Meta Platforms, Inc. (META)
Second Quarter 2026 Results Conference Call – Prepared Remarks
July 29th, 2026

Chad Heaton, VP, Finance

Thank you. Good afternoon and welcome to Meta Platforms' second quarter 2026 earnings conference call. Joining me today to discuss our results are Mark Zuckerberg, CEO and Susan Li, CFO.

Our remarks today will include forward-looking statements, which are based on assumptions as of today. Actual results may differ materially as a result of various factors including those set forth in today's earnings press release, and in our quarterly report on Form 10-Q filed with the SEC. We undertake no obligation to update any forward-looking statement.

During this call we will present both GAAP and certain non-GAAP financial measures. A reconciliation of GAAP to non-GAAP measures is included in today's earnings press release. The earnings press release and an accompanying investor presentation are available on our website at investor.atmeta.com.

And now, I'd like to turn the call over to Mark.

Mark Zuckerberg, CEO

Hey everyone. Thanks for joining today.

We had a strong quarter for our community and business with 3.6 billion people using at least one of our apps each day. We reached several milestones. Instagram reached 2 billion daily actives. Threads crossed 500 million monthly actives, making it the fastest growing conversation app ever. Facebook has reached more than 2 billion daily actives for a while now. WhatsApp just hit an all-time messaging record, peaking at 30 million messages sent per second during the World Cup Final. And also, Kunal Shah just joined as our new head of WhatsApp. He built one of India's most important payment companies and will be a great addition to the team. We also shipped some strong new models from Meta Superintelligence Labs and released new glasses. Overall, the scale and reach of our community across our apps is pretty remarkable, and it gives us a strong platform for delivering new innovations to billions of people.

The opportunity in front of us is massive. First, we are now at a point where our investments in AI are accelerating every major part of our core business. They're improving the experience for people using our apps, driving better performance for advertisers, and helping our teams build new experiences and ship faster. Second, we're developing new personal agents that will be the foundation for our next wave of products and revenue lines in the months and years ahead. Third, we see a large enterprise opportunity to sell to businesses, including APIs, business agents, potentially selling compute directly, and other services that we're building for large customers. I want to spend some time today laying out exactly how our investments are delivering results today and the opportunities that we see over time.

For accelerating our core business, we're seeing strong and promising results in several areas.

In Instagram and Facebook, I'm very optimistic about our work to integrate large language models into our recommendation systems. LLMs add a first-principles understanding of what the content is about and why it is compelling, as well as a deeper understanding of what people are interested in and what their goals are when they're using our apps. This means that we can show more relevant and engaging content that better reflects people's goals and interests.

Our new Muse Image and Muse Video models will also dramatically expand the universe of content that people can discover across our platforms. There are already two large sets of content to draw from -- first from your friends and the people you follow, and second from creators that you don't follow -- but now there is going to be a whole new and nearly infinite universe of personalized content. This is going to make our services a lot more useful and engaging for people.

For ads, we are using LLMs to improve how our systems predict and rank the ads that we show. We've expanded the context we can take into account around a person's organic and ads activity to determine an ad's relevance, driving significant increases in relevance and conversions on both Facebook and Instagram. On a dollar basis, our ads business is reporting faster year-over-year revenue growth than any other company's reported ad business -- and these AI investments are paying off.

We are also seeing a lot of demand for our new AI-powered creative tools. 9 million small businesses on our platforms are now using at least one of our AI ad creative tools, and we're rolling out new end-to-end creative solutions that help advertisers translate performance data into their creative decisions. Muse Image is going to supercharge this. The model can analyze images, improve its own work, and produce better ad variations based on advertiser input. We're getting great feedback on this so far.

We also just launched Meta One, a new subscription offering that provides more tools and AI features across our apps. As demand grows, we're going to offer a variety of different tiers and pricing options there as well.

I'm also excited about how AI is helping our teams speed up product development. Earlier this year we shipped Instagram Instants. We also just launched Forum, a standalone Groups app, and Seller, a standalone Marketplace app. I expect it to become a lot easier to ship new apps, so we're planning to build out more ideas and use our recommendation systems to scale them to the people who will find them interesting, as we've done with Threads.

It's been a little more than a year since we launched Meta Superintelligence Labs and our trajectory is strong. In the last month we shipped Muse Spark 1.1 and Muse Image. Since we rebuilt Meta AI and integrated Muse Spark, we've seen a 60% increase in the number of people interacting with the assistant each day and that continues to grow quickly week over week.

Muse Spark 1.1 is a strong agentic and coding model that is very efficient and excels at computer use, tool use, and multimodal understanding. It's available through our new public API, and we're ramping up distribution through partner channels and more coding agents over the coming weeks. We're also building out features to make it easier for enterprises to adopt Muse Spark and that is an area we expect to continue focusing on.

One reason we're so focused on making Muse Spark great at agentic capabilities is that we think there's a very big opportunity to ship a few types of agents that are aligned with our mission and business.

The first is personal agents. Soon, we will have agents that can work 24/7 on your behalf to help achieve your goals and improve your life, your health, your relationships, your finances -- whatever you want. The first domain that agents have really taken off in is coding, but engineers are more technical and willing to spend time making those agents work. So to build great personal agents, this needs to be a great consumer product that just works out of the box and is easy enough for billions of people to adopt and use. I'm very excited about this and we'll have more to share soon.

As we move towards a future where we're all interacting with multiple agents, I think that WhatsApp and our other messaging surfaces are going to be increasingly important. WhatsApp is already the leading surface where people engage with Meta AI -- and as we build out a platform for more agents across our messaging apps, we're going to innovate on how to deliver a private and secure AI experience. We launched incognito mode this quarter on WhatsApp and the Meta AI app, allowing people to have private conversations with their assistant that even Meta can't see. We're planning to make strong privacy and security a fundamental part of the agents that we're building as well.

I'm also very excited about our progress with business agents. We made Meta Business Agents available globally this quarter on WhatsApp and Messenger, and there are already more than 1 million businesses using them to talk to their customers or complete sales every week. We're rolling business agents out on Instagram now too. One interesting thing about having an agent talk to your customers every day is that it learns over time and can bring all those insights back to you. We're building more agentic capabilities to summarize all these conversations, digest what happened overnight, and surface what customers are asking for. Soon, it'll go further, including suggesting ways to grow your business, giving you competitive intelligence and real-time insights into what's working and what's not. Over time, we'd like to build this into a business-in-a-box service that can help you start and run a whole business using Meta's platforms.

In terms of how we'll monetize these, we have a mix of subscriptions, volume-based pricing, and I expect that we’re going to evolve more of these products to be like our ad systems where businesses only pay us when we achieve results for them. Over time, that will let us run an efficient auction over our compute, similar to how we do that for our advertisers today.

As AI usage in our products and businesses continues to ramp, we continue to invest aggressively in infrastructure to meet the demand. Yesterday, as part of our Meta Compute effort, we announced a new strategic venture with BlackRock to develop a new 1GW data center in El Paso, Texas.

Overall, we expect that a significant portion of our compute is going to go towards training our models, growing our core business, and delivering personal agents and new products. But we also expect to grow a large business serving large customers as well. We've built our API, we're rolling out business agents, we're getting a lot of offers for compute at a significant premium over what we paid for it, and we have more coding and productivity tools on our roadmap as well. We'll have more to share on all of this soon.

As we get closer to personal superintelligence, we're also going to need hardware that allows you to seamlessly interact with it. Glasses are the ideal form factor since they can be with you throughout the day and they can assist you without pulling you away from the moment. Our glasses remain one of the fastest-growing consumer electronics of all time, and we continue adding to the line-up.

We just released our own line of Meta Glasses in collaboration with EssilorLuxottica, including a style that we designed with Kylie Jenner. They're the first glasses to ship with Muse Spark out of the box so that they can understand what you're seeing and give even more helpful answers. Early sales have been strong, exceeding our expectations. We're going to have more to share on our glasses line-up at this year's Connect conference on September 23rd, so I encourage you to tune in then.

Before I wrap, I want to mention that I just published an op-ed about why I'm so optimistic that we're building a positive future for everyone. At Meta, we've always built technology to put power in people's hands so they can connect with the people they care about and shape their world in the ways that they want. It's why we've always focused on making our products affordable and accessible for everyone. And it's a strategy that has served both our community and our business well.

As we enter this next chapter, that same philosophy is guiding how we approach developing AI. We're the only major company building AI with the primary goal of putting superintelligence directly into people's hands. Rather than centralizing superintelligence, we're focused on distributing it widely and giving everyone the ability to direct it towards what matters to them. That's the way that society has always made progress. I think that these are the right values for building a positive AI future. And if we help build this, then I think we'll continue to build a very strong business as well.

That's everything I wanted to cover today. AI is improving our core business -- it's making our apps more relevant, and delivering better results for businesses. We're starting to deliver more novel products, and we'll have a lot more there soon as well. We're investing aggressively because the potential is huge and we know that there are many ways to deliver value here. As always, I'm grateful to all of you for being on this journey with us. And now, here's Susan.

Susan Li, CFO

Thanks Mark and good afternoon everyone.

Let's begin with our segment results. All comparisons are on a year-over-year basis unless otherwise noted.

Q2 Total Family of Apps revenue was $60.4 billion, up 28% year-over-year.

Q2 Family of Apps ad revenue was $59.4 billion, up 27% or 26% on a constant currency basis.

In Q2, the total number of ad impressions served across our services increased 14%. Impression growth was healthy across all regions, driven by growth in engagement and users, as well as ad load optimizations. The global average price per ad increased 12% year-over-year, driven by ad performance gains, improvements in macro conditions relative to Q2 of last year, and currency tailwinds. This was partially offset by strong impression growth, particularly from lower-monetizing surfaces and regions.

For the first time, quarterly Family of Apps other revenue reached $1 billion and grew 73% year-over-year, driven primarily by WhatsApp paid messaging and subscriptions revenue.

Within our Reality Labs segment, Q2 revenue was $431 million, up 16% year-over-year due to strong growth in AI glasses revenue, partially offset by lower Quest headset sales.

Moving now to our consolidated results.

Q2 total revenue was $60.8 billion, up 28% or 27% on a constant currency basis.

Q2 total expenses were $42 billion, up 55% compared to last year, and included $2.4 billion in charges related to legal proceedings and $1.2 billion in severance expenses in connection with the May 2026 headcount reduction. Year-over-year growth was primarily driven by increases in employee compensation, infrastructure costs, legal-related costs, and third-party AI token costs.

Excluding the previously mentioned severance expense, growth in employee compensation was driven by technical hires we've added over the past year, particularly AI talent.

The growth in infrastructure costs was driven by higher depreciation, data center operating costs, and third party cloud spend.

We ended Q2 with over 75,000 employees, down 3% from Q1. This total includes approximately 8,000 employees impacted by the May 2026 headcount reduction. We expect the majority of impacted employees will no longer be captured in our headcount by the end of Q3 2026.

Second quarter GAAP operating income was $18.8 billion, representing an 8% decline year-over-year and a 31% operating margin. Excluding the Q2 legal charges and severance expenses, our second quarter operating income would have increased 9% year-over-year.

Our tax rate for the quarter was 16%.

Net income was $15.8 billion or $6.18 per share.

Capital expenditures, including principal payments on finance leases, were $31.1 billion, driven by investments in servers, data centers, and network infrastructure.

Free cash flow was $784 million. We ended the quarter with $90.3 billion in cash and marketable securities and $83.7 billion in debt.

Turning now to the business performance. There are two primary factors that drive our revenue performance: our ability to deliver engaging experiences for our community, and our effectiveness at monetizing that engagement over time.

On the first, we continue to see significant gains from our content recommendation initiatives.

On Instagram, global time spent this quarter grew double digits year-over-year, largely driven by improvements to our Feed and Reels recommendations. On Facebook video time spent increased 9% globally year over year and over 10% within the US & Canada, where it was driven by ranking improvements.

We are finding that LLMs are increasingly capable of delivering ranking and recommendations gains. First, they make our existing systems smarter by understanding what the content is actually about and generating better training data. And second, LLM-powered agents are also helping with engineering development by evaluating content quality, detecting trends, and testing ranking changes.

Earlier this year we reached a milestone of every public Reels and Feed post on Instagram being automatically processed through an LLM and analyzed across dimensions from topics to tone, and we're working towards including more surfaces on Facebook as well. These signals can then be passed to downstream applications across ranking, recommendations, and content policy enforcement, which is a key building block toward greater personalization. This quarter we also began using our Muse family of models to conduct content understanding across signals like video topic classification and summarization, and we've seen positive early results.

Finally, our recommendations are also becoming more personalized, surfacing more fresh content while giving people more direct control over what they see.

On Reels, we shipped our largest single-release ranking improvement to date, combining faster inference with a new architecture that draws on deeper user history to improve predictions. This drove a 15 basis point increase in sessions on Instagram, with particular strength in reshares and time spent, which are both strong indicators of better content-to-user matching. We are now bringing this to Feed, where early results look comparable.

We are also getting new content to people more quickly. Investments we've made in more real-time infrastructure and modeling improvements on new videos are allowing our largest ranking models to now identify high-quality new Reels at creation. On Instagram Feed, over half of all recommended content is now less than one day old, more than double from a year ago.

We're also giving people more direct control of the content they see. Today, Instagram users can visit the "Your Algo" page, which lets users write natural-language prompts to tune their recommendations. Similarly, on Facebook, we launched "Shape Your Feed". Early results show over 80% retention among users who engage with it.

Looking forward, we are executing on our longer-term efforts to develop the next generation of our recommendation systems. This includes building foundation models that are designed to power organic content and ads recommendations simultaneously, as well as developing LLM-native recommender systems. We hit our first research milestone this half by continuously pre-training a large-scale model with recommendations data and observing healthy scaling laws in the process. We are encouraged by this milestone and expect continued progress in the second half of the year.

Turning to the second driver of our revenue performance: increasing monetization efficiency.

The first part of this work is optimizing the level of ads within organic engagement.

Here, we continue to enhance our systems to show ads at the optimal time and location. In Q2, we also expanded availability of ads on our newer surfaces, including completing our global ads expansion on Threads. On WhatsApp, we have introduced support for more types of ad destinations and advertiser performance goals in Status and continue on track toward our global rollout.

Moving to the second part of increasing monetization efficiency: improving performance for the businesses who use our services.

Within our ads systems, we're delivering performance gains as we deploy more complex and predictive models.

This quarter, we introduced Meta Generative Recommender, a paradigm shift in how our ads system works. Rather than scoring every possible ad individually, we are now using LLMs to reason about ad content and user preferences together, and predict the best ad for each person. This makes our ad matching more intelligent and more precise, which compounds performance gains for advertisers. We deployed the first generative model into our ads retrieval system and saw notable improvements in ad performance. Early pilots using LLMs to better understand user preferences drove a 1% increase in app event conversions on Instagram.

In Q2, we also advanced our user understanding models to analyze ads and organic activity and simultaneously improve both user experience and advertiser performance. Combined with our GEM model for ads ranking and sequence learning, these advancements generated an 8.3% increase in ad clicks and a 15.7% uplift in conversions on Facebook.

We're also leveraging AI to empower businesses to more easily manage their campaigns, develop ad creative, and engage with their customers.

Our AI-powered Advantage+ end-to-end solutions continued to grow, reaching over $75 billion in annual revenue run-rate this quarter. We are working to deepen adoption as advertisers who leverage multiple tools see compounding performance gains.

I'll share one example of how Advantage+ is making performance marketing meaningfully easier and more effective for SMBs. Underneat, an online apparel brand in India, had been setting up each campaign manually across Facebook and Instagram. After adopting Advantage+ sales campaigns layered with Advantage+ Audience, Placements, and Budget optimization, they saw a 13% incremental lift in purchases and a 16% increase in add-to-cart conversions.

Adoption of our Gen AI ad creative tools continues to scale, with over 9 million small businesses using at least one AI creative tool. Image generation, which now lets advertisers produce more creatives at scale from existing content, including a new ability to create images from video assets, saw adoption more than double this quarter. We also introduced a new end-to-end creative solution that gives advertisers the AI infrastructure to translate real performance signals into their next creative decision, while preserving brand identity and tone. We are building with agency integrations from day one so teams can diagnose, generate, and scale high-performing creative without leaving their existing workflows. Looking ahead, with the rollout of Muse Image, we expect to further expand advertisers' ability to generate high-quality, on-brand creatives at scale.

With Meta Business Agent, businesses are better able to serve their customers through our messaging apps by responding to inquiries, recommending products and handling support around the clock. Earlier this month, we also introduced the Meta Business Agent Platform, which gives Enterprises the infrastructure to build, customize and deploy their Business Agent at scale on WhatsApp. The platform provides larger businesses with enterprise-grade controls, guardrails and measurement built in so they can define rules and offer personalized experiences, starting within the messaging apps that their customers already use.

Movida, one of Brazil's largest car rental companies with nearly 400 locations, deployed a Business Agent on WhatsApp to handle the entire booking flow, from vehicle selection and pricing to payment, in a single conversation. Returning customers could complete a reservation in as few as three messages. In a one-month period, Movida reported a 44% increase in daily bookings through WhatsApp when compared to the same period in the prior year, and that 85% of conversations in the channel were resolved entirely by the AI agent without human assistance.

We are also building out additional ways to monetize our ecosystem. Two additional revenue streams are subscriptions and monetizing our competitive models through an API.

Meta One is an evolution of our subscription portfolio to create more value for everyday users, businesses and creators, so they get more features and AI tools to create, connect and stand out. We are excited to bring this to more users and continue to build enhanced tools for our subscribers.

We also recently launched a high intelligence model API at a competitive price and are encouraged by the initial results. We recently made Muse Spark available on OpenRouter for US based developers, broadening its distribution and making it easier for developers to adopt the model. We soon expect to roll out the model API to more distribution channels, make it available in more countries, and open it up for enterprises.

Our approach to building capacity is strongly influenced by several key elements:

First, the broad environment for building infrastructure is dynamic and uncertain in both near-term and longer-term time horizons. The industry has underbuilt historically for the wave of AI adoption, making existing capacity, including our own, extremely valuable. Longer-term, the supply chains need to be built out to support the capacity that we anticipate we and others will need for AI-powered experiences.

Second, we have high confidence in our ability to utilize capacity to scale and build on top of our existing experiences, as well as continue to invest in foundational models that will create substantial new opportunities. Consequently, our current plans are geared towards maximizing 2026 and 2027 capacity. When we have had incremental capacity in the past, it has proven extremely valuable in scaling experiences like Reels, and we are confident that this will be true in this time frame as well.

Longer-term, it is harder to predict the exact usage scaling curves, but we believe that our distribution advantages will give us the opportunity to serve AI products that are valuable for everyone: both our 3.6B users and millions of businesses. This should be true regardless of whether our models are on the frontier, but we believe that being on the frontier will unlock new markets and opportunities for which we may need additional compute.

Therefore, our longer-term capacity strategy aims to give us the flexibility to continue growing compute in 2028 and beyond by laying down data center and network foundations to accommodate future server decisions. The long-lived nature of these assets inherently provides the flexibility that will make it possible to adjust our investment to the pace of AI adoption. In addition, we have been making strategic investments in areas like our internal custom silicon effort, which will provide long-term strategic flexibility and supply chain leverage. This will be helpful in driving better returns on those long-term investments.

Finally, we believe that overall industry capacity is going to remain tight for the foreseeable future. As we’ve said earlier, we strongly believe that the models, consumer experiences, and enterprise offerings that we are building will be the best and highest ROI use of our infrastructure. Those enterprise offerings have the potential to take multiple forms as Mark mentioned – agentic tools, our API, or monetizing compute directly given outsized market demand. We expect that remaining nimble about these opportunities will help us fund our buildout more efficiently while preserving our strategic flexibility to have the compute when we need it and provide us multiple pathways to generate returns on invested capital.

In funding these infrastructure investments, the strength of our balance sheet gives us the ability to attract capital from a wide range of markets to supplement the cash flow generated by our business. Our announcement with BlackRock yesterday is an example of the partnerships that we can structure to complement our approach to building infrastructure capacity.

Moving now to our financial outlook.

We expect third quarter 2026 total revenue to be in the range of $61-64 billion. Our guidance assumes foreign currency is an approximately 1% headwind to year-over-year total revenue growth, based on current exchange rates.

Turning to the expense and capex outlooks.

We are raising the lower-end of our expense outlook to incorporate the $2.4 billion charges related to legal proceedings recognized in Q2. We now expect full year 2026 total expenses to be in the range of $165-169 billion.

We continue to expect to deliver operating income this year that is above 2025 operating income.

We anticipate 2026 capital expenditures, including principal payments on finance leases, to be in the range of $130-145 billion, narrowed from our prior outlook of $125-145 billion.

Absent any changes to our tax landscape, we expect our tax rate for the remaining quarters of 2026 to be between 15-17%, an increase from our prior outlook of 13-16%.

Finally, we continue to monitor active legal and regulatory matters that could significantly impact our business and financial results. For example, we continue to see scrutiny on youth-related issues in several markets and have a number of youth-related trials scheduled for this year in the US, which may ultimately result in a material loss.

In closing, our business momentum continued in Q2, with strong execution across our core ads and engagement initiatives. We're also progressing in our efforts to bring personal superintelligence to everyone, with exciting model releases and we expect to build on that momentum over the course of this year with new products.

With that, Krista, let's open up the call for questions.