

Q2 2026 Earnings Call

Company Participants

- Anat Ashkenazi, Chief Financial Officer
- James Friedland, Head of Investor Relations
- Philipp Schindler, Senior Vice President and Chief Business Officer
- Sundar Pichai, Chief Executive Officer

Other Participants

- Brian Nowak, Analyst, Morgan Stanley
- Douglas Anmuth, Analyst, JPMorgan Chase & Co
- Eric Sheridan, Analyst, Goldman Sachs Group, Inc.
- Kenneth Gawrelski, Analyst, Wells Fargo Securities, LLC
- Mark Shmulik, Analyst, Bernstein Institutional Services LLC
- Michael Nathanson, Analyst, MoffettNathanson LLC
- Ronald Josey, Analyst, Citigroup Inc.
- Ross Sandler, Analyst, Barclays Bank PLC
- Shweta Khajuria, Analyst, Wolfe Research, LLC

Presentation

Operator

Welcome, everyone. Thank you for standing by for the Alphabet Second Quarter 2026 Earnings Conference Call. At this time, all participants are in listen-only mode. After the speaker presentation, there will be a question-and-answer session. (Operator Instructions)

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

James Friedland {[BIO 1556447 <GO>](#)}

Thank you. Good afternoon, everyone, and welcome to Alphabet's Second Quarter 2026 Earnings Conference Call. With us today are Sundar Pichai, Philipp Schindler 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 {[BIO 15004624 <GO>](#)}

Hi, everyone. Thanks for joining us. We have exciting momentum. Alphabet revenue grew 24% year-over-year. Our AI investments are redefining what's possible across every part of our business. Our momentum starts with Search where people are adopting one seamless search experience across AI overviews and AI mode. We saw 17% revenue growth in Search and other and YouTube ads grew 13%. Cloud revenue grew 82%, powered by strong demand for AI infrastructure and AI solutions.

And cloud backlog grew to \$514 billion. It's great to see the wide adoption of Gemini Enterprise with nearly 90% of Fortune 100 using it. Today, I'll expand on our AI progress, our global product footprint led by Search and YouTube, Google Cloud's leadership, and our long-term bets. First, AI. Yesterday, we announced new models, Gemini 3.6 Flash and 3.5 Flash-Lite, which are cost-effective and highly efficient. We are seeing tons of demand for our workhorse Gemini Flash series because it hits the sweet spot of performance and cost. We also launched Gemini 3.5 Flash Cyber, which I'm really excited about. Paired with our core vendor agent, it finds and fixes vulnerabilities and delivers performance at the frontier comparable to far bigger cyber models.

Gemini 3.5 Pro is currently in testing, and our team is already building the next generation of models. We have started our most ambitious pre-training run yet for Gemini 4 and are excited by the progress we are seeing at the frontier. Demand for our models is translating to strong token usage across developers and enterprise customers, and we continue to be supply constrained, a sign of momentum and rapid adoption. More than 9 million developers are building each month with our models across our APIs and key developer products. Our model APIs are now processing approximately 22 billion tokens per minute. That's up from 16 billion just a quarter ago.

Additionally, this quarter, we launched Omni. It allows users to create anything from any input, starting with video. Since launching at IO in May, there's been a 40% increase in daily active users creating videos on the Gemini app. Our Gemma family of open models small enough to run on local devices are hugely popular. These models have been downloaded over 900 million times and our latest Gemma 4 models have been downloaded over 300 million times since launching in April.

Our Agentic development platform Antigravity allows anyone to build in the Agent First era. It has more than 2.4 million weekly active users. Antigravity is a powerful tool for users and enterprises, and it's completely accelerated how we build internally. As just one example, a team in Chrome is now on track to accelerate delivery by 8x, compressing a two-year timeline into three months through model-driven refactoring.

Moving on to our global product footprint, which brings AI to more people than any other company. Starting with Search, AI continues to drive an expansionary moment with new experiences resonating with users and driving growth in queries. As a big football fan, I was particularly excited to see search usage hit an all-time high during the World Cup this year. This really highlights how much people turn to Google in moments that matter. We continue to make Search more helpful and intuitive. We recently brought together AI0 views and AI mode into one seamless search experience that combines our frontier capabilities with the best of the web. And we are continuing to incorporate more frontier capabilities into search with agents, personal intelligence, and notebooks.

Our AI-powered features are driving increased search usage. Since expanding AI mode globally last October, we have surpassed 1 billion monthly active users. And just like AI overviews, AI mode is driving an incremental increase in search queries overall, and we

are now sending billions of clicks to websites every week through AI features in Search. As we serve more of these queries, we have continued to drive efficiencies. Thanks to our engineering and hardware optimizations, this quarter, we reduced the cost of AI mode responses to its lowest level since launch, even as we have brought more advanced AI capabilities. Philipp will also talk more about how we are innovating the ads experience.

Next, the Gemini app, which now has 950 million monthly active users with daily active users tripling in the last year. Users love new agentic features like Daily Brief and our personalized agent Gemini Spark, which is now available in the US and internationally. We have been shipping helpful new features like this at an incredible clip.

On to YouTube, month-after-month, when a big event happens in the world, the world turns to YouTube. It's been incredible to see that over 1.7 billion unique viewers globally watched World Cup-related videos on YouTube during the FIFA World Cup 2026. We are also bringing the power of conversational AI directly into the YouTube experience. Ask YouTube uses our Gemini models to let people ask complex questions about individual videos, get quick takeaways, and jump straight to moments in those videos. The early engagement is encouraging. More than 140 million users engaged with Ask YouTube on the watch page in June 2026. And we are bringing that Ask experience to the wider search experience on YouTube.

Next, Cloud. Our continued momentum is driven by our integrated AI portfolio consisting of chips, models, data, security and agent platforms, all designed to work together. Gemini continues to be a key driver of growth and is deeply integrated across all of our cloud products, including Gemini Enterprise, Data Analytics, Cybersecurity and Google Workspace. We are seeing strong diversified demand across products, customers, geographies and industries.

Our product differentiation is driving expansion in three ways. We are winning new customers, more than doubling our acquisition velocity year-over-year. We are deepening our relationships with existing customers who are expanding their usage and exceeding their commitments by more than 50%, also an acceleration over last quarter. We are driving growth with partners with transactions on Google Cloud Marketplace growing over 7x year-over-year. One of the strongest parts of our growth comes from the rapid adoption of our Gemini Enterprise platform, it's differentiated with easy-to-use tools to build agents and automate processes, connectivity to enterprise systems, cost management and governance tools.

In Q2, Agent development kit, our framework for building and deploying Enterprise AI agents reached nearly 70 million total downloads. As I said, nearly 90% of Fortune 100 are using Gemini Enterprise. We have customers like PepsiCo for AI and analytics solutions, Intel to streamline core processes, HSBC for Wealth Management, Bell Canada for customer engagement, Macy's for commerce experiences, and SIGNAL IDUNA for knowledge management.

More broadly, Gemini is transforming how millions of businesses use AI to build custom agents, automate processes, improve cybersecurity, manage customer relationships, streamline data analytics, collaborate effectively, and more. All of this momentum is driving growth in our paid token usage. Nearly 500 cloud customers have each processed more than 1 trillion tokens in the last year, and usage is so much deeper than that.

Over the last 12 months, more than 2,000 enterprises consumed over 100 billion tokens. We are also seeing strong interest in our AI-powered security platform, which is differentiated because it integrates threat intelligence, cyber response prioritization with Wiz, AI-automated scanning, code remediation, and monitoring.

Today, 90% of Fortune 100 are Google Cloud Security users. Nearly 90% of Wiz customers are using AI-powered security features, and we have seen a more than 45% quarter-over-quarter increase in the number of AI workloads scanned and protected by our security platform. Our security tools are being used to protect critical infrastructure, including financial services organizations such as Morgan Stanley, telecommunication providers such as TELUS, healthcare organizations such as Texas Children's, software companies such as Atlassian, and several government agencies.

And with our new Google AI threat defense, we are really excited to bring our new cyber model and CodeMender to help our customers defend against AI threats. All of this is powered by our leading AI infrastructure. We offer the industry's broadest range of accelerators from Google and NVIDIA, including the new NVIDIA Vera Rubin platform and TPU 8t and 8i, which delivers strong price performance. On top of this, let me mention three unique offerings.

First, our Virgo network, which is designed to meet the needs of modern large-scale AI workloads. It allows customers to connect 1 million AI accelerators across multiple data center sites into a unified supercomputer. Second, our software stack has native support for JAX, PyTorch, vLLM, and SGLang, enabling workload portability across GPUs and TPUs.

And third, our new agent-optimized Axion CPU provides 30% better performance per dollar compared to peer offerings. We are seeing strong growth in demand for our AI infrastructure offerings from leading labs such as Ineffable Intelligence, next-generation AI builders, including Kakao, financial services like Deutsche Borse Group, pharmaceutical companies such as Pfizer and Roche, and robotics and spatial intelligence companies such as World Labs.

We are continuing our strategy of investing in promising companies and internal efforts that bring pioneering AI research into the real world. We have seen exciting progress from our Other Bets. Waymo introduced its newest vehicle, Ojai, to public riders. This is the first vehicle powered by the sixth generation Waymo driver and will welcome more riders in the coming months. Wing has safely completed more than 1 million home deliveries and we continue to grow our presence through partnerships with Walmart, DoorDash, and Papa Johns.

In health and drug discovery, Isomorphic Labs raised over \$2 billion to power its AI drug design engine, scale globally, and advance its drug candidate pipeline. It's truly been an extraordinary first half of 2026 with much more to come. Thanks to all of our employees and partners worldwide for their incredible work this quarter.

With that, Philipp, over to you.

Philipp Schindler {BIO 21014597 <GO>}

Thanks, Sundar, and hello, everyone. I'll start with a review of Google Services, highlighting our momentum across Search, YouTube and partnerships. Google Services revenues were \$95 billion for the quarter, up 15% year-on-year, primarily driven by

Search. Search and Other delivered 17% growth, with retail and finance driving the largest contributions. YouTube advertising revenues grew 13%, driven by direct response and brand. Network advertising revenues were down 1% year-on-year. Strong performance in Search and YouTube underscores how our investments in AI translate into measurable value for users and advertisers.

Starting with Search and Other. Revenues reached over \$63 billion for the quarter. Sundar highlighted the momentum we see in Search, which directly impacts our advertising business. We continue to accelerate the deployment of Gemini across our entire ads infrastructure to boost performance in three areas mentioned before; ads quality, advertiser tools, and AI user experiences.

First, ads quality. At Google Marketing Live, we showcased how Gemini improves query understanding, allowing us to find relevant ads for longer searches previously difficult to monetize. The core engine of our search ads relies on a dual prediction, delivering immediate utility for the user while maximizing measurable value for the advertiser. Gemini completely supercharges this capability. We use Gemini's advanced reasoning to decode the nuances of longer, more detailed queries. With shopping ads, for instance, we drove a 20% improvement in showing highly relevant ads, helping shoppers immediately find the best match.

Second, advertiser tools. Take AI Max. It has become the core building block for advertisers to fully participate in our new AI experiences. It's out of beta, and 0.5 million advertisers have already adopted it. Those who adopt our AI-powered campaigns like AI Max or PMax see an average of 15% more conversions or value on search at a similar ROAS. AAA Auto Club Enterprises used AI Max to personalize creative assets and captured growth from increasingly detailed insurance searches. This led to a 17% improvement in conversion volume and an 11% decrease in cost per lead. We see strong adoption of our generative AI creative tools, which makes creative development easier, especially for SMBs. In fact, over half of our SMB customers globally use AI to create or optimize their creatives.

Third, new AI user experiences. At IO, we showed how search, including AI overviews and AI mode works as one connected experience. As you saw from our results, our performance across this holistic system is robust. We continue to be encouraged with monetization performance on queries that show AI overviews, even as we've expanded AI overviews to more commercial queries. Across AI overviews and AI mode, people are asking more specific and detailed questions, providing opportunities for more relevant ads.

Within AI mode, we continue to test and deploy a range of new ad formats. For text ads, we improved performance by adding contextual side links based on the conversation. Direct offers is gaining momentum with partners like IHG Hotels and Resorts soon servicing special offers during trip planning. In highlighted answers, our latest experience placing clearly marked sponsored links inside list responses is showing early user traction.

Looking at commerce. In collaboration with the retail industry, we established the open source Universal Commerce Protocol UCP as the new standard for Agentic Commerce. Merchants are rapidly adopting UCP with Target and Steve Madden now live, while new members have joined the UCP Shopping and Food Tech Councils to help steer its vision.

We also announced Universal Card, allowing shoppers to add items from different retailers across Google services and buy in a single checkout.

Now on to YouTube. As a FIFA-preferred platform partner, we helped fans around the world tune into the games. As Sundar mentioned, over 1.7 billion unique viewers watched World Cup-related videos and over 550 million watched them on their televisions. This made the FIFA World Cup 2026 the most viewed World Cup in YouTube history. Advertisers connected with fans through FIFA channel takeovers, Gemini-powered soccer-themed sponsorships, and game day mastheads and of course, a huge congratulations to Spain for taking home the trophy.

More broadly, on YouTube, we introduced an exclusive slate of creator shows making it easier for brands to tap into the fandom of creators. We launched custom sponsorships to put brands at the heart of the world's biggest moments as they unfold on YouTube, using AI to dynamically surface videos tailored to brands' desired moment. Brands continue to partner with YouTube creators to engage new audiences.

Kate Spade reached Gen Z through a first-of-its-kind YouTube creator campaign and partnered with creators Ellie Thumann and Hannah Meloche. By leveraging multiple format videos on YouTube, the brand drove a 3.25% brand lift in purchase intent. This unique conversions of brand equity and commercial action makes YouTube a powerful full-funnel platform. SMBs continue to fuel our direct response growth through campaigns like Demand Gen. They leverage our AI to scale visual storytelling across YouTube, Shorts, and now Google Maps.

To attract high-value shoppers across North America and Europe, outdoor brand Arc'teryx used Demand Gen to achieve a 70% better return on their ad spend compared to other paid channels. Looking at monetization across YouTube, we're driving sustained growth across our key priorities. In the living room, we see continued momentum across both brand and direct response. With the launch of Buy with Google Pay, viewers can complete purchases directly on their CTV, turning the TV screen into a stronger performance surface. Shorts continues to deliver high-performing opportunities for social and video buyers alike.

Beyond advertising, YouTube subscription business, which thrives across living room screens is growing faster than ads, particularly driven by YouTube Music and Premium. Internally, we leverage Gemini to transform how we work and how we serve our customers. 83% of our sales team uses Gemini assisted tools weekly, driving up to a 20% higher win rate when using customized pitch narratives. Our ads customer support teams use Gemini-powered agentic solutions that now autonomously address 75% of support queries, freeing them to solve our customers' most complex challenges. Similarly, our new agentic solutions for SMBs have expanded our reach by hundreds of thousands of new customers year-to-date, as always, ending with partnerships.

We're seeing more businesses benefit from the unified strength of Google's AI stack. Booking Holdings expanded a multi-year cloud commitment and are partnering closely to advance our AI-powered ad formats. It's also deploying Google's AI technology to enable new customer experiences like agentic dining reservations on open table, helping restaurants get discovered and booked right when it matters most.

In closing, I'd like to thank Googlers everywhere for their contributions to our success and our customers and partners for their continued trust. Anat, over to you.

Anat Ashkenazi {[BIO 19888043 <GO>](#)}

Thank you, Philipp. My comments will focus on year-over-year comparisons for the second quarter unless I state otherwise. I will start with results at the Alphabet level and will then cover our segment's results. I'll end with some commentary on our outlook for the third quarter and full year 2026. We had an outstanding second quarter, delivering our 12th consecutive quarter of double-digit revenue growth.

Consolidated revenues were \$119.8 billion, up 24% or 23% in constant currency. Total cost of revenues was \$45.9 billion, up 18%. Tech was \$16.2 billion, up 10%. Other cost of revenues was \$29.8 billion, up 22%, driven by increases in depreciation, inventory costs, primarily from the sales of TPU systems to customers, and content acquisition costs largely for YouTube. Total operating expenses were up 27% to \$33.1 billion. R&D expenses increased by 32%, driven by compensation from investments in AI talent as well as depreciation.

Sales and marketing expenses were up 18% due to investments to support the Gemini app and search and compensation. And G&A expenses increased 24%, primarily driven by compensation and charges for certain legal and other matters. Operating income increased 30% to \$40.8 billion and operating margin was 34%. Other income and expenses was \$98 billion, representing a substantial increase from the prior year, primarily due to unrealized gains in our equity securities portfolio.

Net income and earnings per share increased significantly, primarily due to the unrealized gains in OI&E I just mentioned. We generated strong operating cash flow of \$39.1 billion in the second quarter and \$185.7 billion for the trailing-12 months. CapEx was \$44.9 billion in the second quarter, with the vast majority of the spend in technical infrastructure to support our investments in AI. Approximately 60% of our investment in technical infrastructure this quarter was in servers, and 40% was in data centers and networking equipment.

We had negative free cash flow of \$5.9 billion in the second quarter, driven by our investments in CapEx. Free cash flow was \$53.3 billion for the trailing 12 months. We ended the quarter with \$242.5 billion in cash and marketable securities, which includes \$87.1 billion of marketable equity securities. Long-term debt was \$98.2 billion. And as mentioned in our press release, our Board of Directors declared a quarterly cash dividend on our common stock of \$0.22 per share, which is payable in September.

Turning to segment results, Google Services revenues increased 15% to \$94.5 billion, reflecting strong growth in Search and Subscriptions. Total advertising revenues were up 14% with growth across all major verticals. We experienced strong ad growth related to the World Cup, particularly in YouTube ads. Google Search and other advertising revenues increased by 17% to \$63.3 billion, with retail and finance driving the largest contributions.

YouTube advertising revenues increased 13% to \$11.1 billion, driven by direct response advertising as well as brand with strength in the living room as Philipp mentioned earlier. Network advertising revenues of \$7.3 billion were down 1%. Subscription platforms and devices revenues increased 15% this quarter to \$12.9 billion due to strong growth in both YouTube subscriptions, particularly YouTube Music and Premium and Google One, which was driven by demand for AI plans.

Google Services operating income increased 20% to \$39.5 billion and operating margin was 41.8%. The Google Cloud segment delivered outstanding results in the second

quarter, driven by our enterprise AI products and services. Cloud revenues were up 82% to \$24.8 billion, driven primarily by GCP, which grew faster than cloud overall. Core GCP, AI solutions, and AI infrastructure were all important drivers of growth.

We also began to recognize revenues from TPU system sales, which we delivered to customer data centers for the first time in Q2. Cloud revenue growth accelerated meaningfully even after excluding the impact of TPU system sales. Cloud operating income was \$8.8 billion, more than tripling year-over-year, and operating margin increased from 20.7% in the second quarter last year to 35.6%. Google Cloud's backlog increased by more than \$50 billion sequentially, reaching \$514 billion in the second quarter.

The increase was driven by strong demand for our enterprise AI offerings. The majority of the backlog is related to typical GCP contracts from a broad mix of customers and we expect to recognize just over 50% of the total backlog as revenue over the next 24 months. In Other Bets, revenues were \$382 million and operating loss was \$1.8 billion as we continue to expand Waymo's business and invest in key Other Bets. And in Alphabet-level activities, the operating loss was \$5.8 billion, driven by shared AI R&D expenses.

Turning to our outlook and business performance expectations in the second half of 2026. First, in terms of revenues, at the current spot rates, we would expect a slight FX headwind to our consolidated revenue in Q3 compared to a 1 percentage point FX tailwind in Q2. This impact will be seen primarily in Search and YouTube ads. In Google Services, as Philipp mentioned, we continue to benefit from innovation in search by improving the user and advertiser experience with AI. At the same time in Q3, we will begin lapping an acceleration in search performance that began in the third quarter last year.

In Google Cloud, we're seeing significant demand for our products and services, which we expect to drive strong growth. As I mentioned earlier, we started delivering TPU system to customer data centers in the second quarter. We continue to expect to recognize a relatively small portion of the revenues from our existing TPU system sales agreements this year, ramping as we exit 2026. We anticipate the vast majority of the revenues from these agreements will be realized in 2027. And given the supply-constrained environment, we plan to expand the use of third-party capacity in Q3 as a bridging strategy while we build out more internal capacity. This strategy allows us to keep growing our customer base and capture greater overall value. However, it will create modest margin pressure in the near-term as we utilize this capacity.

Moving to investments, we are updating our full year 2026 CapEx guidance range to \$195 billion to \$205 billion, up from our previous estimate of \$180 billion to \$190 billion. The increase in the range is primarily due to an acceleration in the delivery of capacity to meet growing demand. As we previously shared, we continue to expect our CapEx to increase significantly in 2027, and we'll provide more details at a later date.

In terms of expenses, the significant increase in our investments in technical infrastructure will continue to put pressure on the P&L in the form of higher depreciation expense and related data center operations costs such as energy. We also expect to continue hiring in key investment areas such as AI and cloud, and we are investing in marketing to support our AI products.

And finally, we expect the free cash flow will remain under pressure driven by our investments in technical infrastructure, which enables us to capitalize on the AI

opportunity and continue to drive attractive returns. Q2 represented another strong quarter. Our teams continued to deliver impressive innovation, executing with discipline and velocity. I want to take this opportunity to thank our employees for their contribution to our performance. Sundar, Philipp, and I will now take your questions.

Questions And Answers

Operator

Thank you. (Operator Instructions) Our first question comes from Brian Nowak with Morgan Stanley. Your line is now open.

Q - Brian Nowak {BIO 16819013 <GO>}

Thanks for taking my questions. I have two, one for Sundar, one for Anat. Sundar, with more time in GenAI products and tools in the market and investment continuing to step up, can you just kind of give us some perspective on how your view on the size of the overall GenAI ROIC opportunity and the timing of the ROIC has changed now versus one year ago?

And then Anat, you talked quite a bit about supply capacity constraints. Maybe a similar question for you around forward CapEx. How is your budgeting philosophy and these constraints you're putting on to forward spend, how are those changing as you sort of think about the right amount to spend in 2027 to close this capacity constraint?

A - Sundar Pichai {BIO 15004624 <GO>}

Look, thanks, Brian. I mean, it's a good question. I do think it feels like we are in very early innings of what feels like secular shift across multiple areas. In our core information businesses, just the possibilities, when I see what all you can do with the absolute frontier capabilities, there's still a lot of work ahead to translate all that into experiences for our consumer users. You can think about end-to-end agentic experiences to really meaningfully do a lot more for them.

So, all of that looks like extraordinary opportunities with extraordinary returns for executing well on those opportunities. Similarly on the enterprise side, as you can see from our demand, which is reflected in our growth rates, et cetera, again, in my conversations with many CEOs, many companies, they're all still barely scratching the early stages of what's possible here. So, I think we used to talk about cloud itself, very -- small percentage of overall workloads and enterprises have shifted to cloud. Now think about what percentage of workloads have -- are really AI-native and AI-enabled. It again feels like very, very early.

So I would say from a ROIC standpoint, I think we are taking a full-stack approach. We are seeing momentum across consumers and enterprises and developers and so on. So it feels like, if anything, over the past year, we have gotten more bullish on the opportunities ahead. I'll turn it over to Anat.

A - Anat Ashkenazi {BIO 19888043 <GO>}

Thanks for the question on how do we think about forward-looking CapEx in the context of supply constraints. So, we're still in a supply-constrained environment. I think we've said this now for multiple quarters in a row. And we are seeing very strong demand, both

from external cloud customers as well as across the business. And our goal is to invest as long as we see an attractive return on that investment, as Sundar mentioned earlier. We do take a long-term view. So we take a multi-year view at what the needs are as well as focus on next year in the near term and building aggressively to meet those demands.

And as you've seen, while we have increased our capacity quite significantly over the past three years, the demand still outpaces that investment. And we are, just as the rest of the industry, working in a supply-constrained environment. So we're working hard to do that. We do have a benefit of having a -- the full stack approach. So we're able to drive operational efficiencies, technological efficiencies within our technical infrastructure organization, so that we can deliver more compute. But as long as we see these attractive opportunities to invest, we will continue to invest.

Q - Brian Nowak {BIO 16819013 <GO>}

Thank you, both.

Operator

Our next question comes from Doug Anmuth with JPMorgan. Your line is now open.

Q - Douglas Anmuth {BIO 5591566 <GO>}

Thanks for taking the questions. One for Sundar and one for Anat. Sundar, can you just talk about your confidence level that the Gemini models can remain at the frontier? I know you had the Gemini, the Flash model releases this week, but Google perhaps doesn't make maybe quite as much noise as some of the other leading labs or perhaps have quite the frequency. I'm just curious how you'd address any concerns about that ability to continue developing leading models? And then just related, can you talk about your plans for coding and how you look to close the gap in that part of enterprise?

And then, Anat, just following the recent equity and debt raises, can you just help us understand how you think about optimal capital structure and how you view the cost of debt versus equity? Thanks.

A - Sundar Pichai {BIO 15004624 <GO>}

Thanks, Doug. Look, I think the frontier is incredibly a dynamic space, and it's fiercely moving forward. At any given moment when you take a snapshot, it feels dynamic. We are -- we've had clearly frontier models. There are many attributes on which we are still at the frontier. There are areas where we have acknowledged we need to improve coding and agentic coding is an example of that and I think the teams are very, very focused on it with 3.6 Flash. And by the way, Flash is our workhorse model, and we see tremendous demand for it now given it generally hits a sweet spot of performance, cost, reliability, latency, et cetera.

And we've incorporated Gemini and Flash particularly in our entire portfolio of solutions, be it in cybersecurity, data analytics, if you think about an area like customer service, you need really good voice quality, you need live-streaming, you need the ability to reason on that or if you're a professional services firm, you need high-quality summarization, content generation, et cetera.

So, Flash does very well on all of that. In terms of agentic coding, we are iterating and you will see us make continued iterations. 3.6 Flash, for example, compared to 3.5 Flash jumped over 10 points in DeepSuite as a benchmark and it is more token efficient doing so. So we are clearly -- we are using it internally. We are testing it with many customers in coding. So, we see the progress just in six weeks from the prior version to the new version, and you will see more continued iterations on that as well.

In terms of the frontier, we are both very committed and very confident of being at the frontier. For the next generation of frontier, you're going to need much larger base models. We are now training Gemini 4, and we are being very ambitious with it. I am very excited by the progress I'm seeing internally on Gemini 4, and I'm confident that people will be pleased when we are putting it outside. But we will need Gemini 4 as a larger base model to compete at that frontier level. And so we are focused on executing on that well.

A - Anat Ashkenazi {[BIO 19888043 <GO>](#)}

And on the question regarding our capital structure and recent equity and debt raise. So, as we think about our investment needs, and as I said, we look at the next year and multiple years out, we first look at how much we can support based on cash from operations. And as you've seen in our results today, we continue to generate very healthy strong cash flow from operations. So that's our first source of funding. And then we look at debt and most recently, we did the equity raise. And we have expanded our debt portfolio quite significantly over the past 12 months. If you've seen a year ago, we were at about \$16 billion and we went to about \$100 billion across multiple currencies and geographies. So we're expanding that portfolio.

But we also want to make sure we have a resilient not just growth outlook, but also a resilient balance sheet and a strong balance sheet, healthy balance sheet, which is what the rationale behind expending into the equity markets. At this point, we're not planning to go back to the equity markets with the exception of, as you recall, part of our equity offering was the ATM or at-the-market offering that we will do to address the stock-based comp or the tax on SBC, which we'll do for some period of time. But we look at those three dimensions, the cash flow from operations, debt levels, and equity and a way to balance it in a way that ensures that we continue to maintain a healthy balance sheet.

Q - Douglas Anmuth {[BIO 5591566 <GO>](#)}

Thank you, both.

Operator

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

Q - Eric Sheridan {[BIO 22465717 <GO>](#)}

Thanks so much for taking the questions. Maybe two, if I could, both on TPUs. Sundar, can you talk about some of your key learnings as you scale your TPU efforts in terms of both the demand for TPUs and how you think about balancing the external demands for TPUs versus the internal demand for custom silicon inside the organization as you look over the next couple of years?

And then, Anat, if I could just follow up, you made some comments in your prepared remarks about the impact of TPUs on Google Cloud. I didn't know if you could give us a

little bit more granularity about the size of TPUs as a part of the revenue backlog and how to think about the revenue recognition or the potential for margin impact from TPUs in the years ahead? Thanks so much.

A - Sundar Pichai {[BIO 15004624](#) <GO>}

Okay. I think in terms of -- first of all, we are very pleased with our TPU roadmap progress and the value in terms of performance and the edge it gives, and we obviously use it extensively. In terms of allocating TPUs, look, our first priority is making sure we are allocating what we need to compete at the frontier in terms of AGI development. And so that is the foundation for everything we do. But given the extraordinary demand to balance the external demand, even for most cloud customers, we are using both TPUs and GPUs, mainly for serving our model.

So, think Vertex AI, Gemini Enterprise, the momentum we see. So, we're using it for agent workloads, et cetera. So, we are using it for those purposes. To the extent people want it as infrastructure, we are balancing it by increasingly looking at opportunities to put TPUs in their data centers or in other data centers like the project we are doing with Blackstone, et cetera. So, we're using that as an ability to balance and make sure we are allocating as much of the available compute for frontier model development as well as serving now both our consumer business and our enterprise business in terms of first-party models.

A - Anat Ashkenazi {[BIO 19888043](#) <GO>}

And in terms of revenue recognition and how we look at the TPU system sales. So the way to think about it is the following; when we signed the agreements that I've mentioned in the prepared remarks, they would be then reflected in the back -- the cloud backlog. Now the vast majority of the \$514 billion of cloud backlog is the GCP agreements, but the TPU system sales are reflected in that backlog. We start building inventory to be able to sell those systems. So, you see that impact on the cash from operations because we built ahead. Obviously, as we're building that business and ramping up.

And then once we start delivering the sales, generally, that's when we start recognizing revenue. This quarter was a small amount of the total -- that total agreement, it will continue to ramp up throughout 2026, but then you'll see the vast majority of the revenue from that agreement comes through in 2027.

Operator

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

Q - Ross Sandler {[BIO 15948659](#) <GO>}

Yeah, just back to the AI model kind of war that's going on, Sundar. I just wanted to ask about like the speed of model releases and how you can potentially speed up the cadence by which Gemini is releasing models. So, we all saw the third-party compute deal with SpaceX. What else are we doing to kind of speed up the pace of model releases compared to some of the other folks out there? And then you just mentioned the Flash models, which have been very successful to your point. Do you view that as like the right end of the market to be in given how crowded the lower cost end of the model space has become? Just any thoughts on how you see all that playing out?

A - Sundar Pichai {BIO 15004624 <GO>}

Thanks, Ross. Look, I mean, two things I'll say first of all, we are -- we really think about the parade of frontier. We want to make sure for our customers, we are offering a best model -- best models at various price points, right? So you will see us. But it's very important to us to have the best frontier models out there, as well as models which are very performant and low cost. This is why we have Flash-Lite, Flash Pro, et cetera, and with -- so we are committed to being at the full parade of frontier.

On the speed of model releases, I do think you will see us continue to pick up pace. Obviously, you've seen us come at -- we announced 3.5 Flash at IO. We have followed that up with 3.6 Flash, and you'll see continued iterations of that model with further progress in agentic coding, et cetera. We are really focused, putting a lot of effort into Gemini 4. It's a very ambitious effort. We wanted to compete at the frontier level of where the frontier will be when Gemini 4 comes out. And so we are applying a lot of our compute and effort in that direction.

But with that, we are creating a baseline on which -- on top of which you will see us rapidly iterate with subsequent model releases. And so picking up pace and releasing models almost at a monthly cadence is part of our roadmap as we are building Gemini 4 as well.

Operator

Your next question comes from Michael Nathanson with MoffettNathanson. Your line is now open.

Q - Michael Nathanson {BIO 2009061 <GO>}

Thanks. Thank you. One for Sundar and one for Anat. Sundar, just keeping with the model theme of this call, can you talk a bit about what do you think the moats are to this new model war, what is your strategic advantages that fuel your business that even if everyone gets to the same type of model capabilities, what gives you the edge to keep growing like you're doing?

And then, Anat, back to Eric's question, anything you could tell about the margin impact on TPUs going forward, right? Is it incremental margin to the core cloud margin or is it a lower-margin coming through? Thanks.

A - Sundar Pichai {BIO 15004624 <GO>}

Look, Michael, I mean, first of all, look, there are many layers at which we are providing solutions, right? I think part of what the value of our full-stack approach is, customers are coming to us for solutions. So if you take an area like cybersecurity or data analytics, people are deploying solutions in which models are an ingredient in that solution, right? So take cyber, people are using Chronicle and Wiz and with our upcoming CodeMender product, they're using it to detect and patch vulnerabilities and so on or take data analytics, people are bringing what's been in the past siloed data sources, putting it all together and having an intelligence layer on top of it with Gemini Enterprise. And so, the model is just an ingredient in those solutions. So I think it's important to remember that.

Then even within where you're purely on a more model consumption area, people are -- there is what looks as models are increasingly becoming end-to-end orchestrated systems, right, workflows, agent workflows. And to develop all this, it's your ability not just to bring the compute you need to train and serve, the data quality, the environments you have, your ability to continually improve, provide customers with that peace of mind that all the data they have, their data, their trajectories are confidential to them. None of that in any way flows back to the models, their ability to configure and serve all this in a secure way and manage it and provision it and so on.

So, these are big end-to-end things. And so this is what our cloud business as a whole is focused on building it. And when we are seeing demand across all of these components. And so obviously, having our own models allows us to really optimize these solutions and bring an integrated offering, but we will also provide other models as part of these solutions. We also bring infrastructure as part of the solutions, et cetera, and take a whole full-stack approach. So I think we are very, very well-positioned there.

A - Anat Ashkenazi {[BIO 19888043](#) <GO>}

And on the TPU margins, we don't break out margins for any specific products or infrastructure component. And certainly, there are benefits from designing and manufacturing our own chips. But the way to think about it is this is an expansion of our total addressable market. So it expands the opportunities by providing solutions to customers that need those systems in their data centers. But overall, if you think about the cloud margins, obviously, a really fantastic expansion of margin to 35.6% this quarter, outstanding strong operational discipline across the business and leveraging growth in the topline.

And now as you think about Q3 and through the remainder of the year, and I mentioned some of it in the prepared remarks, given the supply-constrained environment we're in, we are planning to expand the use of third-party capacity in Q3 as a bridging strategy while we continue to build out internal capacity. That capacity, given the cost of that capacity will put some pressure on operating margins for cloud, and then the last item we mentioned in the past was the integration of Wiz does create some headwind here in the near-term in 2026 related to the acquisition.

Q - Michael Nathanson {[BIO 2009061](#) <GO>}

Thanks, guys.

Operator

Your next question comes from Mark Shmulik with Bernstein. Your line is now open.

Q - Mark Shmulik {[BIO 21311083](#) <GO>}

Yes, thanks for taking the questions. Anat, Just on these third-party deals and kind of the bridge and capacity, is there a particular objective where the constraints are most severe or is it really broad-based? And I guess stepping back a bit and thinking about capacity allocation across the businesses, has that changed at all? Or is there a way to contextualize how you're now thinking about the trade-off between allocate and compute to search versus model training versus for the GCP business? Thank you.

A - Sundar Pichai {[BIO 15004624](#) <GO>}

Mark, I can comment on it, like -- look again on allocation, I think, look, we -- the baseline with which we start is what it takes to continue AGI development at the frontier. Obviously, the priors on that are just based on where the model frontier moves. And so we start with that as a baseline. And beyond that, we want to obviously, we are prioritizing our core product areas like Search, YouTube, et cetera, as well as Cloud. And within Cloud, we are prioritizing the compute to make sure we can serve our models in the context of Vertex and Gemini Enterprise and our core solutions, be it data analytics and cybersecurity, et cetera, right?

So, our core serving for our core products across consumers and enterprises where the compute is primarily going, and that's how we think about it. I think on the bridge deal, the main thing I would say is, look, there are on the margin, they are very, very large customers of ours on cloud who we are trying to support them through this extraordinary moment and the incremental opportunities they are bringing to us, while a short-term cost over a few months may be very high in the lifetime of the deal as we bring more capacity on is highly ROI positive, right?

So, those are factors we are taking into account. So, are you willing to take upfront six-month deal to be able to serve that customer in what is a multiyear opportunity where the margins and the returns are very, very attractive over that multiyear horizon. So hopefully, that gives some color on how we thought about those opportunities.

Operator

The next question comes from Ron Josey with Citi. Your line is now open.

Q - Ronald Josey

Great. Thanks for taking the question. Maybe I'll switch topics somewhat, and Philipp, ask you a little bit more about monetization from a Search and YouTube perspective, just given the greater signal we have with AI searches, and then you talked about YouTube strength. Talk to us a little bit more on how advertisers are leveraging the greater personalization and targeting that Google has to drive this ROAS. A question we often get is Google is at size and scale, we're still growing 17%. Any reasons what's driving that all-time high on searches? Thank you.

A - Philipp Schindler {BIO 21014597 <GO>}

Yes, thank you so much for the question. Look, our 17% year-over-year growth rate in the Q2 Search and other revenues was really driven by many parts of our business working well together and very, very deep Gemini integration. Maybe to zoom out for a second, all the major verticals contributed to the growth. Retail drove the greatest contribution, not managed, followed by really meaningful contributions from finance, tech, media, entertainment, and so on.

I think it's really important to understand that Gemini supercharges our ability, and this goes specifically to your question to understand what people are looking for and match the right ads. So we're applying Gemini models really across our entire ads infrastructure, whether it's ads quality, advertising tools, ads and new AI experiences. And we're really deeply integrating Gemini into the customer tools to make the campaigns more efficient, which comes on top of this.

And then we have our AI-powered campaigns like AI Max that help advertisers actually adapt and find the opportunities beyond Keyword. Again, that's an ability for us to go deeper and target better and AI Max continues to unlock, we mentioned this billions of net-new searches that weren't really monetizable before.

And then on the YouTube side, high level, the ads growth was driven by direct response and brand advertising. I shared earlier, we continue to see a lot of fermenting in the living room. We have a really exciting ads roadmap as well ahead of us on the brand side. On the direct response side, Demand Gen and Shorts remain really interesting opportunities here. Again, this has targeting components to it, obviously and we continue to innovate really heavily in diverse bonds with shoppable ad formats in the living room, which will help us drive strength in retail as well.

Operator

Your next question comes from Ken Gawrelski with Wells Fargo. Your line is now open.

Q - Kenneth Gawrelski

Thank you very much. Two questions, please. First, how would you assess the forecasted returns on compute capacity investments in '27 compared to prior years '25 and '26 in light of the supply constraint -- the supply chain constraints and supply chain inflation we're seeing? I'm curious as to how you think about the return profile of '27 and kind of future investments versus past investments in capacity?

And the second one, a little different topic here. When you think about Waymo, what are the key factors you would consider when evaluating a change in corporate structure for Waymo? I know you've previously been reticent to talk about this, but the business is scaling clearly, there's leadership there and has a lot of momentum. Under what conditions would it make sense for Waymo to live outside of Alphabet? Thank you.

A - Sundar Pichai {[BIO 15004624 <GO>](#)}

Look, Ken, on your first question, if I understand it, look we are working off a disciplined ROIC framework here. Obviously, we -- to the extent that our input costs going up to us, we reflect that in our ability to price our solutions and see returns there. So all of that is factored into how we are planning. I think our compute capacity investments in '27, I think to the first question I answered, I think we are seeing strong demand indicators, including long-term deals, the existing deals which we have, which are renewing at -- with exceptional demand on a moving-forward basis. And so we are using all that to plan and invest accordingly.

And so, I think if anything, the dynamics look healthier than where we were about a year ago. And so that's what gives us the confidence to undertake those investments. On Waymo, look, I think we are really focused on executing and scaling up Waymo. We have set it up with our bet structure and given the team a lot of support, Alphabet has continued. I think our ability to think and plan long term and invest in the business and give them that long-term roadmap, the confidence in undertaking a long-term roadmap and scaling up, I think it's all hugely valuable. We are really focused on scaling the business right now and executing to that extraordinary potential, and that's what we are focused on.

Q - Kenneth Gawrelski

Thank you very much.

Operator

Our last question comes from Shweta Khajuria with Wolfe Research. Your line is now open.

Q - Shweta Khajuria {BIO 19165741 <GO>}

Okay. Thanks a lot for taking my questions. Could I please ask you two? One is on TPU long-term strategy, longer term, is it fair to think that as TPU sales scale, you'd build merchant silicon business with a software stack that goes with it? And second is on YouTube revenue, given how strong engagement is, what are some of the factors that could drive accelerating growth rate at YouTube on a go-forward basis? Thank you.

A - Sundar Pichai {BIO 15004624 <GO>}

Maybe I'll answer the first one, and Philipp can broadly talk about YouTube revenue growth as well. On the first one on TPUs, look, I think we are -- we've been operating with TPUs for a while. A lot of our cloud serving is done on TPUs and we are also seeing strong demand for TPUs as standalone systems from other customers. And so we'll -- obviously, the industry has constraints too. So, we have to plan and reallocate on a forward-looking basis along the principles I spoke about earlier.

So I don't want to project out too far into the future there. But I do think we are -- obviously, we are -- we will scale up based on the opportunities we see and the demand we see commensurate with the constraints that exist and the allocation needs we have for frontier model development, making sure we can serve our consumer businesses and enterprise businesses well.

And Philipp, you want to add on YouTube monetization?

A - Philipp Schindler {BIO 21014597 <GO>}

Yes. Thank you so much. So look, I mentioned it, the drivers of YouTube ads revenue growth. I think there's a lot of things we're actually excited about when we look at our ads roadmap here in brand, we see ongoing opportunities on connected TVs in the living room. We're building tools to help advertisers find creators scale across screens and then measure the results. And we have a slate of new and returning creator shows that are coming exclusively to YouTube. We see opportunities in direct response, particularly with Demand Gen and Shorts, I mentioned that remain a really interesting opportunity here to expand our base to more small and medium advertisers across more verticals, shorter videos.

Actually, this important one create more opportunities for less disruptive ads, increasing the overall ad effectiveness. So we're going to continue to see a lot of innovation, in my opinion, in direct response with shoppable ad formats in the living room, which again we mentioned that drives retail strength and we announced a few other formats here at Brandcast that we announced Buy with Google Pay to enable CTV viewers to actually complete purchases directly on their TV with two clicks, which is a really interesting opportunity over-time. We launched the affiliate partnership boost to drive incremental sales and the creator earnings, obviously, from YouTube Shopping affiliate commissions. So overall, like frankly very, very exciting times.

Q - Shweta Khajuria {BIO 19165741 <GO>}

Thanks, Sundar. Thanks, Philipp.

Operator

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.

A - James Friedland {BIO 1556447 <GO>}

Thanks, everyone, for joining us today. We look forward to speaking with you again on our third 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.

This transcript may not be 100 percent accurate and may contain misspellings and other inaccuracies. This transcript is provided "as is", without express or implied warranties of any kind. Bloomberg retains all rights to this transcript and provides it solely for your personal, non-commercial use. Bloomberg, its suppliers and third-party agents shall have no liability for errors in this transcript or for lost profits, losses, or direct, indirect, incidental, consequential, special or punitive damages in connection with the furnishing, performance or use of such transcript. Neither the information nor any opinion expressed in this transcript constitutes a solicitation of the purchase or sale of securities or commodities. Any opinion expressed in the transcript does not necessarily reflect the views of Bloomberg LP. © COPYRIGHT 2026, BLOOMBERG LP. All rights reserved. Any reproduction, redistribution or retransmission is expressly prohibited.