you a solution that we've built using Microsoft Azure Stack, Azure Functions, Cognitive Services, and commodity cameras that make this a reality.

In a highly controlled environment like a construction site, being able to track equipment, people, and what they're doing is critical to workplace safety.

So the first thing I want to show you is over to your right where we have a construction yard, and you're going to see Jana is working.

Using our technology right now, both people and object recognition models are currently deployed in that construction yard. And then on the screen there before you, you can see the solution is recognizing Jana in real time, as well as equipment like the green jackhammer on the floor there.

Now, over next door to your left is a workshop space, which is where Ayuba is working. And right now with our technology only people recognition is deployed in that space.

So we're going to start by searching for an object using this phone. Let's say I want to find a jackhammer. I'm going to type "where is a jackhammer." If I could spell it, I will type it. There we go. Now, you can see even though there is actually a jackhammer resting against the bench in Ayuba's workshop, the system doesn't see it, but it does find the one next door, which again is where we have both people and object recognition deployed.

The system is going to let me do things like find who's trained to use that jackhammer, and also who last used it. And so that means I can actually use this technology to ensure my project is meeting health and safety regulations.

And so in order to do that I'm going to deploy an object recognition model in the workshop.

But before I do, take note again of that jackhammer in Ayuba's workshop. It is resting vertically against the bench, which is actually a very unsafe position for such a large, heavy object to be in.

And so I've already gone ahead and I've tagged that jackhammer in a variety of safe and unsafe positions in the system. And so when I update this pipeline, which I'm going to do right now, it's not only going to see the jackhammer now, but it is also going to tell me it's in an unstable orientation, and it's going to notify Ayuba, so he's able to very quickly resolve the issue.

The solution is running more than 27 million recognitions per second across people, objects and activities, both in the yard and in the workshop. And we're making changes to the solution in real time in the cloud, and deploying it to the edge, which is a great example of how we can stretch the cloud to create a mesh of interconnected devices and services.

So now imagine if we could take that a step further, recognizing people and then applying a set of policies to them to ensure workspace safety.

Over in the construction yard we have a new employee, Kristina, who is showing up for her first day of work. So I'm going to go ahead and onboard her into the system by using this phone again and typing a simple command, show me non-employees on the jobsite.

Now, because this is a highly controlled environment, and health and safety is so imperative, there are cameras already deployed on that jobsite. And so they've already captured her photo, and that makes it very easy to onboard her. I can just click on her picture, new person, and go ahead and enter her information. OK, so now she's onboarded.

And now I can create a new policy to enforce that credentialed employees can do things like use jackhammers. So I'm going to type a policy. Only credentialed employees can use jackhammers. OK, the new policy has been activated. And so now she can just go about her job, and I won't get any policy violations, which means our regulations are being met.

Now, if somebody who is not authorized to use that equipment comes onto the construction yard, a workflow is going to notify me and people nearby of a potentially dangerous situation, like that one, for example. You can see the violation's been triggered, and Kristina is able to quickly resolve it.

Now, this was an intentionally simple scenario, but examples like this one and the ones you saw in the video are all based on real customer requirements, and they support even more complex business processes using Cognitive Services, Azure Functions, Azure Stack, and commodity cameras to bring the edge of Microsoft's cloud to any device.

The intelligent edge is the interface between the computer and the real world. And so we have just shown you how you can search the real world for people, objects and activities, and apply policies to them to improve health and safety.

It's early days, but we are really excited for the potential of AI for workplace safety. And with this technology we think you have an opportunity to build even more sophisticated solutions that have the power to transform entire industries. And we'll share more over the next 12 months.

Thank you.

(Applause.)

SATYA NADELLA: Thank you, Andrea.

That was a pretty amazing example of what this new architecture, the ability to distribute intelligence in the cloud, as well as in the edge, can do to deliver something like AI for safety.

But now let's move to end user computing, right? So we've seen a couple of amazing examples of what's happening with these industrial applications, but what's happening with the edge is for sure happening with the devices we use every day.

In fact, just the social complexity, Bill Buxton, who is a UI researcher, always likes to remind me that the fundamental challenge we have as an industry right now is the social complexity of our devices.

I mean, just think about the notifications. How many times do I need to get notified across all my devices?

In order to solve for that you need to be able to have that intelligence that spans all of your devices, you need to have these agents that span all of your devices.

And so you need a new architecture for everything from notifications to agents to task completion. And that's really what we have set out to build. In fact, the journey got started with Office 365. Office 365 was pretty profound in the sense that we took workloads like Exchange or Skype or SharePoint to the cloud.

But beyond the move of an individual workload to the cloud, the more profound shift that happened was underneath. It was the data plane. The people and their relationship with other people is now a first-class thing in the cloud. It's not just people but it's their relationships, it's their relationship to all of the work artifacts, their schedules, their project plans, their documents; all of that now is manifest in this Microsoft Graph.

And with the move to Windows as a service, you even have all of the devices that people use moving to the cloud itself.

So you now have this very rich data graph of people, their activities, their devices, all principled under use security principles and organizational security principles. So when we talked about that Azure Active Directory and the Microsoft Account, those security principles still govern the access to all of this data. So it's really the user's data.

But now there is a new platform in Microsoft Graph that allows every developer as they start constructing these multi-device experiences to be able to access people, their relationships, their activities, as well as their devices.

And you'll see this even tomorrow in the Windows presentation as to how this programming model is going to manifest to the core of what is Windows applications going forward.

And so if you take a scenario like meetings, most people think about meetings as a one-time thing, but the reality is people come together, collaborate before a meeting to prepare for the meeting, and then during the meeting, and then want to be able to access the artifacts from the meeting after the meeting, right? So there is a temporal part to a meeting, there is multiple people, multiple devices.

And when you think about building an application or an experience for meetings, thinking about it in simple terms of a single device experience or a single cloud service is

just not sufficient. You need new programming model abstractions that allow you to think about everything that spans an intelligent meeting.

And to show you this in action, I wanted to invite up onstage Laura Jones and team to show you how intelligent meetings can be built using everything that we have in the Microsoft Graph and Office 365 and Azure, as well as in Windows and other devices.

Laura, take it away.

(Applause.)

LAURA JONES: Today we're going to show you three things that demonstrate our vision for the future of productivity. First, how Microsoft Services work across context and devices. Second, how developers can leverage the Microsoft Graph to build device intelligence right into their apps. And, third, how we're bringing together our productivity platform to optimize the meeting experience.

Hey Cortana, what's my day look like?

CORTANA: I found three events today. First up, at 9:00 a.m. you have a meeting titled Packaging Review. It looks like you're low on fuel, so I suggest leaving 10 minutes early to make your 9:00 a.m. meeting on time.

Also, you've got an update from Expedia confirming flights and a hotel room.

LAURA JONES: Right, that's the upcoming vacation I have, so I'd better let my boss know.

Hey Cortana, ask Time Away to set my out of office from Monday to Friday.

CORTANA: OK, you're marked as out of office from Monday to Friday. Your outgoing reply message is sent and your time off has been entered in the Time Away system.

LAURA JONES: Awesome.

Developers can create Cortana skills that work across platform, like this example with the Invoke speaker where Cortana is working directly with my work time off system to fulfill my request.

Hey Cortana, do I have any other reminders for today?

CORTANA: You mentioned you would provide a list of designers for review in today's meeting.

LAURA JONES: Right. Remind me when I get to the car.

CORTANA: Got it.

LAURA JONES: With the Device Graph I'm able to push this notification to any of my devices. So I could push it to my Xbox, but since I'm headed that way I've gone ahead and sent it to my car. Don't laugh, it gets great fuel economy. So with the help of Cortana I'm fueled up and ready to make my meeting on time.

CORTANA: There's an accident on 520 that will make you late for your 9:00 a.m. meeting. Would you like me to let the other attendees know and connect you when it starts?

LAURA JONES: Yes, please.

CORTANA: OK, I'll let you know.

LAURA JONES: I guess I spoke too soon.

CORTANA: Hi, Laura, just a reminder to provide a list of designers for your packaging review meeting.

LAURA JONES: Ask LinkedIn to pull all the designers from my network.

CORTANA: I found all of the designers in your network. Would you like me to read out the top five or send you a summary?

LAURA JONES: Send a summary to my product review channel.

CORTANA: OK. I added that to your team channel.

LAURA JONES: Because Cortana is aware of the device I'm on, she's able to provide me with contextually appropriate response so I can take the best next steps.

CORTANA: Hi, Laura, your meeting is about to start. Should I connect you?

LAURA JONES: Yes, please.

Hey folks, I'm running a little late but go ahead and get started without me.

RUCHIR ASTAVANS: No worries, Laura, we got your note earlier.

So while Laura joins let me take a second and show you guys around Microsoft Teams. For those of you that haven't seen it, this is Microsoft Teams, the new chat-based collaborative workspace in Office 365. And we're in a Teams meeting.

Now one of the great features in Teams is Tab; using Tabs I can take anything that I access most frequently, whether it's documents, any third-party content, or like this Power BI tab here which has our survey results, and I can pin it in my channels. And from there I can easily screen share it straight into my meeting.

LAURA JONES: Now one of the most disruptive things that can happen in a meeting is when someone enters the room late or enters the call late. Now what if I could simply just walk into the meeting and just like that the Teams app has turned into companion mode because it's aware of my context.

Windows is adding two new components to the Microsoft Graph, activities like this meeting that I'm in or a document that I might be working on throughout the day, and devices like this iOS device. Here we're using Project Roam to adapt my app based on the fact that I've walked into a meeting. Developers can leverage these APIs through the Project Roam SDK to bridge the gap between people, apps, devices and services, and enable truly mobile experiences.

RUCHIR ASTAVANS: Hi, Laura.

LAURA JONES: Hi, Ruchir.

Hi, Fatima. How are things going with the focus group data?

FATIMA KARDAR: Hey, I looked at the data last night. The demographic there was (inaudible) prefer blue than green.

LAURA JONES: OK, good to know. Let's go ahead and pull up the LinkedIn list I sent from the car.

RUCHIR ASTAVANS: Sure, let me share the LinkedIn tab.

LAURA JONES: I worked with Joe before on a similar project. Ruchir, would you mind pulling up some of his current work?

RUCHIR ASTAVANS: Sure. Now in my Teams meeting I can also chat with my channels bot. And developers can use the Microsoft Bot Framework to light up some amazing conversational experiences in Teams meetings. Let me show you one. Adobe Creative Cloud Bot, show me Joe Manzo's (ph) recent projects.

FATIMA KARDAR: That looks cool. Laura, I think that you worked with Joe before. Do you mind making the connection for us?

LAURA JONES: Yeah, sure thing, great.

FATIMA KARDAR: Meetings Bot assign task to Laura to connect with Joe.

LAURA JONES: Great. Now that that's done I wanted to show you guys the latest comps for the packaging. With the Teams companion app I'm able to share video and photo straight into the meeting. Now I think it's a little retro, but let me know what you guys think.

FATIMA KARDAR: Yeah, that's retro.

LAURA JONES: Right, well, hopefully Joe can bring it up to date. Thanks, guys.

RUCHIR ASTAVANS: Thanks, Fatima.

LAURA JONES: So when I get back to my desk and open up my laptop Cortana is waiting there for me with an action item from earlier in the day. When I click on the notification it takes me to Outlook where there's an email with my meeting summary. It has the attendee list, the video links, the tabs use and the bot interaction. Best of all, here's my action items. I could go ahead and mark it as complete right in the email. But I want to dig in a bit more. Developers could add content to this actionable email, or create Outlook plug-ins like this one from Microsoft Stream.

When I click on the plug-in the Stream service gives me a transcription of the meeting. I can see key points, when they were made, and who made them. This way I can get the context that I need to complete my action item. And, of course, the Stream service uses the publicly available Microsoft Translation APIs, so I can get this in the language of my choice.

So what you've seen us show you today are Teams and Cortana, two services leveraging the intelligent edge and Microsoft Graph. With hero experiences like these you're able to turn any device into a productivity device. So I can move fluidly throughout my day no matter where I am. Developers can leverage the Microsoft Graph to create systems of intelligence that will fundamentally change the future and culture of work.

Thank you. (Applause.)

SATYA NADELLA: Thank you, Laura.

So it's a pretty amazing world we can create using this new platform of intelligent cloud and intelligent edge. Hopefully the three examples give you a feel for the depth and breadth of apps you could build. You can build these intelligence cloud apps. You can distribute them using these intelligent edge runtimes. You can have the operational complexity of all of that distributed computing tamed because of this new technologies like serverless compute with support for micro-services and containers. You can extend the Microsoft Graph. You can consume the Microsoft Graph and reason over it; build intelligence into every experience that spans devices. So it's a pretty awesome set of capabilities that these platforms provide and you'll see a lot more of this throughout the rest of the conference.

But I want to close where I started. It's this opportunity and responsibility that we have in front of us. The future of computing is not going to be defined by any runtime or platform technology. It is going to be defined by the choices that you as developers make and the impact of those choices in the world.

And I'm always inspired when I listen to these stories of an individual developer who developed this deep sense of empathy for a scenario and other human beings and then goes after it with this most malleable of resources that we have, software, and changes the world.

Let's roll the video of one such developer.

(Video segment, applause.)

SATYA NADELLA: Ladies and gentlemen, Emma and Haiyan. Thank you so much Emma and Haiyan for being here, sharing your story and inspiring all of us, and showing I believe what that opportunity and responsibility that we all have as developers to build that common ground with our fellow human beings, to have that impact.

Thank you so much again for being here, inspiring us, and I hope that this is something that we as developers can take forward with the choices we make.

Thank you all very, very much. And thank you again.

END