

Google I/O 2016 Keynote Full Transcript

Here is the full transcript of the Google I/O 2016 conference keynote – the company’s yearly developer conference held at Shoreline Amphitheater in Mountain View on May 18, 2016.

Speakers:

Sundar Pichai – CEO, Google

Mario Queiroz – Vice President of Product Management, Google

Erik Kay – Engineering Director at Google

Rebecca Michael – Head of Marketing, Communication Products at Google

Dave Burke – VP of Engineering, Android

Clay Bavor – VP, Virtual Reality at Google

David Singleton – Director, Android Wear

Jason Titus – VP, Developer Products Group at Google

Stephanie Cuthbertson – Group Product Manager, Android Studio

Ellie Powers – Product Manager of Google Play

YouTube Video:

Sundar Pichai – CEO, Google

Welcome! Welcome to Google I/O and welcome to Shoreline. It feels really nice and different up here. We’ve been doing it for many, many years in Moscone, and in fact, we’ve been doing I/O for 10 years but I feel we are at a pivotal

moment in terms of where we are going as a company and felt it appropriate to change the venue.

Doing it here also allows us to include a lot more of you. There are over 7,000 of you joining in person today. And later today, after the keynote, you'll be joined by several Googlers, product managers, engineers, and designers, so hopefully you'll engage in many, many conversations over the three days.

As always, I/O is being live streamed around the world. This year we have the largest-ever audience. We are live streaming this to 530 external events in over a hundred countries around the world, including Dublin, which is a major tech hub in Europe, Istanbul, which is our oldest Google developer group, and even to Colombo, Sri Lanka, which is the largest attendance outside of the US with 2,000 people.

Our largest developer audience on the live stream is from China today with over 1 million people tuning in live from China, so welcome to those users as well.

We live in very, very exciting times. Computing has had an amazing evolution. Stepping back, Larry and Sergey founded Google 17 years ago with the goal of helping users find the information they need. At the time, there were only 300 million people on line. Most of them were on big physical computers, on slow internet connections.

Fast-forward to today, thanks to the rate at which processors and sensors have evolved, it is truly the moment of mobile. There are over 3 billion people connected and they are using the internet in ways we have never seen before. They live on their phones. They use it to communicate, learn new things, gain knowledge, entertain themselves. They tap an icon, expect a car to show up. They talk to their phones and even expect music to play in the living room, or sometimes groceries to show up at the front door.

So we are pushing ourselves really hard so that Google is evolving and staying a step ahead of our users. All the queries you see behind me are live queries coming in from mobile. In fact, today, over 50% of our queries come from mobile phones. And the queries in color you see behind me are voice queries. In the US, on our mobile app in Android, one in five queries – 20% of our queries – are voice queries and that share is growing.

Given how differently users are engaging with us, we want to push ourselves and deliver them information, rich information, in the context of mobile. This is why, if you come to Google today and search for Beyoncé, you don't just get ten blue links. You get a rich information card with music. You can listen to her songs, find information about upcoming show events, and book it right there.

You can come and ask us different queries. Presidential elections or Champions League, and we again give you rich in-depth information. And we do this across thousands and thousands of categories globally at scale. You can

come to Google looking for news as well. For example, if you're interested in Hyperloop, an exciting technology, we give you information with amp right there in search results and they load instantly and you can scroll through them.

Amazing to see how people engage differently with Google. It's not just enough to give them links. We really need to help them get things done in the real world. This is why we are evolving search to be much more assistive. We've been laying the foundation for this for many, many years through investments in deep areas of computer science. We built the Knowledge Graph. We today have an understanding of 1 billion entities: people, places, and things and the relationships between them in the real world.

We have dramatically improved the quality of our voice recognition. We recently started training our data sets with noisy backgrounds deliberately, so that we can hear people more accurately. The quality has improved recently by 25%.

Image recognition and computer vision. We can do things which we never thought we could do before. If you're in Google Photos today, and you search for hugs, we actually pull all the pictures of people hugging in your personal collection. We have recently extended this to videos, so you can say "*show me my dog videos*" and we actually go through your videos and pull out your favorite videos.

Translation. 10 years ago, we could translate, machine-translate in two languages. Today we do that for over a hundred languages and every single day we translate over 140 billion words for our users. We even do real-time visual translation. If you're a Chinese user and you run into a menu in English, all you need to do is to hold up your phone and we can translate it into English for you.

Progress in all of these areas is accelerating, thanks to profound advances in machine learning and AI, and I believe we are at a seminal moment. We as Google have evolved significantly over the past 10 years and we believe we are poised to take a big leap forward in the next 10 years.

So leveraging our state-of-the-art capabilities in machine learning and AI, we truly want to take the next step in being more assistive for our users. So today, we are announcing the Google Assistant.

So what do we mean when we say the *Google Assistant*? We want to be there for our users asking them, "*Hi, how can I help?*" We think of the assistant in a very specific way. We think of it as a conversational assistant. We want users to have an ongoing, two-way dialogue with Google. We want you to help get things done in your real world and we want to do it for you, understanding your context, giving you control of it. We think of this as building each user their own individual Google.

We already have elements of the assistant working hard for our users. I mentioned earlier that 20% of queries on our mobile app in Android in the US

are voice queries. Every single day, people say, *"Okay, Google..."* and ask us questions that we help them with and we have started becoming truly conversational because of our strengths in natural language processing. For example, you can be in front of this structure in Chicago and ask Google, *"Who Designed This?"* you don't need to say *"the bean"* or *"the cloud gate."* We understand your context and we answer that the designer is Anish Kapoor.

Here's another example. You can ask google, *"Who directed the Revenant?"*

[Google: the Relevant was directed by Alejandro Iñárritu]

And you can follow up that with a question, *show me his awards*. Notice that I didn't say the name, which I am glad because I find that name very, very hard to pronounce. And Google could pick that conversation up and return the right answer. This has historically been really hard to do for computers. The reason we are able to do is because we have invested the last decade in building the world's best natural language processing technology. In our ability to do conversational understanding is far ahead of what other assistants can do. Especially if you look at follow-on queries, our studies show that we are an order of magnitude ahead of everyone else.

So today people are using Google and asking us questions in many, many different ways. So we've put together a short video so that you can take a look.

[Video Presentation]

As you can see, users are already looking to Google to get – to help them get things done, but we believe we are just getting started. We believe this is a long journey and given it's a journey, we want to talk to you a little bit about the future. We want to show you the kind of things we aspire to be able to. Let me do that with an example.

Here's a common situation. It's a Friday night. I'm sure many of you can relate to it. Back home, and I want to take my family to a movie. You know, you normally pull out your phone, research movies, look at the reviews, find shows nearby, and try to book a ticket. We want to be there in these moments helping you.

So you should be able to ask Google, *"What's playing tonight?"* and by the way, today, if you ask that question, we do return movie results, but we want to go a step further. We want to understand your context and maybe suggest three relevant movies which you would like nearby. I should be able to look at it and maybe tell Google, *"We want to bring the kids this time."* and then if that's the case, Google should refine the answer and suggest family-friendly options. And maybe even ask me, *"Would you like four tickets to any of these?"* And if I say, *"Sure, let's do Jungle Book,"* it should go ahead and get the tickets and have them ready waiting for me when I need it.

As you can see, I engaged in a conversation with Google and it helped me get things done in my context. And by the way, this is just one version of the

conversation. This could have gone many, many, different ways. For example, when Google returned the results, I could have asked, “*Is jungle book any good?*” And Google could have given me the reviews and maybe even shown me a trailer. And by the way, I saw the movie, it’s terrific. And hope you get to see it as well.

Every single conversation is different. Every single context is different. And we are working hard to do this for billions of conversations, for billions of users around the world, for everyone. We think of the assistant as an ambient experience that extends across devices. I think computing is poised to evolve beyond just phones. It will be in the context of a user’s daily life. It will be on their phones, devices they wear, in their cars, and even in their living rooms. For example, if you’re in one of the hundred different android auto models and you’re driving and you say, “*Let’s have curry tonight,*” we know the Warriors are on tonight and Steph Curry is playing but you know, all you’re looking for is food, and we should be smart, order that food and let you know when it is ready, and maybe even have it waiting for you at your home.

Talking about your home, we’ve already built many, many products for your home. Today, we have sold over 25 million Chromecast devices. So we’ve been thinking hard about how to bring this vision of Google Assistant into your home. Credit to the team at Amazon for creating a lot of excitement in this space, we’ve been thinking about our own unique approach and we are getting ready to launch something later this year. To give you a preview, I’m going to invite Mario from the Chromecast team.

Mario Queiroz – Vice President of Product Management, Google

Thanks, Sundar. Today we want to give you an early preview of how we’re bringing the Google Assistant to the home. Our aspiration is to make the Assistant useful and enjoyable in one of the most important places in your world: where you spend time with your family.

When I walk into my house, I want to be able to continue to have access to the Google Assistant, but I should be able to interact with it in a hands-free way, simply using my voice without having to take out my phone. This is why we’re creating Google Home, a device which will be available later this year.

Google Home lets you enjoy music and entertainment throughout your entire house, manage everyday tasks more easily, and ask Google what you want to know. All of this will be done by speaking with the Assistant. It will let anyone in the family, kids or adults, have a conversation with Google.

Google Home is unmatched in far-field voice recognition since it’s powered by more than 10 years of innovation in natural language processing. I can continue the two-way dialogue with the Assistant that Sundar mentioned earlier, whether I’m standing nearby cooking dinner or sitting across the room playing a game with my daughter.

With Google Home, we set out to create and design a beautiful product that's warm and inviting and fits naturally in many areas of the home. We're designing the top to blend into your home. We'll give you the option to customize the base with different colors and finishes, including metal and fabric to reflect your home's style, whether it be in the living room, the kitchen, or the bedroom. We're putting a lot of craftsmanship into the hardware design, including the LED's placement and choreography, the speaker that's going to fill the room you are in with music, the microphone system, and the clean face without any buttons.

This is Google Home! We think it will be a beautiful addition to any room in your house. And we're even more excited about what it's going to do for you.

First, music and entertainment are a big part of what makes being at home relaxing and fun. Not long ago, we introduced Chromecast, designed to play your favorite shows, movies, and YouTube videos on the biggest screen in your house. Last year, we added Chromecast Audio to bring the music you love to your best speakers.

Chromecast has been one of the hottest-selling consumer electronics products since the day it launched, and Google Home will build on those experiences. Google Home is a Wi-Fi speaker that streams music directly from the cloud, so you get the highest quality playback. It will deliver rich bass and clear highs all from a beautiful, compact form factor. Of course, you can access songs, playlists, albums, artists, and podcasts from your favorite music services just by asking with your voice. Or if you prefer, you can send music from your Android or iOS device through Google Cast. And unlike other home assistants, Google Cast support allows you to control other speakers in your home, without complicated setup.

So when you want to listen to Coldplay in the living room speakers, you can simply say, "play Viva la Vida " in the living room and it will start playing. Cast support also enables multi-room playback so that you can add one or more Google Home devices to a group of speakers and really blast your favorite tunes. And it lets you control your video content, too. Let's say that you want to watch that episode of Jimmy Kimmel or the trending YouTube video on your TV, just tell Google Home, and the content will appear on the biggest, brightest screen in your house.

Next, Google Home will become more and more of a control center for your whole home. Home is where lots of daily tasks just need to get done. Having access to the Google Assistant makes this a lot easier. It's like having a voice-activated remote control to the real world whenever you need it. You can do the basics like setting alarms and timers or managing to-do lists and shopping lists. We're also designing Google Home to connect your smart home seamlessly. It will support the most popular home networking systems so that you can easily control your lights, thermostats, switches, and more, including our own Nest devices. Further in the future, we'll work with developers to make it possible to control things beyond the home, like booking a car, ordering dinner, or sending flowers to mom, and much, much more, all with just your voice.

Third, Google Home lets you ask Google about anything you want to know. Of course, you can get the basics like the weather or facts that you might find on Wikipedia. But what makes Google Home really shine is that it has search built in. It draws on 17 years of innovation in organizing the world's information to answer questions which are difficult for other assistants to handle. You might ask: how much fat is in an avocado? Or, what is Draymond Green's jersey number? And then follow-up that last question with where did he go to college? Or try something more complex. What was the US population when NASA was established? You'll get immediate, accurate answers from Google each time.

And the Google Assistant not only knows a lot about the world, but it will stand apart in how it can also get to know you over time. With your permission, of course. It can help you retrieve your travel itinerary, your daily schedule, your traffic to work, your package delivery information, and much more. And as Google keeps getting better, so will Google Home.

So that's Google Home. A beautiful, smart, voice-enabled assistant for the whole family. Enjoy music and entertainment throughout your entire house, manage everyday tasks effortlessly, and ask Google what you want to know. It's early days, but we want to give you and show you how we envision the Google Assistant coming to life at home. We created a short video to bring the product into a family setting to capture what it might be like in the future to have your personal Google around the house. Let's roll the video.

[Video Presentation]

We're really, really excited about what's ahead. Google Home will be available later this year. In the meantime, to stay up-to-date with the latest news, please sign up at google.com/home. We wanted to share this early preview at I/O so that we can work with partners in a more open way to deliver awesome experiences at launch. We'll have a lot more to share soon with the developer community about how to begin to integrate with the Google Assistant.

And with that, let me ask Sundar back on stage.

Sundar Pichai – CEO, Google

Thank you, Mario. It's really exciting to see the Google Assistant come to life with Google Home and to help people get things done. To do this well, as Mario mentioned, we really need to work with developers and third parties so that we can provide these actions for our users. We already do this a lot at Google today. Using Google products, you can already book a movie ticket with Fandango, get a car with Uber, listen to music on Spotify, book a restaurant with OpenTable, maybe record a ride with strava, and many, many more such use cases. So we are thinking about this thoughtfully, and we're working on a comprehensive way by which third-party developers can interact with the Assistant, and we'll be sharing a lot more in the upcoming months.

I talked earlier that we are working hard at core use cases on mobile. One

such important use case is Photos. Last year, thanks to the advances in computer vision, we approached Photos with a new perspective. We announced Google Photo at Google I/O last year with the goal to help users find and organize their photos and videos, keep them safe, and share them effortlessly with their family and friends on any device. We've seen tremendous adoption with Google Photos, and just in the past year, we have seen tremendous adoption, and today we are at over 200 million monthly active users.

Our computer vision systems have automatically applied over 2 trillion labels. This is what allows us, when you search for Pomeranian, to find the right picture. And by the way, over 24 billion of those labels are for selfies. We even have Pomeranian selfies. Google Photos shows what's possible when you approach an existing area from a new perspective.

Another core use case on users' phones is communications. It's an exciting area, and there is a lot of innovation. But given our advancements in machine learning, we wanted to approach this core use case with a new perspective. Erik Kay is going to join to talk to you more about it.

Erik Kay – Engineering Director at Google

Thanks, Sundar. You know, communications is all about sharing life's moments. That great restaurant I found, the winning shot in overtime, my daughter's recital. Today I can share moments like these the second they happen, just by pulling out my phone and sending a message. Communications is such an important part of our lives, and it's an important focus for Google as well.

What makes me personally excited about communications is the potential for innovation when you combine the power of mobile with advancements in machine learning. So today we're giving you a look at what we've been up to with two new communication apps that show what's possible when we bring Google technology to this essential human activity.

The first is a new messaging app called Allo. Thank you.

Allo is a smart messaging app. It learns over time to make conversations easier, more expressive, and more productive by bringing all the richness of Google right into your chat. Allo is based on your phone number, so you can easily get in touch with anyone in your phone book.

This morning, we're going to walk you through three areas that make Allo really special. First are some ways to express yourself and keep the conversation going. Then, we'll talk about what it means to have the Google Assistant built right into your messaging app. And, finally, we'll tell you how Allo keeps your conversations private and secure.

So let's get started. On the stage with me to help demo is Amit. He leads product management. And on the screen behind me you can see Amit's in a chat – in a conversation with his friend Joy. So we designed Allo to help you express yourself and keep the conversation going. So, of course, there's a

great selection of stickers which we've sourced from independent artists and content producers from around the world. But expression is more than just emojis and stickers. So we've added more features that help you say what you mean.

First, let's look at whisper shout. We wanted to give you a way to add more emotion to your words. Sometimes you want to get your point across big and bold. Other times you want to say things a bit more softly. Whisper shout lets you express how you really feel by making your replies very big or very small.

Amit, can you show everyone how that works? So Amit has typed "yay" and throws a smiley face in there. Now watch, rather than tapping the *send* button, he slides it down to whisper and slides it up again to shout. Down to whisper and up again to shout. I think it's pretty cool, too. No more yelling in all caps to get your point across.

It's also fun to add emotion to your photos. And Ink lets you get creative with the photos you send. Amit's picked a photo of his adorable baby girl and wrote *ahoy* on the little sailor. Posted it. It's that simple. I use Ink all the time. It's really fun.

Another way Allo helps you express yourself is by letting you type less, a lot less. We've taken a page from the inbox playbook and built Smart Reply right into your chat conversations. This is especially powerful in messaging when you're on the go. So now you can keep your conversations going with just a quick tap.

So you can see when Joy asks: *dinner later?* That Amit has offered two Smart Reply suggestions: I'm busy and I'm in! Okay. I promise I won't act out any more emojis. So Allo uses machine learning to suggest replies on the fly, anticipating what you might want to say next. Now, these aren't just canned replies. Allo learns over time and will suggest responses based on how you like to express yourself. The more you use Allo the better the suggestions will become. So the suggestions you see are unique to you. You can say the things you want without having to type a single word. And since messaging isn't just about texts, smart replies contain stickers and emojis, too, because as they say, an emoji's worth a thousand words. Do they say that?

Now I want to show you something really cool. Allo even offers smart replies when people send photos to you. This works because in addition to understanding text, Allo builds on Google's computer vision capabilities to understand of the content and the context of images. In this case, Allo understood that the picture was of a dog, that it was a cute dog, and even the breed of the dog. In our internal testing, we found that Allo is 90% accurate in determining whether a dog deserves the "*cute dog*" response.

So let's try something even harder. When Joy sends a photo of pasta, we're able to identify the precise details of the image and create smart replies mentioning both linguini and the clams. What's really cool here is that we don't just identify what's in the image. Smart Reply actually creates a

conversational response, something like “yummy,” the kind of thing you would actually say in response to a photo of food. This is only possible because we’ve married our strengths in computer vision and natural language processing. And if you think about it, there’s a lot of complex technology at work here just to help you say something as fun – as simple and natural as “*I love linguini.*”

So, that’s a little how Allo helps you express yourself and keep your conversations going. The intelligence behind Smart Reply also gives you a taste at how assistive technology can make your message experience simpler and more productive. The Google Assistant built right into Allo takes it even farther, so I’m pleased to introduce one of our leads, Rebecca, to tell you more about the Assistant in Allo.

Rebecca Michael – Head of Marketing, Communication Products at Google

Thanks, Erik. As you heard earlier, the Google Assistant is an ongoing dialogue between you and Google that helps you get things done in your world. It’s also designed as an ambient experience. It’s there for you whenever you need it. And in messaging that really means bringing the Google Assistant right into your conversation with friends. So I’m going to show you how the Assistant can help in Amit’s and Joy’s conversation.

So they’re planning a dinner and Joy now says she would like Italian food. The Assistant intelligently recognizes that they could use some tips for Italian restaurants nearby and you can see its proactive suggestions at the bottom of the screen there. Tapping this brings up restaurant cards that everyone in the chat can see. These are powered by Google’s Knowledge Graph which means that Allo can help with all kinds of information in the real world. So there’s some back and forth about which restaurant to go to. And it looks like they’re leaning towards Cucina at 7 o’clock.

So Amit taps on the restaurant card to bring up more details. Looks good. Serves pasta which we know Joy fancied. So Amit can now go ahead and make a reservation right there through OpenTable. The Assistant prompts him to confirm the number of diners and what time they want to eat. And now they’re booked and ready to go.

So what we’re seeing here – what we’re seeing here is completely new. In the past, Amit would have had to leave the chat to do a Google search, return with some restaurant options, switch back again to share the options, go out again to make the reservation at OpenTable and then come back in to share the details with the rest of the group.

In Allo, Amit and Joy can choose and reserve a restaurant right there in the chat and in a natural and seamless way.

So OpenTable will be just one of the many partners that will be part of the Google Assistant ecosystem, and we can’t wait to share more soon with the developer community about how to get started.

Cool. So we saw some proactive suggestions from Google there. You can also call on the Google Assistant directly at any time just by typing @Google as Amit is doing now. And he's going to ask for *funny cat pics*. Really, Amit? Okay. Cool. So Google obliges, of course, with a lovely lineup of cats from Google Image search. Wow, that chin.

Okay. So you just saw how the Google Assistant can be really helpful in groups. You can also have a one-on-one chat with Google. What we're seeing now is Amit's contact list and Google's appearing at the top there. So let's jump in and have a chat.

Just like with any other conversation, this one picks up right where you left off and the Assistant will remember things like your name and even tell you how it's feeling. So let's try something more interesting. Amit's a big Real Madrid fan and he wants to know how they got on in their last match. So he asks the Assistant: did my team win? It looks like they did. They won their – yeah. Some Real Madrid fans out there. Cool. And so they won their last match on Saturday. Let's see when they are playing next. That's pretty cool. They are through the Champion's League final at the end of the month. We can keep going like this and find more news about the team just by tapping on the suggestions there.

But let's go ahead and tap on their roster. There we go. We have a carousel with all their top talent and should we check out everyone's favorite? Okay. There he is Mr. Ronaldo. Let's see some of his best tricks. So Amit only has to type "*best tricks*" and the Assistant understands the context and that he means Ronaldo's best tricks and responds with a YouTube video. We won't play that video now, but if we did tap it, it would play right there in the chat.

So, let's play a game instead. Amit's asking Google to play a game. And, okay. Google's suggesting the Emoji game. The way this works is Google provides a string of emoji and you have to guess the film title. Are you ready? Yeah? Okay, let's play.

Okay. So looks like we've got a crown and some rings. Any ideas? Lord of the Rings? I hear The Princess Bride over there. Amit, any ideas? Lord of the Rings. Oh, it was too easy. Okay. We could keep going but I think you got the idea. So that was obviously – should we keep going? No. Any ideas? They're good. Okay. So this is obviously a super simple game and there's a bunch more that you can play in Allo. But we really think that the best games will come from the developer community.

So that was the Assistant in Allo. Allowing you to harness all of the richness and knowledge of Google from search, YouTubes, Maps, Photos as well as many of our partners to have fun and get things done.

As you saw, the interface is inherently personal and conversational and the Smart Suggestions tab help to keep the conversation flowing. But most importantly for the first time, you can use Google in your chats with your friends.

And with that, I'll hand it back to Erik.

Erik Kay – Engineering Director at Google

Thanks, Rebecca and thanks, Amit. I'm really excited about the possibilities that Google Assistant brings to messaging.

Now let's talk a bit about privacy and security. We realize that sometimes people want to be incognito. Of course, this applies to chat as well. So, following in the footsteps of Chrome, we've created an incognito mode for Allo. While all messages in Allo are encrypted, chats in incognito mode are end-to-end encrypted. Incognito chats have discreet notifications hiding the sender and the message content from shoulder surfers and prying eyes.

Incognito also offers message expiration, so you're in control of how long your private messages stick around. And similar to when you close an incognito window in Chrome, when you delete an incognito conversation in Allo, it's gone forever and no one can see it again on your device.

With incognito mode, Allo gives users additional controls over their privacy and security and we anticipate adding even more security features to it over time.

So that's Allo. Allo is fast, smart, and secure and it lets you express yourself in fun new ways. And Allo will be the first home for the Google Assistant as it starts to take shape, bringing the richness of Google right into your chats.

Now, allo is all about messaging, but let me talk to you a minute about video calling. Just as technology has helped connect people through messaging, video calling has brought families and friends together from around the world. It can be a magical experience, so why don't we pick up the phone and do it more often? Well, maybe the other person has a different kind of phone. Maybe the connection is a little spotty. And even when it does work, getting a call can often feel intrusive because you're not quite sure why the other person's calling. That's why we challenged ourselves to design a video calling experience that would feel magical every time.

So, I'd like to introduce you to *Duo*, a simple one-to-one video calling app for everyone. Duo is the video companion to Allo. It's fast and performs well even on slow networks. It's end-to-end encrypted. It's based on your phone number allowing you to easily get in touch with the people you care about, and it works on both Android and iOS.

But here's a feature that makes Allo really special. We call it *Knock Knock*. Knock Knock shows you a live video stream of the caller before you've even picked up. Not only can you see who's calling but what they're up to and why they're calling, a smile, a beach, a newborn baby can all draw you into the moment, making calls feel spontaneous and natural. And once you pick up, Duo puts you right into the call seamlessly transitioning from the live preview to the live call. Knock Knock totally changes what it's like to receive a

video call.

So let's take a look. And almost right on cue, I'm getting a call from my daughter Ava. I use Ava for all my demos. So as you can see, and Elena apparently, too, is popping in there, I haven't even picked up yet but Ava is right there smiling and making funny faces. I can tell she is really eager to talk. So let's answer it.

[Ava: Hi, dad.

Erik Kay: Hi, Ava, hi, Elena.

Ava: Are you done with your presentation yet?

Erik Kay: Almost. I'll be done soon. Could you wave to everybody?

Ava: Hi, everybody.

Erik Kay: All right. Bye, I'll talk to you soon.

Ava: Bye, Dad.

Erik Kay: Oops!]

So you just saw how Knock Knock drew us into the conversation with Ava and Elena before we even picked it up and how fast and smooth it was.

Now, Knock Knock only works when the video is instant. There can't be a gap between the phone ringing and the video appearing. This is a really hard technical problem but we're uniquely positioned to tackle it. Duo was built by the team that created WebRTC, the open source platform that now powers much of today's mobile video communication. And we know this technology like no one else. We used a new protocol called QUIC that allows Duo to establish an encrypted connection in a single round trip. We obsessed over every last detail of video transmission, hand tuning and optimizing codecs, bandwidth probing, encryption and more. The result is the unique feel of Knock Knock. You're right there in an instant.

So the other thing that really sets Duo apart is how reliable it makes video calling. Duo works wherever you are, whether it's New York or New Delhi, Buenos Aires or Butte, at home or on the road. Duo proactively monitors network quality multiple times per second, allowing it to degrade gracefully when bandwidth is limited and to seamlessly switch between Wi-Fi and cellular. And, of course, this all happens automatically behind the scenes, so you don't have to worry about it. All this technology combines with a smooth interface that fades away allowing you to focus on each other in beautiful HD video and audio. This is what makes Duo a video calling experience that feels unlike anything you've used before.

Now, I'd like to show you a short video that shows how Duo brings magic back to video calling.

[Video Presentation]

You can really see how fun and spontaneous video calls become with Duo. We're incredibly excited to be connecting people all over the world with these two new apps: Allo, a smart messaging app with expression and the power of the Google Assistant. And duo, a simple video app with the spontaneity of Knock Knock. Both Allo and Duo will be available this summer on Android and iOS.

Now to talk to you a bit about Android, I would like to invite to the stage our resident rock star Dave Burke.

Dave Burke – VP of Engineering, Android

Thanks, Erik. Hey, I'm Dave. It's amazing to be here at Google I/O 2016 at the Shoreline Amphitheatre. All right. What a great venue! I think this is going to be the very closest I get to my childhood dream of being a rock star on the stage. There's something wrong about nerds being allowed on a rock star stage like this.

Anyway, so far we've talked about the scale of mobile and how we're thinking about evolving our products to be smarter and more assistive through machine learning and AI and a key driver of this scale is Android. As you've heard, this year marks the tenth anniversary of our first developer conference. It's also been 10 years since we started working on Android.

So how are we doing? Well, Android is the most popular OS in the world. More than 600 Android smartphones have launched in the last year alone. Everything from Disney Princess phones to metal unibody devices tricked out in titanium. There's a phone for everyone. And as Android continues to expand into new screens like on the wrist and in the car and on TVs and connected devices, there's increased opportunity for developers to reach users whether they're at home or on the go.

Android Wear. There are now 12 partner brands with iconic watch makers like TAG Heuer and designers like Michael Kors.

Android TV. There are now millions of new Android TV devices, growing rapidly, with media content and games from the biggest names in the industry.

Android Auto. More than 100 car models and stereos have launched with another hundred on their way by the end of the year. And, of course, Google Play. There were 65 billion installs in the last year alone and I'm just in constant awe of all the amazing apps and services that you're creating that's fueling this.

So let's talk about what's new in the platform. With the N release, we wanted to achieve a new level of product excellence, so we set about redesigning and rewriting many fundamental aspects of how the system works. Now, a lot of the features in N were inspired by users. How they use their phones, what they've told us, and how we think we can make their day-to-day experience better and

more useful.

This year we decided to do something a little different by releasing early developer previews of the N release before Google I/O. We want to share our work in progress with you as we build it, so we have more time for your feedback. Also, getting the platform out earlier means there's more time for app developers and device makers to be ready for the release later this year. The response to the N developer preview has been overwhelming. Many of you are already developing on the N preview release on a daily basis and it's just humbling to be part of a project of this scale. So thank you for all of the feedback so far.

Now, often one of the hardest parts of creating an Android release is coming up with the name. And I have no idea why, but this year the N dessert name is proving trickier than all of the others. So for the first time ever, we're going to be inviting the world to submit their ideas to www.android.com/N. And we're looking forward to your input but please don't call it NameyMcNameface. I should add that we all reserve the right to pick the winner.

All right. In the meantime, let's jump straight in and talk about some of the biggest changes in N around performance, security, and productivity.

Let's start with performance. We've improved performance in N in two key areas: graphics and runtime. In recent Android releases, we extended the OpenGL standard to bring advanced graphics capabilities, usually found on desktop and game consoles, to mobile.

With N, we're making our biggest leap forward with the introduction of Vulkan. Vulkan is a modern 3D graphics API designed to give game developers direct control of the GPU to produce incredible graphics and compute performance. And we made a concerted effort to work with the industry on Vulkan, so you can use the same APIs and graphics assets and shaders on the desktop as well as mobile.

Because Vulkan has a lower CPU overhead than OpenGL, game developers are able to squeeze in more effects per frame while still maintaining a high frame rate. Let's take a look at a Nexus 6P running a new version of the *Need for Speed* game by Electronic Arts and there are a bunch of really nice improvements in this version, thanks to Vulkan. You'll notice the beautiful graphics in Reflexions and materials on the car thanks to physically based rendering. Also check out the realistic motion blur effect which is computed for every object at the side of the road. And there's a really nice water surface effect on the road. And the shaders for these are pre-compiled ahead of time and can now run anywhere. So that's graphics performance.

We've also spent a lot of effort working on improving the Android runtime. First, we've made major optimizations to our compiler. The compiler in N performs anything from 30% to 600% faster on major CPU benchmarks like Dhrystone. Second, we've added a new just-in-time or JIT compiler, and JIT compilation means that app installs are much faster – 75% faster in N. So now

users can get up and running in your apps much more quickly. And also because JIT is more selective about which methods it compiles, we're also able to reduce the amount of storage needed for app code by a full 50%.

Now, unlike conventional JIT systems, the Android runtime uses profile guided optimization to write compiled code to flash the next time you run the app, so this improves performance and reduces battery consumption. In summary, the new JIT compiler improves software performance, makes installs faster, and reduces the amount of storage you need for apps on your device.

Let's talk about another big area of focus for us: security. We designed Android from the beginning with a multilayered defense-in-depth security architecture. And android employs the latest cutting-edge security technologies, things like SELinux, Verified Boot Integrity, and Full Disk Encryption.

With N, we're continuing to strengthen our defenses in three key ways. First, N introduces file-based encryption. By encrypting at the file level instead of the block level, we're able to better isolate and protect individual users of the system.

Second, we learned the importance last year of hardening the security of the media framework. Especially since it's accessing files from anywhere on the internet. So in N, we've split out key subsystems into individual SELinux protected processes, things like codecs and file extractors. By improving the security of the media framework, we improve the security of the entire device.

Third, and this is something that's really cool, N automatically keeps your phone up-to-date with the latest version of the system software without you having to do anything. Like Chromebooks, new Android devices built on N have two system update – two system images. So when an update is available, your phone will automatically download the new system image in the background. So the next time you power up your phone, it will seamlessly switch into the new software image. You're no longer asked for your password when the phone powers up, thanks to file-based encryption and a new feature called Direct Boot. and that pesky *"Android is upgrading dialogue"* is finally gone, thanks to the new JIT compiler. I think the best feeling the software industry is actually deleting code, by the way.

All right. This approach to software updates is one of the most loved features of Chromebooks and I'm really excited to bring it to mobile as well. So that's some of the ways we're improving security mechanisms in the platform but let's not forget about all of the security services that Google provides to keep all Android devices safe. In fact, when you think about the scale of Android and Google Play and the number of devices and apps out there, we're providing one of the most comprehensive mobile security solutions in the world.

Let's take a look at a few examples. Google Chrome protects users when they're surfing the web through a system called safe browsing. Safe browsing

warns users ahead of time when they're about to go to a site that we know contains malware or is known to be deceptive. Today, we're protecting over 1 billion mobile Chrome users.

Another example of how we protect users is through the Google Play store. All Android apps undergo rigorous security testing before appearing on the store. We review every app to make sure it meets our policies. We also run an app security improvement program with developers to identify potential security vulnerabilities. For example, we've worked with key banking and ecommerce applications to ensure they're using https properly to protect against man-in-the-middle attacks.

Google Play itself is built on a state-of-the-art cloud-based infrastructure we call SafetyNet. SafetyNet – with SafetyNet Google's expert systems and machine learning models analyze billions of signals every day to predict bad behavior. If an app steps out of line, Google Play will block or uninstall the app, no matter where it was installed from. And the scale of SafetyNet is extraordinary. Every day, we test over a billion devices and over 8 billion installed apps. And all of this happens under the hood to keep you safe and secure, no matter what version of Android you're on.

Let's move on. A third area of focus for us is our continued effort to improve productivity. And we've taken a close look at how people multitask on Android, to understand what's working for them and where we can improve. And we've particularly focused on the recent apps screen. And what we learned from our user research is that over 99% of the time, people only select an app within the last seven. So we decided to simplify by automatically removing apps in the list that you haven't used in a while. This then makes it much easier to find the app that you're looking for.

Also, based on popular demand, we've finally added a *Clear All* button at the top. Yeah. Feels good! But my absolute favorite feature is something that we call Quick Switch. You can now flip to the previous app you were in just by double-tapping the Recents button from anywhere. You can think of it like a simplified Alt-Tab. And it's amazingly useful in so many situations. For example, let's say I'm in a phone call and I'm trying to coordinate an event. I can flip over to the Calendar app I was just in by double-tapping the Recents button at the bottom right. From there, I can check my schedule and then flip back to the dialer by double-tapping the Recents button again. It's pretty cool.

Now, many of you have also asked for the ability to display more than one app at the same time. So we've invested a lot of effort in redesigning our window management framework in N. And we're introducing two powerful new windowing modes in this release: Split-Screen and Picture-in-Picture. Split-screen is designed for tablets and phones and it's really simple to use. So, for example, let's say I'm watching a video on YouTube to learn how to make the best nachos. I can long-tap on the Recents button to enter multi-window and from there launch something like Google Keep, for example. Now I can update my shopping list for ingredients while I'm watching the video.

The second mode, Picture-in-Picture, is designed for Android TV and it's a great way to let you keep watching something while you perform another task. For example, let's say I'm watching a live TV program on retro gaming and they're talking about Pac-Man and I want to see if I can install and play the game myself. I can put the live content into Picture-in-Picture mode to keep watching it and then go ahead and perform a voice search for Pac-Man. This will then give me an option to install the game from the Play Store, all at the same time as watching the content. It's pretty cool.

Notifications is another area we've worked on to improve productivity in Android and it turns out that today over half of the notifications shown in Android originate from messaging applications. So we decided to make some changes to really optimize for this use case. We've added a new *direct reply* feature which lets you quickly reply to a message like so. You no longer need to launch the app to fire off a quick response, so it's a real time saver.

We've also added a feature to give you more control over notifications. With N, you can long-tap a notification to change its visibility. For example, you can block notifications from a given app or set them to show only silently. So now you're able to choose which application – which notifications are important for you.

One other area we've worked on to improve your productivity in Android is your ability to express yourself with Emoji and Android is the first mobile platform to support the new Unicode 9 Emoji standard. And with this addition are more human-looking glyphs and support for skin tone variations. Unicode 9 also brings 72 new emoji glyphs, so now you can let your friends know, for example, when you're dancing like a left shark while juggling and eating avocado toast in order to win first prize in that selfie contest, basically my typical Friday night. Not!

But more seriously, we're really committed to this space and we're continuing to work with the Unicode consortium on the next generation of emoji. In particular, you may have seen some of our suggestions around better representing women in professional roles, so thank you for all the support for that so far.

All right. Let's wrap up. Android N is the best version of Android yet. I have to say that. But it's actually true. We've made it faster and more performant with the powerful new JIT compiler and Vulkan 3D graphics. We're continuing to harden our security and provide the first truly seamless software update system for mobile. And we're making users more productive, with better multitasking, brand-new multi-window support and improved notifications. In fact, there are over 250 major new features in N. Everything from Java 8 language support Lambdas to data saver, setting suggestions and much, much more. You can check out the "*What's New in Android*" session later today to learn even more.

We're still putting the final touches on the N release and we expect to launch it later this summer, but if you can't wait until then, I'm happy today to announce that we're publishing our first beta-quality release

candidate for you to try out on your main phone or tablet. You'll be able to opt into the new beta program at www.android.com/beta and run N on your Nexus 6, 9, 5X, 6P, Nexus Player, and Pixel C.

Now, there's one more area in N that we've been working hard on that we haven't talked about yet. And to tell you more about what it is and how it fits into our bigger plans, let me invite up Clay Bavor. Thank you.

Clay Bavor – VP, Virtual Reality at Google

Thank you, Dave. I'm Clay Bavor, and I lead the Virtual Reality team at Google. And, yeah, just to get right to it, virtual reality is coming to Android N. So it all actually started at a Google I/O two years ago with Cardboard. And since then Cardboard has done some pretty amazing things. There are millions of them out there in the world, in all shapes and sizes. We've enabled thousands of developers to build their first VR app, and users have installed over 50 million Cardboard-enabled apps. We think that's pretty good for what is, after all, just a piece of cardboard.

Now, we love Cardboard, and for us, it represents so much of what we think VR should be about. It should be mobile. It should be approachable, it should be for everyone.

But we knew it was just a start because there's a limit to how much you can do, how immersive of an experience you can create with some cardboard and with phones that were really only meant to be phones. We wanted to create something that has the best attributes of Cardboard but which is also comfortable, richly interactive and far more immersive. But to create that kind of immersion, you have to solve – to make your brain say, *yep, I'm somewhere else*, you have to solve a lot of really hard problems across all parts of the VR experience. You have to design a system that's capable of rendering at very high frame rate and resolution. To make the experience really comfortable you have to minimize what's called motion to photon latency. That's the delay between when you move your head and when the picture updates to reflect that motion. And you need to solve for how you interact with things in VR. And when you nail those things it just feels like you're there.

Well, we've been working on these problems and more. And what we've built won't be available until this fall, but we'd like to introduce you to it today. We call it Daydream.

Daydream is our platform for high quality mobile virtual reality, and in it are all of the ingredients you need to create incredible immersive VR experiences. Now over time Daydream will encompass VR devices in many shapes and sizes, but today it's about how Daydream will enable high quality VR on Android smartphones.

And there are three parts to it. The smartphones themselves, including VR optimizations to Android N, are reference designed for a headset and a controller, and apps, both how you get them through Google Play and the apps

themselves. We've designed and built each part in concert with the others with a focus on getting the end-to-end user experience just right.

So let's start with smartphones. Now the first thing we did was look at what it takes to build a smartphone that's great at being a smartphone but also at being the core of a VR system. And with input from the major silicon vendors and smartphone manufacturers, we've created a set of phone specifications for VR. And we call phones that meet these specs Daydream-ready. And the specs include things like high performance sensors for accurate head tracking, displays with a fast response time to minimize blur, and powerful mobile processors. And if a phone meets these specs, it will be capable of delivering a great VR experience.

But the smartphone itself, it's only part of the story. The operating system, the software, it needs to be able to make use of all of these capabilities, all while keeping latency to an absolute minimum. So we've introduced what we call VR Mode as part of Android N. We've worked at all levels of the Android stack to optimize it for VR. And we focused in particular on performance and latency, which we brought down to under 20 milliseconds. By adding things like single buffer rendering, and VR system UI, so notifications and alerts come through properly in VR. And all of this makes for a really comfortable VR experience that we think users are going to love.

Now, it's important, these improvements are part of the core of Android N, so the entire ecosystem can benefit. And what that means for developers is there are going to be a lot of Daydream-ready phones. In fact, Samsung, Alcatel, Asus, Huawei, HTC, LG, Xiaomi, and ZTE all will have smartphones that are compatible with the Daydream-ready spec and several will be available this fall. So that's phones with Daydream-ready phone specs and the VR optimizations as part of Android N.

Let's turn to headsets. Now this is obvious but a VR headset, it's something that you wear on your head. And because it's something that you wear, there are so many things you need to get just right. It has to have great optics. It has to be comfortable. The materials need to feel good. And it needs to be really easy to put on and take off. We've taken what we've learned in all of these areas and we've created a reference design for headsets that will work seamlessly with Daydream-ready phones. And we're sharing this design with partners across the ecosystem and there will be several of them coming to market with the first available this fall.

Now, when it comes to VR, everyone thinks about headsets, but the controller, how you interact with VR, it's just as important. We wanted to create a controller that's optimized for VR, that's both powerful and intuitive. And so we've been working on a controller for Daydream. Looks like this. And if we actually zoom in a little bit, you can see the controller itself, it's very simple. There are few buttons and a clickable touchpad so you can scroll and swipe. But hidden inside the controller is the magic. We've built in orientation sensors so it knows where it's pointing, how it's turning. And you can do some pretty awesome things with it. Let's have a look.

[Video Presentation]

So as you can see, the controller, it's super flexible, and the developers we've shared it with absolutely love it. Now the controller, too, will be part of the reference design that we're sharing with partners, with the first available this fall.

Okay. So we've talked about smartphones and operating systems, headsets and controllers. But ultimately, that's not what VR is about. It's about what you can experience. So let's turn to apps, what you can do in VR.

Now, first, if you're a developer, you know that there's a lot upstream, from someone using your app or playing your game. Users have to find it in a store, buy it, install it, launch it. Well, this will all work seamlessly in Daydream. And that's because we've built Google Play for VR. Users will be able to browse and search and buy and install VR apps in VR. And once you've installed an app, you can keep coming back to it from what we call Daydream Home, which gives you access to all of your favorite games and apps.

Let's actually talk about some of those apps, the things you can do and the places you can go. Our partners like the New York Times and The Wall Street Journal and CNN are bringing their VR apps to Daydream so you'll be able to experience the world's news like you're actually there. Hulu, Netflix, HBO, even IMAX are bringing their libraries to Daydream, so you'll be able to watch shows and movies in a virtual cinema or an immersive 3D film in a virtual IMAX theater. Here is a shot from just one of the dozen IMAX films that will be available. And I don't know about you but I'm pretty fired up about hanging out with astronauts in VR.

So, something else that's going to be awesome in Daydream is games. We have been working with the likes of Ubisoft, and CCP, NetEase and Electronic Arts, and these amazing developers are creating games that take advantage of all that we've talked about, and there are some really neat things in the works.

We've also been working on some of our own apps. Google Play movies is coming to Daydream, complete with high definition DRM video support. That means you'll be able to watch movies and TV shows from Play but in a virtual movie theater. Street View is coming to Daydream, so you'll be able to walk the streets of the world without having to fly around the world. And Google Photos will support VR photos so you can step inside and relive favorite moments.

And there's one more. YouTube. We've rebuilt YouTube from the ground up for VR. In it is voice search, discovery, your favorite playlists; again, all in VR. And we've added spatial audio, improved VR video streaming so you'll be able to step inside the world's largest collection of VR videos and experience places and concerts and events like you're actually there. And by the way, you'll also be able to watch every single standard video currently on YouTube, but in a very different way. And we think people are going to love it.

So that's Daydream, our platform for high-quality mobile virtual reality. Daydream-ready smartphones with VR optimizations as part of android N, a comfortable headset and a powerful, intuitive controller and some amazing apps and experiences, all designed in concert and open and at the scale of Android.

Now, Daydream arrives this fall but you can get started developing for it today with the latest Android Developer Preview. And we'll go into that, and more, tomorrow here at 9:00 a.m. So that's it for VR and Android. To tell you about wearables in Android, I'd like to turn it over to David Singleton. Thanks.

David Singleton – Director, Android Wear

Thanks, Clay. We launched Android Wear at I/O two years ago and since then, we've partnered with 12 brands to pair distinctive styles with the latest in wearable technology. And the result is an impressive collection of over 100 different beautiful designs. So whether you admire the heritage of TAG Heuer, or the iconic designer Michael Kors, spend your time hiking a trail or riding a wave, you can wear what you want. And Android Wear works with Android and iPhones, because no matter what phone you carry in your pocket, you should always be able to wear what you want on your wrist.

And today I'm sharing a preview of our biggest platform update yet. Android Wear 2.0. Over the past two years, we've learned a lot about what people want and don't want from a watch. We know that the most important role of your watch is helping you stay connected to what matters, to important, timely information, to the people you love, and to your health, all from your wrist. And that's why we're evolving the platform to build even better experiences for the watch face, messaging, and fitness.

Android Wear already has thousands of watch faces you can download and now we're making them even more useful by letting any watch face, show data from any app. So now you can mix and match the styles you love with the information that's most useful to you.

Let's take a look. Jeff Chang from the Wear team has the newest LG watch with LTE. And here you can see a watch face from Zuhanden. Jeff has customized it with his calorie count from Lifesum, stock information from Robinhood, and his tasks from to-do-list. If he wants to see other tasks for today, he simply taps right there on the watch face and sees a reminder to call mom.

And watches are uniquely suited to connect us to those people we love. You'll never miss a call from your child's school or a message from a close friend. And that's why we're redesigning key experiences on the watch to be even more intuitive and enabling new ways to respond to messages designed just for your wrist. This includes Smart Reply, that knows the context of your message, best-in-class handwriting recognition and a new keyboard, all powered by Google's machine learning, and here's what it looks like when Jeff gets a message from a friend. He taps to reply, chooses handwriting input and now he uses his finger to sketch big letters on the watch and the text does

recognize automatically and now she knows with one more go – she knows that he'll be there at 3:00 p.m.

Many of you want your watch to be like a personal coach, helping you stay aware, motivated and connected to your body. First, we're improving the fitness experience with Google Fit platform's automatic activity recognition. Second, your apps can exchange data using the Google Fit API, so information like calories consumed in a nutrition app can sync with calories burned in a running app and finally, we're expanding the ways you can enjoy listening to music while you work out. Even when you leave your phone behind.

When you want to go for a run, you can just go with the watch you're already wearing. No need to strap your phone in an awkward armband. And thanks to the hardware sensors on your watch and automatic activity recognition, apps like Strava will start tracking your time and distance when you start running. And if you enjoy music while working out, you can launch Spotify right from your watch face. And the best part of all of this, you don't even need a phone. In fact, when Jeff was showing you the demos, his phone was turned off! And everything you saw here today, from sending messages to streaming music worked on just his watch.

With Android Wear 2.0, apps can be standalone. That means the apps on your watch can have direct network access to the cloud via Bluetooth, Wi-Fi, or a cellular connection. And that means a fast and richer on-watch app experience for both Android and iPhone users. With standalone apps, watches with cellular support become even more powerful. You'll be able to make calls, get help from Google, and use your favorite apps right on the watch, no matter where your phone is or even if it's on or off.

Starting today developers can download a preview of Android Wear 2.0 and everyone will be able to enjoy these exciting new watch experiences in the fall. It's time for us to reimagine what's possible for wearables together. And we can't wait to see the incredible things that you will build.

And now here's one example from the wider Android developer community.

[Video Presentation]

Jason Titus – VP, Developer Products Group at Google

Hi, I'm Jason Titus. And I lead our Developer Products Group here at Google. And as you can see, software developers are changing our world in a big way. In fact, it's hard to find an aspect of our lives that hasn't been touched from how we learn to how we get around, to how we meet people. You are constantly finding ways to use new and emerging technologies to improve our lives.

And despite all of the innovation that has already occurred, the opportunities ahead are greater still, but it's really hard to go from an

idea to building a great product to getting it in user's hands. So I'd like to share some of the things that we've been doing over the last year to make this process easier on any platform.

Let's start with the web. The shift to mobile for the web platform is well underway, with over – we've recently hit a milestone with over a billion people using Chrome on mobile every month. We've been working on several initiatives to try and make the web work better on mobile devices. I want to talk about two of them. First, we've implemented powerful new web standards in Chrome that enable a new class of website to gain app like behavior. Like working reliably on even the worst networking connection or sending notifications to re-engage users, we call websites that use these features progressive web apps, because they get progressively better depending on the capability of the web browsers, and they can lead to dramatically better user experiences.

Second, we created an open source project called Accelerated Mobile Pages or AMP, to make it simple to create extremely fast mobile websites using existing web standards. On average, AMP pages load four times faster and use one-tenth as much data. These things load almost instantly. So we're making it easier to develop mobile websites, but we're also investing in native development. We want Android Studio to be the quickest, most reliable way to build Android apps. I'd like to invite Steph up to tell you more.

Stephanie Cuthbertson – Group Product Manager, Android Studio

Thanks, Jason. I'm Steph. I work on the Android team. Android Studio is our official IDE. It's purpose-built for Android. And it was only three years ago, right here at I/O, that we showed it for the first time. Since then, it has built a ton of momentum. 92% of the top 125 apps and games now use Android Studio and millions of developers worldwide.

Our Android engineering team sees that almost all the professional developers that we connect with have used it and switched over from our Eclipse tools. Your feedback has been awesome because it's helped us focus on the right things. We care a lot about making it great. And there's much more to come.

This morning, we'll release our latest, which is Android Studio's 2.2 preview, focusing on speed, smarts, and Android platform support. First, speed. Already with 2.1 we made building and running changes 10 times faster than it was only six months ago. Our new Instant Run feature drove most of this. Now when you hit "run" we deploy changes directly into the running app. The emulators are three times faster and push speeds are even faster than that. So all of that means the emulators are now actually faster than the physical device that's probably in your pocket right now. But we want to keep speeding you up to launch earlier.

One of the things that you told us is that we could not possibly make build speeds too fast. So with 2.2 we're going to accelerate build speeds again. Plus we know Instant Run is a long-term bet so we'll keep expanding coverage.

Another thing we want to do is make it easier to write tests to help drive up app quality. So with test recording, now as you just tap through your app, we'll record – we'll write all of your Espresso test code for you. Almost – it's basically as if – it's as if you had handwritten the code yourself and you can run those tests locally or you can run them with our IDE-integrated cloud test lab, to make sure your app runs well on many Android devices.

Now, one of the things the engineering team is most excited about is that with 2.2 you're going to be able to build layouts faster and they'll run faster, too. So let me show you this. 2.2 includes a rewritten feature-rich layout designer, with new constraint layouts. So what you're going to see here is you'll be able to use the layout window. It's almost like a sheet of paper. So you can position your widgets and once you're happy, you'll see here we save you a ton of time because we're going to automatically add all the constraints for you, doing a bunch of actually very cool math under the covers.

You'll see the UI here is now going to adapt so you can try it on different Android devices. It will adapt to different orientations. But what you can't see is behind the scenes your app will run faster, too. So many of you here are professional Android developers and you know building rich UI usually requires nested layouts and those are hard to performance-tune at times. With constraint layouts, there is no nesting required, so overall you get the same rich UI, less work, and better performance by default.

Second, smarts. We want to help you write better Android apps. The 2.2 preview includes a new APK Analyzer so you can figure out, *"Hey, what's making my app so big?"* This is really useful for all of you who are targeting emerging markets. You've talked to us a lot about this. You also had asked for a layout inspector so we have a new layout inspector in 2.2, so you can find out what's inside your Android layouts.

When you run code analyses, you'll find a bunch of new quality checks, and that's in addition to the 260 that are already there, all designed to eliminate common classes of errors that we see on Android. And we integrated the latest version of the wonderful IntelliJ IDEA 2016.1.

The third thing I want to talk about is Android Platform Support. So we developed the IDE right alongside the platform, so we can bring you the very best. 2.2 includes, to name just a few, support for N's new Jack compiler which is bringing you wonderful Java 8 features like Lambdas, default methods. And for all of you developers who are working on graphics-rich like games and apps, you know C++ is critical and it's because of your requests we've been subtly building out C++ support.

Now, we already have C++ editing and debugging, but in 2.2 you'll find a major change based on your feedback. In addition to Gradle, for build we'll now support Cmake and NDK build. And both of those build systems will work with the debugger and then on top of that, we made C++ debugging better also.

So I would really appreciate it, we would really appreciate it, if you

download and try all of this today and give us feedback. With that, I would love to pass it back to Jason. Thank you very much.

Jason Titus – VP, Developer Products Group at Google

Thank you, Steph. It's great to see Android Studio getting so much faster. And there is one more product I'd like to talk to you about.

18 months ago, we acquired Firebase, a great mobile backend-as-a-service (BaaS) for storing your apps data and syncing it across iOS, Android and the Web. Since then, its usage has grown to over 450,000 developers. We would really admire its great developer experience and we wanted to figure out how to bring that to other areas of app development. We've been working closely with companies big and small all over the world to explore how we could make our offering better.

Today, we are announcing the next generation of Firebase. Firebase is now a suite of integrated products to help you build your app, grow your user base, and earn money. This is the biggest, most comprehensive developer update we have ever made. So at the heart of the new Firebase is a mobile analytics tool we've built from the ground up called Firebase Analytics. It's inspired by much of the work that we've done in the last 10 years with Google Analytics, but it's designed specifically for the unique needs of apps.

It's great for app developers for a couple of ways. First, it gives you rich insights into what users are doing inside of your app. And it also tells you where they're coming from with rich cross network attribution; all of this in a single dashboard.

Second, we've created a new feature called Firebase Audiences, which integrates all of Firebase together. With Firebase Audiences, you can group users based on the criteria that matters to you most, and then take action with notifications, experiments, and even re-engagement campaigns on Adwords. And the best part, Firebase Analytics works across Android and iOS and it is completely free and unlimited. Surrounding analytics are over a dozen other major features. I'm going to cover some highlights.

First, there's cloud messaging and notifications, which is built on the world's most popular cloud-to-device messaging platform, with over one million apps sending over 170 billion messages every day. And now that it's integrated into Firebase, you can send targeted notifications to your audiences without writing any code. And, like analytics, it is completely free and unlimited.

Now, app quality is important to all of us. If your app crashes, it's bad for your users, and it hurts your business. Within Google, teams like Chrome build and deploy across multiple platforms at scale. So we've built on their infrastructure and created Firebase crash reporting. You can use it to quickly identify bugs and issues, and take action to reduce impact on your users.

Another aspect of building a high-quality app is being able to tune and experiment with features. So we've taken the same infrastructure that we use for our own apps and created Firebase Remote Config which lets you create experiments and test app configurations at scale.

Finally, to help you grow, we've created Dynamic Links. A Dynamic Link is just a regular URL whose behavior can be configured depending upon where it's tapped. It persists through the app install process so you can drive installs while still maintaining a superior user experience. In this example with NPR, with Dynamic Links, we were able to take what was 13 taps down to only 4. So as I mentioned, there are over a dozen features of the new Firebase, and while each of them is interesting on its own, the power of Firebase is in how they're integrated together.

Here's a quick example. First, with crash reporting, you can see crashes as they're happening. Then in analytics, you can look at what the impact is on your business. Once you fix the crash, you can actually use notifications to reach back out and invite users back into your app. Or use remote config to offer them a coupon. And all of this is integrated into our Google Cloud platform, so that if you want to do deeper, more customized analysis of your analytics data, you can export it to BigQuery, Google's fully managed, petabyte scale data warehouse.

So I've only skimmed the surface, but we're going to go into detail later on this afternoon, and we have over 30 sessions on Firebase throughout I/O. The new Firebase is available today through one easy to install SDK across Android, iOS, and the Web. Please give it a try. We're really interested to hear what you think and we can't wait to see what you build next.

And with that, here is Ellie to give you a preview of a new Android effort that we're working on for you.

Ellie Powers – Product Manager of Google Play

Thanks, Jason. Hi, everybody. I'm Ellie from the Android team. And I'm here with Ficus Kirkpatrick, our Engineering Director. We hope that Firebase will make it a lot easier for all of you to build great apps. And we're also thinking hard about how Android apps should evolve and do more. We'll be showing you a sneak peek of a new project. We think it's going to change how people experience Android apps, and we'll be bringing it to Android users and developers over the next year.

Developers like all of you here have built amazing Android apps. They unleashed the full power of Android devices, seamlessly combining the camera, sensors, smooth animations, and more. But you've really told us that you want to be able to bring users into your apps more quickly. We want to make it easier for users and developers to connect; for users to access a wider range of apps and for developers to reach more people.

With the web, you can just click on a link and land on a web page. That's one

click and a few seconds. What if you could run any app with one tap? That's what we're working on. We're evolving Android apps to run instantly, without installation. We call this Android Instant apps. We're going to share a quick preview of what we've been working on, and I can't tell you how much I've been looking forward to this moment for a really long time.

So let's say my friend Michael sends me this link to Tasty on BuzzFeed video. And let's be completely clear. Let's be honest with each other. I do not have the BuzzFeed video app installed on my phone. So what we're going to do is we're going to tap the URL, and it's going to take me right into BuzzFeed video's Android app, without installing it.

Ficus, go! What's happening here is Google Play is fetching only the pieces of the app that we need right now – there. We are in the Android app and I didn't even have to install it. It was pretty fast too; right? So here you can see there's a bunch of different videos showing to make a whole bunch of different recipes, and the videos start playing automatically. Remember, it's a real Android app, right? And I can just swipe and go to the next video really quickly. The app was able to open so fast because it's been split up into modules. When Ficus tapped the link, Google Play downloaded only the code that was necessary to display this screen. And if I want to keep BuzzFeed video on my home screen, it's simple to install the app right here.

Here's another example. B&H photo and video has a beautiful Android app but I don't have it on my phone because I don't shop for cameras every day. Now if I'm searching for something specific, like a camera bag I can still get that same experience. With one tap, the app opens up right to the bag I want to buy. Technically, this is a deep link to the Android activity B&H wrote to display this product page. And that's all Google Play needed to download. I can also swipe here and see more details about the bag.

Now, when I add it to my cart, the animation there, it was pretty sleek and at checkout time Android Pay works, just like if I had the app installed. I don't have to pull out my credit card or type in my name and address. With Android Instant apps, I'm already signed in and I am ready to pay. So it's going to take me two taps, not two minutes.

Finally, let's see how Android Instant apps can help me when I'm out and about. So I walk up to a parking meter and I need to pay. I'm in a real hurry today and I don't have time to install a parking app, but what if I could just tap my phone, and with NFC it could bring up the parking app immediately? All I have to do is choose how long I want to park. I'm already done. And now Ficus can run off the stage to his meeting and even add more parking time later if he really needs to.

So that's Android Instant apps. As a user, it's totally seamless, from launching the app, to signing in, to making payments.

Now, as a developer, you'll update your existing app. It's the same Android APIs, the same project, and the same source code. And it can take less than a day of work, depending on how your app is built. You'll modularize your app

and Google Play will download only the parts that are needed on the fly, as we saw here.

We're really excited to give developers more ways to get their apps into the hands of users. In addition to discovering your apps in the Play Store and installing them, Instant apps will provide another on-ramp. People can use your Instant app directly, and as we showed earlier, if they want to install it, that's easy, too. Most importantly, you will be in control of the experiences you build. And when you do make the update, your app will be just one tap away for over a billion people.

Oh, yeah, I almost forgot to mention one really important detail. So this demo that Ficus did, he actually did that on a phone running KitKat. You should know that Instant apps is going to be compatible all the way back to Jelly Bean. And we really want for all of you to be able to try this, and we're really going to make it available to you just as soon as we can. But it's a big change in how we think about apps. We want to get it right. And so that's going to take some time for us. We're working now with a small set of developers and we'll be gradually rolling out access and actually rolling out Instant apps to users later this year.

We are so excited about all of the opportunities that this will open up, and we can't wait to see what you are going to build when your app is just a tap away.

And with that, I'll hand it back to Sundar.

Sundar Pichai – CEO, Google

Thank you, Jason and Ellie. Firebase is the most comprehensive developer offering we have done till date. I'm glad it's available today and look forward to hearing your feedback.

We talked a lot today about machine learning and AI. We think there's an opportunity to accelerate computing by working on this with everyone else. And so we're trying to do that in two ways. First, we are opening up core components of our machine learning systems. Last year, we open-sourced TensorFlow so that developers can embed machine learning and deep neural nets with a single API. In 2015, it was the most forked project on GitHub, and it is the number one machine learning project on that site.

Last week we open-sourced a natural language parser which is also built on top of TensorFlow. We are doing these things so that we can engage the external community and work on this together with everyone.

Second, for developers and companies out there, we are also exposing our machine learning capabilities through our Google Cloud platform. We already have a cloud machine learning platform underway, and you get access to computer vision, speech, language, and translation APIs. And we are working on bringing many more machine learning APIs so that you can get access to the same great capabilities we have inside at Google. We believe this will be one

of the biggest differentiators for the Google Cloud platform over time.

And, by the way, when you use Google Cloud platform, you not only get access to the great software we use internally, you get access to specialized hardware we build internally. And talking about that, for machine learning, the scale at which we need to do computing is incredible and so we've started building specialized custom hardware. We call these Tensor Processing Units or TPUs.

TPUs deliver an order of magnitude higher performance per watt than all commercially available GPUs and FPGAs. And when you use the Google Cloud platform, you get – you can take advantage of TPUs as well. TPUs are what powered AlphaGo, DeepMind's AlphaGo, in its game against Lee Se-dol. Go is an ancient Chinese board game, has a simple 19-by-19 grid but it is one of the most complex games humans have ever designed. It has more possible board configurations, many more possible board configurations than there are atoms in the universe.

Beating Go for computers was widely considered to be the grand challenge for AI and most people thought it wouldn't happen for another decade or so. So we are really thrilled that AlphaGo was able to achieve this milestone recently. One thing worth calling out, in the second game there was a move 37 by AlphaGo. It changed the course of the game and is now widely considered one of the most beautiful Go moves ever seen in tournament play. It was not just an intuitive move but a very creative move. We normally don't associate computers with making creative choices and so to us, this represents a significant achievement in AI.

By the way, Lee Se-dol has won every single game since he played against AlphaGo and he has even replayed some of the moves he learned from AlphaGo in that game.

As a state-of-the-art capabilities in machine learning and AI progress, we see them becoming very versatile and we think it applies to a wide range of fields. I want to give you two examples. First, robotics. At Google, a bunch of 20%ers decided to help robots – train robots to pick up objects. You can see it behind me. This is not new. You usually do it by writing control system code. You program the robots with rules. But this time they decided to use deep learning techniques. So you create a continuous feedback cycle so that the robots can learn hand-eye coordination by themselves using deep learning. As you can see, they keep doing it and they keep learning it over and over again and over time, they get better and they even learn natural and useful behaviors. So, for example, the robot is nudging the stapler away to pick up that yellow object. We didn't write that rule. The robot learned it automatically using deep learning. So it's an amazing example of what machine learning can do.

Another example in healthcare, diabetic retinopathy is the fastest growing cause of blindness. It affects 4.2 million people in the US and many more worldwide. To detect it, you need to do a scan, a scan of the eye like you see behind me and a highly trained doctor can detect it. If detected early,

the treatments are effective. If detected late, it causes irreversible blindness. And it's very, very difficult to have highly trained doctors available in many parts of the world. So we set out to work with a small team of engineers and doctors and, again, use deep learning and started training on eye scans. And over time, our computer vision systems have gotten really good at detecting diabetic retinopathy early. This is still early, and there's a long road ahead, and we'll work with the medical community to get it in the hands of as many people as possible. But you can see the promise, again, of using machine learning.

When I hear about advancements like these, I'm reminded that we live in an extraordinary period for computing. Whether climate change, healthcare, or education, the most important struggles already have thousands of brilliant and dedicated people working to make progress on issues that affect everyone.

Now, consider what the best climate change researchers, doctors, or educators can do with the power of machine learning assisting them. As you have seen today, I'm incredibly excited about the progress that we are making with machine learning and AI. We believe that the real test is whether humans can achieve a lot more with the support of AI assisting them. Things previously thought to be impossible may, in fact, be possible. We look forward to building this future together with all of you.

Thank you for joining us at Google I/O.