

Google I/O 2026 keynote in 35 minutes

35 MIN · YOUTUBE · [HTTPS://YOUTU.BE/OMhKgQmeMhI?si=97U890pxy-EBZEEQ](https://youtu.be/OMhKgQmeMhI?si=97U890pxy-EBZEEQ)

<https://youtu.be/OMhKgQmeMhI?si=97U890pxy-ebzEEq>

SUMMARY

Google's recent announcements at I/O 2023 highlight significant advancements in AI and user experience across various products. Key features include the introduction of Ask Maps and Ask YouTube for more intuitive search capabilities, enhanced voice interaction in Google Docs, and the launch of Gemini Omni for dynamic multimedia creation. Additionally, new tools like the universal shopping cart and personalized AI agents aim to streamline user tasks and improve digital interactions.

- *Ask Maps and Ask YouTube***: New features allowing users to ask complex questions and receive tailored, digestible responses.
- *Docs Live***: Voice-activated document creation and editing, enhancing productivity for users.
- *Gemini Omni***: A multimodal AI tool for creating and editing videos and images with conversational prompts.
- *Gemini Spark***: A personal AI agent that manages tasks and integrates with various Google services for seamless user experience.
- *Universal Shopping Cart***: An intelligent cart that tracks prices and deals across platforms, enhancing online shopping.
- *AI Search Enhancements***: New search box features for multimodal queries and dynamic, interactive search results.
- *Daily Brief Agent***: A personalized digest summarizing important tasks and information for users.
- *Audio Glasses***: New wearable technology enabling hands-free interaction with AI for navigation and task management.

Hello, Shoreline and hello to everyone watching from around the world. Recently, Maps got its biggest upgrade in a decade, including a new feature called Ask Maps. People are using it to ask more complex and much longer questions. Here is a real query from a parent. My kid just fell into the duck pawn and the wedding starts in 30 minutes. Where can I walk and buy her a new dress? I'd like to hear how that turned out.

We're also bringing this conversational AI to two more products. First, Ask YouTube. People come to YouTube every day to ask a lot of questions. It's a lot of great videos. Sometimes it's hard to know where to start. Ask YouTube entirely reimagines the experience. Say you want to teach your three-year-old how to ride a pedal bike and they already know how to ride a balance bike. Just ask YouTube. You will see a couple of differences in results. The information is digestible and easy to navigate. You get an overview and helpful tips. You will see videos that best match your interest. So, if you want to try a specific method of teaching, you can go deeper there. And best of all, it jumps right to the part of the video most relevant for you. It remembers the context so you can follow up with questions like should I buy one with handbrekes or pedal brakes making it an ongoing conversation. It even lays out the information in a table so it's easy to compare.

We are starting to test Ask YouTube now and it'll roll out broadly in the US this summer. With Docs Live, you can just verbally brain dump whatever is on your mind and let Gemini do the rest. Let's see it in action with a demo from our product team. This is all in real time, not sped up. All right, let's try this out. So, I just remembered I'm doing an alumni talk for my high school's career day tomorrow. Uh, I need to come up with some talking points to explain what I do for a living as a software engineer, but I'm not really sure where to start.

Oh, actually, can you just pull my resume from Drive? Although, that might be boring. Um, maybe can you come up with some funny analogies so it'll be more of an engaging talk for the students? Oh, and also I think the school sent me an email. I think the subject is something like career day logistics. Maybe just grab the details from there, throw them at the top of the doc so I know where to go and what time to get there. Let's just update those requirements and turn this straight into a draft.

This is cool, but it's a little dense. Maybe format the analogies as a table so it's a little easier for me to scan and also add a note to tell the story about how my brother inspired me to become a software engineer sort of at the top of my doc and bold it so I don't miss it. Yeah, that looks great.

In the future, you'll be able to create new docs and edit them directly, all with your voice. Docs Live is rolling out for Pro and Ultra subscribers this summer, and the same powerful voice capabilities will come to Gmail and Google Keep, too. >> I'm excited to announce Gemini Omni. Our new model that can create anything from any input. It combines Gemini's intelligence with the best of our generative media models for a new level of world understanding, multimodality and editing.

Models like VO, Nano Banana, and Genie are able to create extremely realistic videos, images, and interactive simulations. Although not perfect, they already demonstrate some impressive notions of intuitive physics. And with Omni, we've now made even more progress. It's a step change in simulating things like kinetic energy and gravity. Previous systems would have found these concepts difficult. Gemini's world knowledge and reasoning really shine in Omni. It can translate complex ideas into highly accurate videos. So for example, you can give it a simple prompt like make a claimation explainer of protein folding and get this.

>> Proteins start as chains of amino acids. They fold into patterns like the alpha helix and flat sections called beta sheets forming a perfect three-dimensional shape. >> But the initial generation is just the start. The creative process is rarely a single step. It's usually iterative. Just like Nana Banana redefined image editing, Omni gives you a more natural way to edit video with conversational language. What's really cool is you can give it your own videos, for example, this selfie, and change reality in a really fun way.

You can easily adjust the details and style or even add elements and the whole scene to reflect your new idea. We are now going a step further and adding content credentials verification across products. This will show you if the origin of the content was AI or a camera and if it's been edited with generative AI tools. In this example, Gemini can tell this photo was captured with the Pixel camera and then edited with Google Photos. We want more people to have easy access to these tools. So we are expanding both synth ID and content credentials verification to search and chrome.

You can simply circle to search or right click in Chrome and ask was this generated with AI and you'll get a clear response along with other helpful context. For example, this image was making the rounds on social media last year. It's obviously fake. I don't eat hamburgers. It might not be as clear to everyone else. That's where these tools can be really useful. Of course, this only works at scale if more partners decide to watermark their own AI generated content. Nvidia signed on to Synth ID last year. And today, I'm thrilled to announce that OpenAI, Cacao, and Leven Labs are adopting Synth ID, too.

And today I'm excited to introduce Gemini 3.5 Flash, our first in a series of models combining Frontier Intelligence with action. Two things I would highlight. First, when compared to 3.1 Pro, Flash is better across the board. almost all benchmarks. It's made huge progress in coding and look at that extraordinary jump in GDP val a benchmark that captures many real world economically valuable tasks.

Second, 3.5 flash is a very capable model at the frontier and comparable to the best models but much much faster which is why when you look at the intelligence versus output speed it's in a whole league of its own in the top right quadrant. When looking at output tokens per second, it's four times faster than other Frontier models and it's incredible delight to use. And now anti-gravity is massively expanding its suite of agentic capabilities, surfaces, integrations, and product features. To start, we're launching a full CLI experience, an anti-gravity SDK, native voice support with Gemini audio models, and integrations with many services and platforms like Android, Firebase, and Google AI Studio.

All of this is available for you to try today. But most importantly, at the core is anti-gravity 2.0, a new standalone desktop application that delivers fully on that original glimpse of a truly agent optimized experience. The new anti-gravity is unabashedly agent first. Focusing on the core agent conversations, agent produced artifacts, and multi-agent orchestration. The anti-gravity agent harness, the invisible framework for Gemini to perform real world tasks, has become much more powerful with new core primitives such as sub aents, hooks, and asynchronous task management.

And underpinning all of this are the Gemini models with Gemini 3.5 Flash having been co-optimized with the anti-gravity harness. Of course, being engineers, we were curious to see how far we could push the limits of what was possible with these agents and models. So, using the new anti-gravity and Gemini 3.5 Flash, we asked our agent to take on what we consider to be a highly complex and impressive task. Build a working operating system from scratch. But clearly, this isn't a real OS unless I can play Doom. Let's take a peek and see if it works. Moment of truth.

Amazing. That never gets old. Today, anti-gravity 2.0 is available globally for everyone. Gemini 3.5 Flash is available for everyone today across our products and APIs. We are also excited for 3.5 Pro. We're using it internally. It's showing great improvements. I know you can't wait to get your hands on it. Give us until next month to get it to you. Introducing Gemini Spark. It's your It's your personal AI agent that helps you navigate your digital life, taking action on your behalf and under your direction. It runs on dedicated virtual machines on Google Cloud. And it's 24/7.

And yes, you can close your laptop. It's powered by Gemini 3.5 and the Google anti-gravity harness which allows it to perform longrunning tasks easily in the background. Spark will integrate seamlessly with tools starting with our own and in the coming weeks with third party tools through MCP and you can work with Spark however is most convenient in the Gemini app or soon through email and chat. So, watch this.

Start a few threads for me. The first one, find all the upcoming meetings with Sundar and turn them all hot pink so I don't miss them. The second one, last night I met our new neighbor, John. Write a note to him and his family. Invite them to our block party because they weren't on our list originally. The third one, create a document with the top things my wife and I need to do for the kids before the end of the school year. Categorize it by deadline and priority. Make it easy to digest. I don't want to miss anything.

All right. So, we'll send that in. And you can see at the speed of my voice, it's taking that one task and it will capture all of that context as fast as I can talk. It starts out as a single thread here. And in the background, it's actually going to go through and break those down into individual tasks. Now, I can just put my phone away and get on with my day. And Spark works in the background for me. We'll check in later, see how it's doing. Because we're prioritizing safety, we're rolling out Gemini Spark deliberately to trusted testers this week and as a beta for us Google AI Ultra subscribers next week.

We want this new type of help to be in as many hands as possible. So to do that, we're introducing a new ultra plan for \$100 a month. And for those of you that need maximum limits, we're dropping the price for our top tier ultra plan from \$250 a month to \$200 a month. And there's so much more to come. Later this summer, Gemini Spark will operate directly within Chrome, acting as your agentic browser across the web. It can take action on your tasks under your direction. We're also building a dedicated home base for your agents on your phone, Android Halo, which is coming later this year. I'm excited to announce we're launching a brand new intelligent search box. This new search box puts our most powerful AI tools right at your fingertips, and you can ask across modalities with text, images, files, videos, and search reasons across them all. Now, this is the biggest upgrade to our iconic search box since its debut over 25 years ago and is starting to roll out today. Next, we're making it even easier to continue the conversation with search, bringing AI overviews and AI mode into one seamless AI search experience, so you can flow effortlessly from your question to your response on the main search results page to follow-ups in AI mode. And your context stays with you. And as your conversation gets deeper, your links and your sources get even more relevant to what you might want to explore.

So you continue to get the best of AI and the best of the web. And I'm excited to share this new seamless AI search experiences live today across desktop and mobile worldwide. We're entering the era of search agents. Now to start, you can set information agents to work for you 24/7 in the background. They can find you exactly what you need, exactly when you need it, and help you take action. You can spin up multiple agents and search simultaneously to get updated and make progress on all those things that matter to you. And these will work with and alongside Gemini Spark to help you get more done. Now, let's say you're apartment hunting. You can do a total brain dump of what you're looking for with all your criteria like location and natural light and availability and your agent will continuously scan the entire web across sites, social and forums. Or if you're a sneaker man, you can just ask to be updated when any of your favorite athletes announce sneaker collabs or

drops. And it monitors everything from blogs to our shopping graph so you don't miss out. Now you'll be able to put information agents to work for you this summer. Just ask search to keep you updated on whatever you want to know.

>> We're bringing anti-gravity and the agentic coding capabilities of Gemini 3.5 Flash right into Search. So Search can build you the ideal format exactly for your question, completely custom on the fly. We're talking dynamic layouts, interactive widgets, entire experiences, all created just for you. This is a gentic coding at the scale of search.

Let me give you an example. Say I'm a college student trying to wrap my mind around astrophysics. I can go to search now and just ask, "How do black holes affect spacetime?" And check this out. I get this interactive visual right in the AI overview. search now gets for a concept this complex. I need to interact with it to really understand it. But this is still kind of 101. So I'm going to follow up and say now show me how two orbiting objects like binary black holes create gravitational waves. Search dynamically builds a brand new interactive visual in real time completely custom for my specific question.

Generative UI with anti-gravity is rolling out to search this summer for everyone free of charge. >> I am excited to announce the universal cart, a truly intelligent shopping cart. It works across merchants and across services. You'll be able to add things to your cart when you're browsing search, chatting with Gemini, watching YouTube, or even reading your Gmail. The moment you add a product, your card goes to work for you in the background. It finds deals, uh, looks at price drops, gives you insights on the price history, and alerts you when something comes back in stock. It all runs on our Gemini models, so your card gets even smarter as the models improve. So, just think of it as shopping with superpowers. I am excited to announce that we're rolling out the universal card in the US across search and the Gemini app this summer with YouTube and Gmail to follow.

>> Today, I'm excited to announce that we've completely redesigned the Gemini experience from the ground up. From the moment you open it, we've greeted you with a stunning new design language we call neural expressive. We've added fluid animations, vibrant colors, new typography, haptic feedback throughout the app. But we all know that good design isn't just about how it looks. Good design is about how it works. So, we've involved the entire experience. We've made it easier to discover and generate those gorgeous images, videos, and music with built-in templates that you can easily remix. And we've completely transformed the Gemini Live experience. It now opens up immediately and in line and soon you'll be able to pick a regional dialect that resonates with you.

You've got a right good mix of different accents knocking about like this one from Liverpool. Neural Expressive is rolling out globally on Android, iOS, and the web starting right now. Gemini Omni is coming right into the Gemini app for paid subscribers today. It's going to let you bring your ideas to life using any combination of text, images, and video inputs. And as I've been using it, it feels like the nano banana for video moment is here. It's never been so easy to create and edit videos. Let's look how this plays out in the real world. I want you to meet Sashu. She is working on a new song and she wants to create a quick video teaser. So she shares the raw video.

She adds some reference visuals to it. And now this is the coolest part. She can transform the style of her video. maybe even switch the camera angle to a 360 degree shot. And Gemini puts it all together for her. Let's take a look at what it looks like. All done within Gemini. Google AI Plus, Pro, and Ultra subscribers around the world can try Gemini Omni today right in the app. The third update today is about how agents are coming to Gemini. This is a big shift for Gemini because agents don't just answer questions, they proactively work on your behalf. And to show you what this means, I want to introduce one of our newest out-of-the-box agents called the daily brief. It's a personalized digest that's designed to be your first stop every morning. Here's how it works. You can see here that it's synthesizing information from across my inbox, my calendar, my tasks. It's finding the most important things for me to be aware of. And I totally could forget these things like that returning that book. It's not just summarizing the data though. It's organizing it by topic. It's even suggesting the next step. And with this travel info, I can just take the next step right in line.

All of it is super concise in this morning digest that's built for skimming. And I can't wait for you to try it out. The daily brief is rolling out today also to Google AI plus pro and ultra subscribers starting right here in the US. You'll remember at the beginning of the show we sent a few tasks off. So let's go check in on them, see how they've completed. This was our school year planning checklist and I asked it remember to create a doc of all the things I didn't want to forget between now and the end of the year. And so I'll open this one up. And what's amazing about it is it takes advantage of all the Google Docs formatting. So I can immediately click in and see the checklist here for our various kids all subdivided like this and easily to go one by one. The date, the activity, the color coding all integrated in one spot.

Pretty incredible how much time this can save. Now, I know I've promised just three updates, but we have one more. Last month, we dropped the Gemini app for Mac OS. Here it is on the screen. It's gorgeous. This was a small team that built this native app from scratch using anti-gravity. They did over a 100 features in less than a 100 days. Now, two weekends ago, we were hacking on the Mac app and we came up with something cool and we wanted to sneak it into the show. You all want to see it live?

>> Yeah. >> All right. So, what you can do with this on Gemini on Mac OS is actually take a look at a bunch of documents like this. You'll be able to select them all and then long press the function key and just dictate the email to the kennel. So, it works something like this. Hi there. I need to do a short boarding stay for my two dogs, Louis Cinnamon and Hank, starting this Thursday. Oh, wait, no, actually, it's this Friday. They've never stayed with you before, but they're very social dogs. And also, can you turn these files into a table with their details, allergies, recent vaccines, and make this email sound friendly so we make a good first impression?

All right, I'm going to release the function key. You can see Gemini's thinking at the bottom here on this MacBook. What it's done is because I've selected those files in Finder using its multimodal understanding, it can go through the PDF, it can go through these images of their invoices, and it's all controlled by my voice. So, it can actually take all that complex information and look at that. There it is. It's got a table in line.

It's also so amazing because it corrects. Remember I said Thursday, no scratch that Friday. And it picks up that and automatically cleans up my input. This is the power of what Gemini can do using your voice. These new

voice capabilities and Gemini Spark will be coming to the Mac app this summer as well. Introducing Google Pix, a new product in Google Workspace. Pix is our image creation and editing tool that helps you create just about anything from party flyers to infographics with the creative controls you want.

Watch how easy this is. You start with a base image as your canvas. And what's really cool is that Pix understands what's in your creations and how the objects work together. You can hover over an element and click to remove it. Or you can resize an object to fit the frame. Once the layout is set, you can add or edit text and translate all of it with just a few clicks.

>> Pretty cool. Every output from our creative tools, including Pix, is watermarked with synth ID. And Pix is rolling out this summer. Take a look at this raw footage. I love how this person is walking, his presence, his pacing. Let's not change any of that. With a simple prompt and style reference, Omni allows us to transform the environment, add visual effects, and any other element, all while preserving the original performance.

And now you can even add new characters while maintaining everything else in the scene. Next, let's look at our second big update. A new agent in Google Flow. Until today, Flow could only execute one prompt at a time. Now, your agent can take multiple actions all at once. Starting with just a single image, I can ask the agent to help me find the best camera angles for this scene.

It analyzes what's happening in the image, concepts, the most compelling angles, and then boom, a single image becomes 16 unique videos. The agent can also handle large-scale edits like transforming all of the scenes from early morning to late at night. Its understanding of context is precise. The desert sky goes completely dark and the headlights turn on, illuminating the dust. It's a true collaborator helping you create and edit at scale.

Our next update is Flow Tools. Now you can vibe code any creative tool you could think of right in Flow. Custombuilt by you for your unique creative process like designing video effects, handdrawn animations, or layering text. You can start building, sharing, and remixing tools. Today, Google Flow Music brings the same creative control to help artists create original songs. For months, one of our teammates had a piano riff in his head. Let's listen to his original recording.

It's a cool foundation, but he wanted to turn it into a demo to guide his band. So, he recorded his piano into flow music and prompted it for an R&B direction with a female vocal to inspire his band singer. Let's take a listen. Jackie on display. Tell me why would I try to fly.

Now, this isn't his final track, but it helped his band decide what to record next. And all of these features in Google, all these new features in Google Flow and Google Flow music are available today. >> Today, I'm excited to announce that our first audio glasses will arrive this fall. They are designed to give you all day help with Gemini that is spoken into your ear privately rather than shown on a display.

And these glasses let you stay handsfree and heads up for things like listening to music, taking photos, making

calls, or tapping into your phone apps, all without reaching for your pocket. And yes, they're going to pair with both Android and iOS devices. These are the first two designs of a bigger collection that's coming this fall. Okay, you know what time it is now. Who's up for a live demo?

No pulling out my phone. All I have to do now is ask. >> Hey, can you navigate me to the place I met my friend Gianna at last week? >> Hey Nisha, I've set your route to the Redwood Grove Nature Preserve from last week's hike. Want to add a stop on the way to grab your afternoon cold brew? >> Yes, Gemini. I would love that. >> Okay, I'm starting walking navigation with a stop at Koopa Cafe. Turn around and head towards Bill Graham Parkway.

>> With glasses, you can allow maps to further understand your context and what's in front of you. So, you get detailed directions like the coffee shop is coming up on your right. Speaking of coffee, Nisha, do you want to show how Gemini intelligence can allow you to order that cold brew up ahead? >> That's a really good idea. Gemini, can you actually put my usual order in at that coffee shop we just talked about? >> Sure. I'll order you a nitro cold brew for pickup from Koopa Cafe on Door Dash.

>> So on the screen behind me, you can actually see Nisha's phone that's in her pocket. Gemini is able to launch apps like Door Dash, then click through all the different option screens automatically to order her coffee. And in a moment, it's going to be ready for uh Nisha's confirmation. >> I have prepared your order for the Nitro Cold Brew from Koopa Cafe. Would you like to confirm? >> Yes, please. And add a tip for 20%, too.

Hey, Gemini. Any important messages I miss? >> Yes, your family group chat decided to meet for dinner at 7 today. And there is a reminder to say your daily affirmations. You are strong. You are those affirmations a bit later, but why don't you add that dinner to my calendar? >> Sure, I'll add an event for family dinner at 700 p.m. to your calendar right after your team celebration. >> So, we saw how glasses work on your phone, but what about when they're attached to your watch to give you a glanceable display?

>> It's not a Google IO without an audience selfie, but we're going to do it with a twist this time. Can everyone strike their favorite pose? Gemini, take a photo of this amazing audience, but turn it into a cartoon and add a big blimp in the sky that says Google IO 2026 on it in fun colors. >> Everyone strike a pose. Okay, if this goes well, everyone go bananas in the audience because Nano Banana Unglasses is just awesome. And just in a few seconds, you'll even see that seamless preview on her watch. Any second now. Drum roll.

Here we go. AGI is now on the horizon and it will be the most profound and impactful technology ever invented. If built right, it could propel human progress and flourishing beyond our imaginations. We're in a moment of immense promise, but also enormous responsibility. It's important that we are clear-eyed about the potential challenges and use all the tools at our disposal to ensure the safety of our agentic systems and ultimately AGI itself.

One area of risk that has gained a lot of attention recently is cyber security. Google has invested in this area for decades and we're bringing our frontier capabilities and deep expertise to help secure the world's code bases. We

have tools like our code security agent code mender which automatically finds and fixes critical software vulnerabilities. Today we're inviting a select group of experts to test a new code API and we'll be launching it more broadly soon. I'm excited to announce Gemini for Science, which brings together a number of powerful AI tools to help accelerate research.

Gemini can already assist in solving complex problems, but our new labs prototypes streamline daily scientific tasks. Whether it's staying on top of newly published papers, transforming research goals into usable code, or generating new hypothesis. Another powerful tool for science is simulation. AI simulations are going to be critical for understanding and predicting dynamic systems that are simply too complex to model directly today. An amazing example of this is Alpha Earth Foundations. It's the closest thing we have to a digital twin of the planet that could help address problems like deforestation and food security.

I've always believed the number one application of AI should be to improve human health. At Isomorphic Labs, we're modeling molecular interactions to massively accelerate the development of new medicines supported by leading industry partners. We're now in pre-clinical stage with multiple projects including potential treatments for immune disorders and cancer. Our mission our mission is to reimagine the drug discovery process with the goal of one day solving all disease.

Something that would have seemed impossible just a few years ago, but I truly believe is now within reach. Google's cutting edge research and products will help unlock AGI's incredible potential for the benefit of the entire world. When we look back at this time, I think we will realize that we were standing in the foothills of the singularity. It will be a profound moment for humanity. This technology will be a force multiplier for human ingenuity and usher in a new golden age of scientific discovery and progress, improving the lives of everyone, everywhere.

We look forward to building the future with all of you. Thank you and enjoy the rest of Google IO.