

This transcript is provided for the convenience of investors only, for a full recording please see the [Q3 2026 Earnings Call webcast](#).

Operator: Welcome, everyone.

Thank you for standing by for the Alphabet Third Quarter 2026 Earnings Conference Call.

At this time, all participants are in a listen-only mode. After the speaker presentation, there will be a question-and-answer session. To ask a question during the session, you will need to press *1 on your telephone.

I would now like to hand the conference over to your speaker today, Jim Friedland, Head of Investor Relations. Please go ahead.

Jim Friedland, Head of Investor Relations: Thank you.

Good afternoon, everyone, and welcome to Alphabet's Third Quarter 2026 Earnings Conference Call. With us today are Sundar Pichai, Philipp Schindler, Thomas Kurian, and Anat Ashkenazi.

Now I'll quickly cover the Safe Harbor.

Some of the statements that we make today regarding our business, operations, and financial performance may be considered forward-looking. Such statements are based on current expectations and assumptions that are subject to a number of risks and uncertainties. Actual results could differ materially. Please refer to our Forms 10-K and 10-Q, including the risk factors. We undertake no obligation to update any forward-looking statement.

During this call, we will present both GAAP and non-GAAP financial measures. A reconciliation of non-GAAP to GAAP measures is included in today's earnings press release, which is distributed and available to the public through our Investor Relations website located at [abc.xyz/investor](#).

Our comments will be on year-over-year comparisons unless we state otherwise.

And now I'll turn the call over to Sundar.

Sundar Pichai, CEO, Alphabet and Google: Thanks, Jim. And hello, everyone.

Q3 was a strong quarter for Alphabet. Our deep full-stack approach is driving momentum across the business with many opportunities ahead.

Let me give you a quick snapshot of the quarter.

Alphabet revenues grew 25% year-over-year. We're now at a half-trillion-dollar annual run rate. That's remarkable growth.

Search has strong momentum. Our AI investments are driving increased usage and overall query growth, including commercial queries. In fact, our year-over-year query growth rate is the highest it's been in the last five years.

Google Cloud revenues grew 103%. Our AI infrastructure and solutions are the major drivers for cloud growth and our customers see us as the leading provider here, so there's huge demand. We are working hard to bring more capacity online to meet this long-term demand. At the same time, we're seeing strong demand for our TPU systems across a wide range of customers, and we expect that to continue as we scale. I'm excited to welcome Thomas to the call and he'll share more about Cloud's momentum.

On our AI models, customers are using our models more than ever. Our first-party model APIs are now processing nearly 32 billion tokens per minute, 40% more than last quarter and more than four times what we processed a year ago.

And with Gemini 4 Argon and now Carbon, which we announced today, we're excited to be leading at the frontier.

The Gemini App is now at one billion users and is the fastest-growing product in Google's history.

Lastly, we've surpassed 150 million YouTube Music and Premium subscribers globally, including trials.

Today, I'll talk more about our AI progress, our global consumer platforms led by Search and YouTube, then our long-term Bets.

Starting with AI, highlighted by our Gemini 4 series of frontier models.

We're seeing strong demand and usage for Argon which we launched two weeks ago. In fact, on our developer platforms, usage in the first week ramped up three and a half times faster than Gemini 3, even with capacity constraints.

Argon is proving to do really well in areas like finance, legal, cybersecurity, and science.

Developers and Cloud customers are already using it for things like untangling legacy codebases, accelerating scientific research, robotics, discovering product improvements, and domain-specific tasks requiring in-depth reasoning.

And at Google, just in the last few weeks, teams across Chrome, Ads, and Search have used Argon to write and optimize code much faster, shipping 35% more code with 30% fewer errors.

Argon is just the start of the Gemini 4 series.

Today we announced Carbon. It's our strongest frontier model with state-of-the-art cybersecurity and long horizon capabilities. It ranks number one on the recently updated Artificial Analysis Intelligence Index at 59.

Carbon outperforms Argon across almost every domain, like full-stack web development, multimodal understanding, complex knowledge work, and scientific reasoning. We're rolling Carbon out first to a subset of trusted cyber defenders before we expand access; and it's being tested by the U.S. government.

Beyond our Gemini 4 series, yesterday we announced Genie 4, our most advanced world model yet that simulates safe, realistic, interactive worlds with native audio in real time. This is an important building block for training robots and agents.

This rapid progress underscores our incredible model momentum.

We're at the point where AI is helping us increase our training and launch velocity. I would expect the next twelve months to see continued acceleration because of RSI. We have more models coming soon, including more efficient versions to help developers optimize for speed and cost and to deploy across our billions of users.

Let's turn to Search, which is bringing the benefits of AI to more people than any other product in the world. Just like our shift to mobile over a decade ago, we are focused on building the best Search user experience, followed by monetization.

We are building one seamless experience across AI Overviews, AI Mode, and traditional search. And people are using our AI features more, searching in entirely new ways. Their queries are longer, personalized and multimodal. Queries in AI Mode are on average three times longer than traditional searches, and one in five AI Mode searches globally use voice, images or documents. We've now crossed over 50 billion visual searches per month.

One area I'm excited about is helping people get things done, be it booking reservations, filling in forms, scheduling appointments, or making purchases. As AI becomes more personal and starts taking action for people and businesses, trust and security matter more than ever. People already put their trust in Google to get things done in their daily lives; and our leadership here remains an important differentiator.

We focused first on getting these agentic experiences right with the Gemini agent for enterprise, because the bar for security, safety, and governance is extremely high.

As a next step, we're bringing these features to Search. And also, of course, to our billion users on the Gemini app. People are already using the Gemini app to do more than ask questions; they're using it to do things, including complex, multi-step agentic queries. Our Gemini 4 series has been a really big step forward here and has already

been handling these complex tasks really well.

Agents will operate as a unifying layer across our products, so you'll see us extend these agentic experiences beyond Search and the Gemini app to Chrome, Android and our Workspace products as well.

On to YouTube, which continues to lead in the living room as the number one streaming platform in the U.S. in terms of watch time. It's also seeing good subscriber growth, as I mentioned.

At Made on YouTube last month, we announced more than 30 products and features for creators and viewers. One great example is custom feeds, where you can simply pick a topic you love and get a tailored video home feed based on that topic. Philipp will say more on YouTube.

On to our long-term bets. Waymo is on a great trajectory. We've now driven more than 300 million fully autonomous miles. Waymo now operates in 15 major U.S. cities and just recently opened up service to the public in San Antonio, San Diego, and Tampa, with planned expansions into Canada, Germany, Singapore, and Spain.

We're also seeing strong progress across our other long-term initiatives. Isomorphic is using Argon to support drug design research. We demonstrated "magic state cultivation" for the first time in quantum computing, and launched our TPUs into space with Project Suncatcher.

So, it's been a busy quarter and there's much more to come. I want to thank our employees for their work this quarter.

And over to you, Philipp.

Philipp Schindler, SVP and CBO, Google: Thanks, Sundar, and hello, everyone.

I'll highlight our momentum across Search, YouTube, and Partnerships.

We continue to accelerate the deployment of Gemini models across our entire ads infrastructure to boost performance in the three areas I've highlighted before: Ads Quality, Advertiser Tools, and AI User Experiences.

First, Ads Quality.

The integration of Gemini models into our ads infrastructure is improving how we decode user intent and drive ads relevance. For example, in Q3, we made the biggest update to product listing ads in a decade, adding new AI capabilities to help users make more confident purchase decisions. Already, we've seen satisfaction increase by 19% and dissatisfaction drop more than half.

We are also seeing AI models enable better predictions of clicks and conversions, stronger overall performance, and deeper personalization of ads. For instance, through several innovations in personalization modeling, we delivered 18% incremental advertiser value and an 8% increase in user sentiment on Discover Ads in 2026.

Overall, AI is accelerating our ads development and drives the majority of our product launches.

Second: Advertiser Tools.

We are ensuring every advertiser, from global brands to small businesses, can seamlessly participate in our AI surfaces. AI Max and Performance Max continue to see strong adoption. To give marketers even more control, we're now introducing a way to test different budgets and ROI targets across multiple Search campaigns in a single AI Max A/B test.

To drive better performance, we integrated our AI creative tools directly into AI Max. In Q3 alone, Gemini models helped advertisers generate over 1.5 billion campaign assets, up over 7x year-over-year. For example, by using AI Max to capture long-form and multilingual queries, Telstra saw a 52% increase in conversions and a 36% reduction in cost per order, showing up in key moments however their customers search.

Third: AI User Experiences.

Sundar mentioned the distinct behavioral changes we are seeing as users ask entirely new categories of questions.

We think of Search holistically, a unified product spanning traditional results and AI powered Search.

On AI Mode, many of these queries are upper-funnel queries, allowing us to capture new high-value commercial intent. In addition, the majority of AI Mode queries are incremental, creating more monetization opportunities over time.

Additionally, our user research shows that people are comfortable with ads on these surfaces and find them highly relevant. Click-through rates are already similar between sponsored and organic product list items within AI Mode.

Based on these collective results, we are now scaling high-quality ads within AI Overviews and AI Mode to more users and queries. This includes taking our monetization footprint global for the first time. Today we are announcing the expansion of ads in AI Mode into international markets, alongside multi-turn ads in AI Mode in the U.S.

As always, our approach is deliberate -- prioritizing the user experience and engagement first, and then scaling the formats that drive proven ROI for advertisers.

That discipline is what gives us a strong, durable foundation for the future.

As Sundar mentioned, we are in the early stages of agentic experiences across our surfaces, and excited about the commercial opportunities ahead. The Universal Commerce Protocol, UCP, is being rapidly adopted by retailers. Today, UCP enables direct checkout for hundreds of thousands of brands across Google, including Chewy, Nike, Tapestry, and many more. We are combining this with the Agent Payments Protocol as we build for the future of commerce.

Our focus spans two fronts: unlocking agentic capabilities that act on behalf of the consumer, and designing practical tools that give shoppers confidence and control.

We launched Universal Cart this quarter to support that second track, mirroring how people actually shop today across Google. This allows retailers to maintain their relationships with customers, while helping to capture sales and avoid abandoned carts.

Moving to YouTube. We view YouTube as a flywheel of creators, viewers and advertisers.

First Creators. Every day, more than 20 million videos are being uploaded. We're making creation easier with new generative AI tools in YouTube Studio, plus a conversational editing assistant in Shorts and YouTube Create.

Looking at viewership. At Made on YouTube, we introduced Shorts series where creators can structure their Shorts into seasons and episodes. Sequential playback makes it seamless for viewers to discover new Shorts and keep watching.

The third part is monetization, starting with ads. Our focus remains on delivering value across the entire funnel. In Q3, Living Room accounted for over one-third of total YouTube ads revenues. Small and medium advertisers continue to adopt Demand Gen driving Direct Response.

On YouTube, Wayfair's Demand Gen campaigns drove a 2.3x higher return on Ad Spend compared with traditional Video Reach Campaigns.

Our YouTube Creator Partnerships program connects brands with over three million creators to drive highly effective campaigns.

Brands like Lexus continue to partner with creators like Colin and Samir, driving an 11% lift in brand consideration.

More than 1.3 million channels are enrolled in YouTube shopping globally. Building on this momentum, we recently announced that Amazon is joining YouTube Shopping. U.S. creators can now tag Amazon products directly in their videos, Shorts, and livestreams, driving direct conversions.

Importantly, our subscription ecosystem continues to complement our advertising business. Subscriptions continue to grow faster than Ads, led by YouTube Music Premium. YouTube Premium has more than 150 million subscribers. With Premium and Premium Lite, viewers have choice. As more people opt into premium experiences, it's a win for viewers, creators and our business. Subscriptions help fund the creator ecosystem and drive the platform forward.

Finally, Strategic Partnerships. Businesses are choosing Google because we bring together Cloud, AI, Ads, and consumer platforms to accelerate their transformation. We expanded our partnership with American Express to support the ongoing modernization of their contact centers and collaborate on agentic commerce across Search and Wallet. And Unity signed an expanded, multi-year Cloud deal to use our full AI stack to power their next-generation game engine.

Our commercial priorities are clear: expanding advertiser ROI, systematically advancing our AI monetization efficiency, and scaling our ecosystem partnerships to support the next era of discovery.

In closing, I want to thank Googlers everywhere for their incredible execution, and our customers and partners for their continued trust.

Thomas, over to you.

Thomas Kurian, SVP and CEO, Google Cloud: Our strong momentum is driven by our integrated AI portfolio of chips, models, data, security, and agent platforms, all designed to work together.

We have strong diversification across our product lines, industries, and customers.

Google Cloud now has 18 product lines with more than \$1 billion each in annual revenue, and our growth this quarter was driven by strong growth in Google Cloud platform, with strength in both AI and Core GCP. TPU System sales were also a meaningful contributor.

Over the past 12 months, more than 2,600 customers have used at least 100 billion tokens; and nearly 75, at least 10 trillion tokens on Gemini Enterprise.

Nearly 90% of the Fortune 100 use Gemini Enterprise, and customers who used Gemini Enterprise in Q3 2025 more than quadrupled their spend with us year-over-year.

One of the important ways in which customers are using Gemini is for knowledge work. Across all of our AI services, we have sold more than 100 million seats for knowledge work.

Gemini Enterprise now includes the Gemini agent that is deeply integrated across our entire Cloud portfolio. It can research and answer questions, write and run code, and automate process flows, with one agent, one API, and a single prompt window.

It gives companies strong security and cost controls, and the widest choice of frontier models.

Customers include:

BNP Paribas, who is deploying Gemini Enterprise to more than 65,000 employees to streamline workflows, while maintaining strict enterprise-grade security and data governance.

Tata Steel, who deployed more than 300 specialized AI agents in just nine months to enhance manufacturing safety, optimize supply chains, and cut service times for customers by 50%.

Ulta Beauty, who built the Ulta AI shopping assistant, resulting in a 3x higher conversion rate than traditional browsing on their site.

Gemini is also driving acceleration for our database, analytics, and cybersecurity products. Gemini can now help users find, analyze and visualize information from structured and unstructured data.

Gemini uses our Knowledge Catalog to find the right sources and definitions of information, improving accuracy, lowering token costs, and making responses more reliable.

Gemini uses our Lakehouse to access information from multiple systems with zero copy.

To respond to agentic analysis, BigQuery improved query speed over 40% year-over-year, while reducing query processing cost by over 50% year-over-year.

Customers include:

Deutsche Telekom, which uses BigQuery, Cloud Storage, and Knowledge Catalog to build a unified data lakehouse across 40-plus legacy on-premise systems.

Etsy, which migrated three petabytes of data to a unified lakehouse. By using BigQuery and Iceberg, they now perform critical data joins up to 60% faster.

The growing risk of cybersecurity threats is also driving strong demand for our AI-powered security platform. Our Gemini Flash Cyber models are frontier models that automate vulnerability discovery across 20 programming languages and patch 2.6 times more reliably than other models. And there's even better performance to come as

we roll out Argon and Carbon.

Our security platform integrates threat intelligence, cyber-response prioritization, automated scanning, code remediation, and monitoring.

Gemini in Google Threat Intelligence autonomously analyzes millions of external events every day with 98% accuracy.

And more than 95% of Wiz customers are adopting AI-powered features, and growth continues.

These security tools protect critical infrastructure, including financial institutions like Banorte and Jack Henry; and telecommunications providers such as AT&T.

All of this is powered by our leading AI infrastructure, spanning our TPU systems, Axion CPUs, the latest NVIDIA GPUs, and our Agent Compute Engine, which is specifically designed for AI Agents. TPU systems are specialized systems built for AI, combining hardware, compute storage and networking, with systems components and optimized software for training and inference.

TPU 8t and 8i offer up to 2.7 times better training performance and 80% better inference price performance respectively, versus our previous generation of TPUs. To make it easy for developers to port workloads to TPUs, we recently worked closely with Meta to deliver Torch TPU, which allows users to train, serve, and deploy Pytorch models on TPUs using the tools they're already familiar with and rely on.

Nearly 100 customers use TPUs today, and we are seeing strong and growing demand from many others. In addition to our work with Anthropic, both directly and via Broadcom, and public reference customers like Apple, who are scaling pre-training of their foundation models on the latest generation of TPUs, we're also working with new frontier model companies, including Mirendil and Physical Intelligence; capital markets firms including Citadel Securities and Jump Trading; universities like Caltech; and public sector service providers, including Cassava and Cirrascale.

Given continued strong and accelerating demand, we've recently entered into partnerships to deliver TPU systems with both Crux AI and Oracle, with more in the works.

Over to you, Anat.

Anat Ashkenazi, SVP and CFO, Alphabet and Google: Thank you, Thomas.

My comments will focus on year-over-year comparisons for the third quarter, unless I state otherwise.

I will start with the results at the Alphabet level and will then cover our segment results.

I'll end with some commentary and our outlook for the fourth quarter and full year 2026.

We had a strong third quarter, delivering our thirteenth consecutive quarter of double-digit revenue growth.

Consolidated revenues were \$127.6 billion, up 25% in both reported and constant currency.

Total cost of revenue was \$51.4 billion, up 24%. TAC was \$16.4 billion, up 10%.

Other Cost of Revenues was \$35.1 billion, up 32%, driven by hardware costs primarily related to the sales of TPU systems, increase in depreciation, content acquisition costs, and technical infrastructure operations costs.

Total operating expenses were up 20% to \$35.6 billion.

R&D expenses increased by 36%, driven by compensation, primarily for investment in AI talent as well as depreciation.

Sales and Marketing expenses were up 22%, due to investments in Gemini and compensation.

And G&A expenses decreased 17%, primarily due to lapping of higher legal expenses that occurred in Q3 last year, partially offset by current year legal charges.

Operating income increased 30% to \$40.6 billion, and operating margin was 31.8%.

Other income and expenses was negative \$8.2 billion, due primarily to unrealized loss from our investment in SpaceX from a decline in its stock price in the third quarter.

Net income and earnings per share decreased due to the decline in OI&E.

We generated strong operating cash flow of \$42.3 billion in the third quarter and \$179.6 billion for the trailing 12 months.

Operating cash flow was negatively impacted by the buildup of inventory for TPU system sales, as well as a \$5.2 billion cash payment for a 2018 EC fine following the conclusion of the appeal process.

CapEx was \$51.4 billion in the third quarter with the vast majority of the spend in technical infrastructure to support our investments in AI. Approximately 65% of our investments in technical infrastructure in Q3 was related to servers and networking equipment, which generally have useful lives of six years. The remaining 35% was for data centers, which have useful lives of seven to 40 years.

We had negative free cash flow of \$9.1 billion in the third quarter, driven by our

investments in CapEx. Free cash flow in the third quarter was also affected by the aforementioned inventory bill to support TPU system sales and payment of the EC fine.

Free cash flow was \$19.7 billion for the trailing 12 months.

We ended the quarter with \$265.1 billion in cash and marketable securities, which includes \$90.4 billion of marketable equity securities. Long-term debt was \$125.1 billion.

We have also entered into multiyear agreements to secure capacity for technical infrastructure and inventory, as well as leases and energy. At the end of the third quarter, we had future purchase commitments of \$924 billion and pre-commencement leases and other obligations of \$152 billion.

These dollar amounts reflect the maximum contractual obligations. However, a number of these agreements have variable pricing and volume flexibility, which could result in lower spend.

And as mentioned in our press release, our Board of Directors declared a quarterly cash dividend on our common stock of 22 cents per share, which is payable in December.

Turning to segment results, Google Services revenues increased 11% to \$96.3 billion, reflecting strong growth in Search and Subscriptions.

Total advertising revenues are up 11%, with growth across all major verticals.

As we stated in our outlook for last quarter, FX was a slight headwind on year-over-year growth rates and had a more significant impact on sequential basis.

Google Search and other advertising revenues increased by 12% to \$63.3 billion, with retail driving the largest contribution, followed by finance.

YouTube advertising revenues increased 9% to \$11.2 billion, with Brand driving the largest contribution and notable strength in Living Room.

Network advertising revenues of \$7.5 billion were up 1%.

Subscription, Platforms & Devices revenues increased 11% this quarter to \$14.3 billion, driven by strong growth in YouTube subscriptions, particularly in YouTube Music and Premium and strong demand for AI plans in Google One.

Google Services operating income increased 17% to \$39.2 billion, and operating margin was 40.7%.

The Google Cloud segment delivered record growth in the third quarter. Cloud

revenues were up 103% to \$30.7 billion, driven by strong growth in GCP, with strength in AI, including both AI infrastructure and AI solutions and Core GCP.

TPU systems sales were also a meaningful contributor. GCP accounted for the vast majority of the revenue generated in Q3. The Cloud revenue growth rate, excluding the impact of TPU systems, was again incredibly strong.

Cloud operating income was \$11.1 billion, more than tripling year-over-year. Operating margin increased from 23.7% in the third quarter last year to 36.2%.

The third-party bridging capacity, which we discussed on our prior call, came online only at the very end of Q3 and, therefore, had limited negative impact to our margin in the quarter.

Google Cloud's backlog was \$561 billion in the third quarter, with the majority of the backlog related to typical GCP contracts. We expect to recognize just over 55% of the total backlog as revenues over the next 24 months.

In Other Bets, revenues were \$372 million, and the operating loss was \$2 billion dollars as we continue to expand Waymo's business and invest in key Other Bets.

In Alphabet-level activities, the operating loss was \$7.7 billion, driven by shared AI R&D expenses.

I'll now turn to our outlook for the remainder of 2026. First, in terms of revenues, at current spot rates, we would expect a slightly larger FX headwind to our consolidated revenues in Q4, relative to the prior quarter. This impact will be seen primarily in Search and YouTube Ads.

In Google Services, we will continue to lap the Search performance acceleration that began in the third quarter last year, which will pressure revenue growth in Q4.

In Google Cloud, we expect TPU system sales to continue to scale in Q4, followed by a meaningful revenue ramp in 2027. With respect to Cloud margins, given the supply-constraint environment, we started utilizing third-party compute capacity at the end of Q3 and will ramp usage in Q4. We expect margin headwinds in the near-term as we utilize this capacity.

Moving to investments, our full-year 2026 CapEx guidance range of \$195-\$205 billion remain unchanged. We continue to expect our CapEx to increase significantly in 2027, and we'll provide further details on our Q4 call.

In terms of expenses, the significant increase in our investment in technical infrastructure will continue to weigh on the P&L in the form of higher depreciation expenses and related data centers' operation costs such as energy.

Before we move to Q&A, let me say a few words on how we think about return on invested capital and free cash flow.

We're in a unique position to be able to monetize our AI investments across many products that reach billions of users and enterprises. Our AI investments in Search are driving our highest query growth in half a decade. Google Cloud has generated tens of billions of dollars in AI revenues year-to-date, and Google One is our fastest-growing subscription product, driven by demand for AI plans.

Throughout this investment cycle, we expect to generate an attractive ROIC that remains meaningfully above our cost of capital and anticipate generating strong cash flow from operations.

Our sizable investments in technical infrastructure and inventory build of TPU systems as we scale that business will put pressure on free cash flows.

Q3 was a strong quarter, highlighted by our innovation and disciplined execution across all areas of the business, including the launch of Gemini 4.

I want to take this opportunity to thank our employees for their contributions to our performance.

We will now open the call for Q&A.

Operator: Thank you. As a reminder, to ask a question, you will need to press *1 on your telephone. To prevent any background noise, we ask that you please mute your line once your question has been stated.

And our first question comes from Doug Anmuth with JP Morgan. Your line is now open.

Douglas Anmuth, JP Morgan: Thanks for taking the questions. One for Sundar and one for Philipp.

Sundar, on a recent podcast, you said that 2027 could be the golden age of productivity for Google because of the new class of models and how they'll integrate with your other products. I was hoping you could elaborate on that and just share more how you think about the opportunity set there ahead.

And then, Philipp, just on query growth, the fastest in five years, but obviously some deceleration in Search growth. Can you just help us understand the disconnect there? Is it simply that you haven't been showing as many ads basically in AI Overviews, in AI Mode, and that's now becoming a bigger percentage of the overall queries? Thanks.

Sundar Pichai, CEO, Alphabet and Google: Thanks, Doug. We definitely started seeing a step change I would say with 3.8 Flash internally, as our models got better at

long horizon tasks, agentic coding, et cetera.

But definitely seeing a huge point of inflection once we have started deploying Argon internally. It is one of the most meaningful internal productivity levers we have seen on software engineering, directly tracking high-quality code that's getting checked in, the rate at which it's getting checked in, reducing number of rollbacks.

But more importantly, people are tackling massive problems, as well as we are beginning to see innovation on the product and the consumer experience side from deploying those models.

But all this happened only about four weeks ago. We are now deploying Carbon internally, and Carbon is a meaningful step up from Argon.

So, as we roll it out and make it all widely available to Googlers with more compute, I expect -- that's what I was referring to, and I'm even more excited since I said that in the podcast, what I would view as a golden year ahead in terms of us being able to be more productive; so, excited for the innovation ahead.

Philipp?

Philipp Schindler, SVP and CBO, Google: Doug, on the second part of your question. In Q3, Search grew 12% year-over-year to \$63 billion, and when you look at the change in growth rate from Q2 to Q3, there were a few drivers here. First, there was FX. It was the single largest driver of Search deceleration relative to Q2.

Second, there were vertical-specific dynamics. In retail, some of our APAC advertisers were impacted by the EU eliminating its de minimis exemption. In media and entertainment, spend normalized after an unseasonably strong Q2 driven by the FIFA World Cup, which had more matches, and therefore more advertising budgets allocated in June than actually in July.

To the next part of your question, it's worth noting that the strong growth in AI Mode queries had an impact because of the lower monetization we were seeing in AI Mode. To be clear, we feel very good about our trajectory here. Our ad experiments have shown promising results and we're now scaling ads within both AI Overviews and AI Mode to more users and more queries; including, as I mentioned, international markets and multiturn ads in the U.S.

Douglas Anmuth, JP Morgan: Thank you.

Operator: Our next question comes from the line of Brian Nowak with Morgan Stanley. Your line is now open.

Brian Nowak, Morgan Stanley: Thanks for taking my questions. I have two. Maybe let me start with agentic. You rolled out Spark to some of the users throughout the year. I

know, Sundar, you've talked a lot about shipping more agentic products to the broader base of users. Can you just help us understand a little bit some of the key hurdles you have to overcome technologically to ship out agentic products to the full base of users? Is it model efficiency? Is it capacity? Is the technologies are ready? What do we have to do technologically to do that?

And the second one, to Doug's question on Search. Anat, can you help us better understand the Search performance comparison factor that you talked about impacting Search, and how do we think about that as we kind of go into 4Q-1Q, just kind of looking at the comparisons?

Sundar Pichai, CEO, Alphabet and Google: Thanks, Brian. First of all, on agents, what an exciting and dynamic area. I think when I look at our new models and we see the capabilities of the kind of complex tasks they can take in an automated way and execute, it was very, very high quality. It's clear that the agentic capabilities we have all envisioned are going to be real.

And so, I couldn't be more excited to bring the next generation of agentic experiences to our users.

To your question, I would say a few things. I think the Gemini 4 model series has been important. We are working on distilled versions of our Carbon models, so that we can deploy it across a billion-user scale efficiently, both from a compute, from a latency standpoint, and from a quality standpoint. So that's been an important area of investment.

Definitely making sure you need more than a capable model at scale, you need a robust harness, you need a highly secure runtime, and this is where GCP comes in, being able to give each and every one of our users a secure runtime with 24/7 availability in the cloud, the right ecosystem of skills and connectors. So, all of that is the foundation.

We have been hard at work, and we decided to roll it out to Enterprise first, which we announced with our Gemini agent for Enterprise. Essentially, because they have the highest bar for security governance and controls. It's been well received, and we are in the process of expanding that.

Later this year, we will bring that same experience to consumers, obviously starting with the one billion users in our Gemini app.

We already see people in the app moving beyond Chat to what are task and agent-like experiences. And so, the trust users have in our products, the fact that all their context is right there, I think sets us up well.

But we are really focused as a company on delivering that unified experience across our products, and you should expect to see a lot of progress later this year.

Anat Ashkenazi, SVP and CFO, Alphabet and Google: Thanks. And on the second question with regards to my comments, I believe you were referring to the outlook with regards to Search. A couple of dynamics that I've mentioned. One is on sequential, related to impact of FX headwinds, which we saw this quarter as well versus the previous quarter. So where we are today with current spot rates, we expect that this will be somewhat of a headwind to Q4 results.

And then when we look at the year-over-year comparison, this quarter we started lapping the acceleration we saw in Search growth in 2025, the back end of 2025. We left about three points of growth versus Q3. And then we expect to see that continue in Q4, a couple points of lapping that growth acceleration that we saw at the end of 2025.

Brian Nowak, Morgan Stanley: Great. Thank you both.

Operator: Our next question comes from the line of Eric Sheridan with Goldman Sachs. Your line is now open.

Eric Sheridan, Goldman Sachs: Thanks so much for taking the questions. Based on some of the comments you made recently, Sundar, on compute constraints, can you refresh for us your key priorities for the application of compute as it becomes available across the organization?

And the second part, maybe for you or for Thomas, would be within Google Cloud specifically, how does your view continue to evolve with respect to the application compute, to cloud customer needs relative to external TPU? Thanks so much.

Sundar Pichai, CEO, Alphabet and Google: Thanks, Eric. I'll start, and then Thomas will answer the second half.

Definitely compute allocation is an area where we are spending a lot of time on, because of the forward planning that's needed. I would start with saying we feel fortunate to have an extraordinary set of opportunities ahead of us. We are operating a scaled consumer business that reaches billions of people, a hyperscale cloud business with a lot of momentum, and of course a range of other highly important areas across Google and Alphabet.

We always take long-range planning, and so each of our business areas, some at long-range plans with a very disciplined ROIC framework. And so I would broadly think of all of this as a long-range planning process, not a short-term allocation exercise.

Obviously, our AI models underlie everything we do, from Search and YouTube, to Cloud, to Waymo, and so on. So it's the foundation, and we view it as a highly leveraged approach. The full-stack approach drives a lot of efficiency, as well as gives us a lot of optionality across a wide variety of opportunities.

So, obviously that is the foundation, and we plan for that for the long-term, too. And about that, as I said, we really look at our businesses, and really take a disciplined long-range planning approach.

To us, obviously we are seeing very, very robust demand signals across a wide variety of areas, and so we view it as an extraordinary set of opportunities, and we are investing for that moment.

Thomas?

Thomas Kurian, SVP and CEO, Google Cloud: Eric, on the demand in Cloud, we are always managing our portfolio to ensure we have a breadth of customers; and we have different types of customers from digital natives, AI labs, traditional enterprises, financial and capital markets customers, so it's a broad set of customers.

Many of them have either started with using TPU systems, for example, in our Cloud. We have over a hundred significant TPU customers who start in our Cloud, they learn to use TPUs; and remember, a TPU system is a combination of hardware, system software, compilers, the whole thing, assembled as a system, including optical networking and other things.

We have taken a very deliberate approach on how we're bringing TPUs to customers. The first one is to build the right silicon, so we have 8t for training, 8i for inference. They both offer substantial price performance and performance improvements relative to prior generations.

Second, we've optimized the software. The optimization of the software is not just to make the systems run better and to give you great performance, but also to make portability seamless.

Recently, we introduced, for example, Pytorch, which opens up a much larger market for us by optimizing Pytorch on the TPU.

Third, we are broadening distribution. So, in the past, you could use TPUs in our Cloud. Now we make, for customers who want TPU systems in their data centers, we're making TPU systems available in their data centers.

We're also broadening availability to other providers. For example, we announced a partnership with Oracle, and we've launched, along with Blackstone, an agreement for Crux AI, a joint venture, to offer customers TPU systems.

We are seeing extraordinary demand, and strong growth from both AI labs, capital markets, financial service institutions, and Enterprises. And if you imagine these hundred-plus customers growing to many more, and also scaling their use into TPU systems, we see strong demand over the future.

Operator: Our next question comes from Ross Sandler with Barclays. Your line is now open.

Ross Sandler, Barclays: Thanks. I want to just stick with the TPU resale business topic you just mentioned.

So, Anat, last quarter you kind of said that the revenue will be recognized on system sales, I believe. So if a gigawatt of TPU full system costs around \$35 billion for the compute and all the components around it, what percent of that are we actually recognizing as revenue? And what kind of flow-through does that have in terms of margin profile relative to where the Google Cloud segment currently sits?

And then, I guess the second part of the question would be, as you help companies like Anthropic bring on more of their own kind of first-party infrastructure, what does that mean for the CapEx intensity of Alphabet as you start to kind of hand off some of the future CapEx to the AI labs and may not require as much CapEx intensity on your own -- in your own data centers? Thoughts on that would be great.

Anat Ashkenazi, SVP and CFO, Alphabet and Google: Let me combine these two questions. First, obviously, we don't provide specific financials at the product level, but let me walk through how to think about the financials associated with the TPU system sales.

And obviously, just as a reminder, and Thomas can elaborate on that as well, a TPU system is not just a chip, but it's a specialized system that combines multiple things, including hardware, compute, storage, networking. So we are selling much more than just a chip. Software, training, inference -- there is a significant added value that's being provided to the customer that's purchasing a TPU system.

And the way this works, and this is linked to your second portion of the question, when we sign an agreement with a customer to sell TPU systems, once we sign that agreement, you'll first see that showing up in our backlog. And then we start building inventory. So think about this from a CapEx versus cash perspective.

We build the inventory ahead of getting the revenue stream in. So I think I've alluded to this on the previous call. You see inventory balances go up, where previously inventory balances primarily were some of our hardware products. Now those TPU systems that are not yet sold will be showing up; so you see cash going out, and we're building up the inventory base.

And then as we start delivering and billing customers for these TPU systems, then you see cash coming in. So it's not the traditional kind of CapEx as you're thinking about this, and I think that's the essence of your question, but rather a different cycle of cash flow generation.

Now, given that we are in the process of building up and ramping up that business,

you're seeing that cash needs in our cash flow statements, and you see those inventory builds, to date we've had some system sales recorded as revenue, but expect more of it to come in through 2027.

Operator: Our next question comes from the line of Justin Post with Bank of America. Your line is now open.

Justin Post, BAML: Great, thank you. I have some questions on the agentic transition. You did launch the Universal Shopping Cart. Clearly, people are liking and using agentic use cases.

As more of that happens, how do you plan to manage the Search transition to agentic use? Do you think that could be a headwind?

And second, do you think agentic could actually drive more monetization for your traffic over the long-term?

And then I'll go again at TPU margins. We're wondering, do you think this is a really good business for you, and maybe why? If you can't give us the margins, why you think this is a good way to use all this capacity. Thank you.

Philipp Schindler, SVP and CBO, Google: Justin, we're in the early days of agentic, but maybe directly to the second part of your question, we view agentic experiences as clearly additive. Our focus is building trusted, secure superscale experiences that users love; and quite frankly, on agentic commerce foundational blocks, like you mentioned that universal commerce, the cart, the agents payments protocols, and so on.

When users shop, that's important, there's a part of the journey they don't like, the grunt work, which is frankly fantastic to automate with AI. But we also want to enhance, not just automate away the discovery experience that people frankly love in so many cases. Users like choice and they want to consider things way beyond price, like brand loyalty, customer service, and so on.

So we're really architecting these agentic workflows in a way that we can intentionally design the time and the space for users to explore those choices.

So, simple example, an agent could respond: Here is the insurance product that we found, and here's another strong option, one of which is a higher premium, maybe a lower deductible, and they're currently offering a very special discount.

So, given that there's almost never only one perfect option for a decision, we're actually confident that we can actually provide really interesting, compelling advertising models for agentic use cases. And obviously, you can imagine going beyond this with Subscriptions and other payment methods.

But I'm really actually excited about the advertising opportunities in an agentic world.

Sundar Pichai, CEO, Alphabet and Google: Overall, I would add and say -- look, I would view it as an expansion opportunity, right? Any time, just like a few years ago, being able to incorporate AI chat use cases in Search allowed people to ask more and do more in Search, which is what is playing out in our usage.

I think agents would, again, be an opportunity to expand the types of problems we solve for our users.

I think people always underestimate all the components you need to deliver these complex experiences for users. And Search on the surface looks simple, but there's extraordinary deep levels of investment and things we have built over time, like Knowledge Graph being an example.

And so our understanding of the world, et cetera.

So, I think I'm excited about agents, and I would view it as a new frontier, which brings in a new set of use cases to Search, and so we are excited to work through that.

Thomas Kurian, SVP and CEO, Google Cloud: On TPU systems and margins, we don't break out margins for individual products or infrastructure components. However, we see several benefits from designing and manufacturing TPU systems.

So first of all, to remind you, the TPU system is not just a chip, but it's a complete system.

Given the price performance benefits of TPU systems relative to competitors, growing demand for accelerated compute and the scarcity of accelerated compute, we are very focused on continuing to improve margins as we scale the business.

Two additional points.

First of all, when Anat mentioned that we procure components to build inventory, we only do that after we have signed contracts, so all of the inventory for TPU systems are associated with committed contracts.

Additionally, by putting TPU systems in third-party facilities and in customer data centers, it allows us, one, to broaden our addressable markets. For some industries, they have data in their premise that they cannot move to a cloud, and it allows us to bring accelerated compute to where the data sits.

Additionally, it also offsets the need for us to have to prebuild facilities to host these TPU systems.

Operator: Our next question comes from the line of Mark Shmulik with AllianceBernstein. Your line is now open.

Mark Shmulik, Bernstein: Thanks for taking the questions. I'm going to be on the same couple of topics. But Philipp, just on agentic, if you've kind of just taken a look at the early more agentic style use cases on Gemini AI Mode, and kind of observing some of the 3P integrations with Google properties, anything that's surprised you or changed your view? I know we've talked a lot about commerce, but perhaps like on consumer productivity as well.

And then either Thomas or Anat, just on Cloud growth, is there anything you could share around the rate of growth of Google Cloud excluding the TPU system sales? And just how much of that has been constrained by capacity? As some of that 3P capacity comes online in Q4, is it fair to assume an acceleration of that Google Cloud business? Thank you.

Sundar Pichai, CEO, Alphabet and Google: Maybe I'll take the first part on what we are observing on these early agentic style use cases.

We are definitely seeing it. Just recently in the Gemini app, we have added a Task tab, where people are actually typing in agentic queries, and so we can observe that.

Look, in all of this, users really want to be delighted, so they effectively are looking for -- they have a high bar. What encourages repeat use is your ability to more handle these longer tasks and actually deliver a satisfactory outcome at the end.

Just like having seen this in other areas including coding et cetera, when you're doing it at 70%, it looks like it doesn't work, until you cross a certain threshold, and it really takes off.

And so, we are closely learning from all of it, identifying gaps, laying all the foundations I talked about earlier, and I'm confident it's an area where you will see a lot of progress across our surfaces as we go through later this year and '27.

The one thing, as we transition to TPUs, one other comment I would make on the TPU systems is, look, in the semiconductor industry, economies of scale matter. So I would look at our ability to do TPUs, both for deploying in our data centers and TPU systems in others' data centers, as the ability to scale up the product, which directly translates into better economics.

Thomas Kurian, SVP and CEO, Google Cloud: To address the question on the rate of growth of Cloud, our Google Cloud segment delivered record growth in the third quarter with Cloud revenues up 103% to \$30.7 billion.

The strongest growth contributor was Google Cloud Platform, with strength in both AI and Core GCP. So GCP accounted for the vast majority of revenue generated in Q3. And Cloud revenue growth, excluding the impact of TPU system sales, was again very, very strong.

Demand was also strong across our product lines. So we see three things going on.

One, as people use Gemini, we see revenue growth from Gemini itself; but we also see that as people use Gemini, it drives growth in other parts of our portfolio. For instance, as they build agents, they use compute to perform a task, and so they buy compute. They use storage to store memories of the agent. Database to serve up what's called context for the agent to reason on. And obviously given all the cybersecurity threats, we are seeing strong growth in our cybersecurity portfolio.

So we are confident about our growth and the diversity in our growth. We're also seeing new customers buy into GCP and exhilarating growth in new customer adoptions, so we also see that as a growth factor for us.

And we're always -- at any point in time, we could obviously serve a lot more demand with capacity, but we are managing all of these parameters and growth carefully.

Operator: And our last question comes from Michael Nathanson with MoffettNathanson. Your line is now open.

Michael Nathanson, MoffettNathanson: Thanks. I have two for Sundar.

In an earlier answer to Eric's question about TPU and Cloud, we heard a willingness to work with some of your competitors, or companies that we may view as competitors, in models and cloud. So I wonder how you think about the philosophy here, about working with people that maybe are competing with you and why the strategy is an opportunity, not a risk, longer term makes sense.

And then secondly, on agentic consumer behavior, one of the debates is around safety and trust. Can you share with us what you've learned so far about safety and trust? Are those variables that matter to consumers? And how do you think -- if those variables matter, what are you doing to protect consumer trust and privacy as you build out agentic consumer tools? Thanks.

Sundar Pichai, CEO, Alphabet and Google: Thanks, Michael. Great questions.

On the first one, look, I've long viewed what we build as categories of things which we build as products, and then there are categories of things which we build as platforms. You need to have a clear view on which is which. And when you're building things as platforms, you really work hard to make the platform great and make sure it works for everyone.

We have a long history of operating many platforms, be it a browser, be it an operating system, be it an extraordinarily successful cloud platform. And in all of that, we have worked hard to provide it to all players and be a great customer.

You've seen us work on our models with other companies, including Apple, with whom we compete in many, many areas.

Similarly, we view our TPU systems as an important platform we are building. We want to be world class in how we build it. That means you're putting it against market forces, testing it, and competing for people who would have the most demanding workloads on top of it. And I think that's what makes sure those teams are continuing to be at the frontier. So that's how we think about it.

And over time, I think that always ends up in expanding opportunities, and we have been able to do well taking that approach for a very long period of time.

Second, on agentic behavior. To the earlier question -- and I think I should have emphasized, people really -- when people are approaching these agentic experiences, we clearly see people are concerned about safety and trust. I mean, it is a very, very high factor in how they want to use these things.

Over decades, the trust we have built by delivering these experiences in a safe way across many consumer products with scale in billions of users, as well as the security investments we are putting in end-to-end as we are building our models and agents, deploying them both with extensive testing, the kind of realtime monitoring we do as we deploy these at scale. I think all of that will end up being a huge differentiator, and we already hear about it from our users. And I think they would walk with their feet as this frontier goes through.

Michael Nathanson, MoffettNathanson: Thank you.

Operator: Thank you. And that concludes our question-and-answer session for today. I'd like to turn the conference back over to Jim Friedland for any further remarks.

Jim Friedland, Head of Investor Relations: Thanks, everyone, for joining us today. We look forward to speaking with you again on our Fourth Quarter 2026 call. Thank you, and have a good evening.

Operator: Thank you. Everyone, this concludes today's conference call. Thank you for participating. You may now disconnect.