Meta Platforms, Inc. (META)
Fourth Quarter 2025 Results Follow Up Call
January 28th, 2026

Operator: Good afternoon. My name is Krista, and I will be your conference operator today. At this time, I would like to welcome everyone to Meta's Fourth Quarter and Full Year 2025 Results Follow-Up Q&A Call.

All lines have been placed on mute to prevent any background noise. After the speaker's remarks, there will be a question-and-answer session. If you would like to ask a question, please press star one on your telephone keypad. To withdraw your question, again press star one. We ask that you limit yourself to one question and this call will be recorded. Thank you very much. Kenneth Dorell, Meta's Director of Investor Relations, you may begin.

Kenneth Dorell: Thank you. Good afternoon, and welcome to the follow-up Q&A call. With me on today's call are Susan Li, CFO, and Chad Heaton, VP of Finance. 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 back over to Krista for the first question.

Operator: Thank you. We will now open the lines for a question-and-answer session. To ask a question, please press star one on your touchtone phone. To withdraw your question, again press star one. We ask that you limit yourself to one question. Please pick up your handset before asking your question to ensure clarity. If you are streaming today's call, please mute your compute speakers. And your first question comes from the line of Youssef Squali with Truist Securities. Please go ahead.

Youssef Squali: Hi, great. Thank you, guys, for taking the question. Maybe Susan, this is something you mentioned in passing on the main call. But so, the EU Commission found Meta's consent or pay model in breach of the DMA. Is there a new compliant framework for the EU in terms of product changes, timing, pricing, and any way to bracket maybe the revenue and expense impact for '26, including a potential -- any shift to maybe less personalized ads?

And then just quickly on the operating income guide. Are there any guardrails you can share in terms of the guidance commentary for your '26 versus '25? I guess if I had just the same operating income dollar amount year on year on say 20% top line growth in '26 would imply something like mid-13s -- 30s margins versus 41% in '25.

Could we see that kind of margin compression, or is there a minimum margin level you do want to maintain? Thank you so much, Susan.

Susan Li: Thanks, Youssef. So, on your first question about less personalized ads, we aligned in December 2025 with the European Commission on further changes to our consent model for personalized ads in Europe. So, in the coming days, users in Europe will see the changes to the less personalized ads user flow.

We continue to think that our aligned solution goes far above and beyond what's required by the DMA, and our appeal of the Commission's decision will continue in the courts. And with this said, we can't rule out that regulatory authorities or courts could seek further modifications to our less personalized ads offering.

So, if that were to be the case, it's possible that those changes could result in a materially worse user experience in the European Economic Area in Switzerland. But we're in dialogue with regulators, and we don't have an update on when we will reach full resolution on this. The expected impact from the updated user flows that we will be rolling out this quarter is factored into our outlook. So, that's on the LPA.

On your second question, we don't manage to a target margin. We believe that we are in a very fortunate position as a business to both have very strong revenue growth and just healthy business overall coming into 2026, seeing strong revenue growth.

And also, fortunate to have a lot of very compelling opportunities to both build what we believe is going to be transformative technology, but also user products and experiences that we believe over the long run will be not only very cool technological experiences, but also good business opportunities.

So, we're reinvesting a lot of the revenue into these opportunities now. Again, I think I mentioned this on the call, but they're primarily in the form of either AI infrastructure or talent. Both are pretty dynamic in terms of what is available in 2026. And in fact, infrastructure has both the dimension of when some of our already planned capacity is able to come online, as well as our ability to secure new capacity.

So, it's a very dynamic planning process. The framework of growing consolidated operating profit over time is the one that we're using. And as we've mentioned in the past, that in and of itself, of course, is lumpy. Over the fullness of time, merely growing operating profit would not make us a very attractive investment opportunity relative to other equity investments.

But we believe we have historically delivered above that, and we believe we will over the long run, but that it will be lumpy and there are periods where it could be pressured because of the magnitude of investment opportunities ahead of us.

Operator: Your next question comes from the line of Stephen Ju with UBS. Please go ahead.

Stephen Ju: Hi, Susan. Thanks for taking the question. I think you have not offered formal guidance for 2026. I usually don't look at an annual outlook, but there is an implied range of outcomes for what the revenue might be for 2026, even if we're sitting here assuming a flat operating income dollar number for '26 over '25. So, I mean, just running through the math here, it seems like you're -- that the dollar growth of revenue is going to be somewhere close to $50 billion.

I think '25, you did about $35 billion. A year prior, it was less. So, you're seeing this pretty incredible pickup of the business performing that much better. So, I guess I'm sitting here trying to assume some of the -- think about what the drivers of a lot of the growth in '26 will be. I know you talked about, I guess, some of the pieces that go into this from greater engagement, as well as, I guess, the greater efficiency and the ad serving, targeting, etc.

And I think your shareholders tomorrow are going to be trying to figure out how much of the growth in '26 will be sustained in '27, et cetera. I know you're not giving guidance for '27 or beyond, but just help us understand how durable a lot of this growth is going to continue to be, especially as you continue to unlock a lot of innovation from all this CapEx and OpEx that you're incurring this year. Thanks.

Susan Li: Thank you for the question, Stephen. So, as you said, we are not giving full-year 2026 revenue guidance. But having said that, we have come into 2026 to date against a very healthy macro backdrop. I alluded to this on the first call, talking about how strong the holiday marketing season was. And we've seen the continued macro strength continue into the first weeks of 2026.

We've also seen good results from the set of investments that we made in 2025, specifically to fund work to drive ad performance improvements and organic engagement initiatives. And we have funded a similar set of investments in 2026.

One of the things that always weighs on me in this process is sometimes it's hard to know for certain where you are on the marginal return curve of how these different initiatives add up. Our ability to robustly measure the expected return from any individual project, I think, is pretty reasonable. But there can be more uncertainty when you sum them up. And what we've seen in 2025 is, we've seen good aggregate returns from those efforts.

I think at the high end of the 2026 range, you'd see good aggregate returns also from the efforts we're funding this year. But at the same time, the slope of the curve could be steeper than we think. There could be volatility in the macro landscape. There are really a lot of variables and moving pieces.

So, while again, I feel very good about the business fundamentals and the work that we're doing, there is certainly a wide range of uncertainty through '26 and even more so, obviously, as you go into 2027.

Operator: Your next question comes from the line of Benjamin Black with Deutsche Bank. Please go ahead.

Benjamin Black: Great. Thank you for taking my questions. I wanted to dig in a little bit into the Reality Labs losses. So, you mentioned they should be relatively stable in '26 versus '25 and then moderate from 2027 onwards. Just curious what's the driver of this? Is it associated with somewhat narrow focus on wearables? Is it associated with faster revenue growth? Is it something else? It'd be good to hear a little bit more there.

And then just quickly on cost per action for your advertising partners. I know it's a challenging thing to measure, but it would be great to hear some updates in terms of how that is trending. Thank you.

Susan Li: Ben, on the first question about Reality Labs, we continue to have optimism in the future of VR, and we are investing continually -- going forward, rather, in building future headsets. However, consumer adoption of VR has generally been on a slower growth path than wearables, and we are rebalancing our Reality Labs portfolio to reflect this.

So, we are meaningfully reducing our investment in VR and Horizon this year, but we're growing our investment in wearables to capitalize on the momentum that we're seeing in our position as a market leader. That's the shape of the rebalancing of the portfolio so far.

Now, beyond 2026, our expectation is that we are going to gradually manage Reality Labs towards a lower operating loss envelope, both because we expect to benefit from more mature supply chains as we scale the wearables in VR ecosystems, and hopefully generate higher-margin revenue lines and continue to operate more efficiently over time.

Having said that, that's what we're planning for, but it's certainly possible that either changes are as not linear because of unforeseen changes to our roadmap, or if the market develops differently than we expect. So, I think our expectation is for Reality Labs losses to trend down from 2026 levels, but it's hard to characterize the exact shape of that line.

Chad Heaton: And on the second part of your question about cost per action for advertisers, so overall, we're continuing to drive ad performance improvements. One measure we look at to monitor this is conversion growth. It's a complicated metric because advertisers optimize for different conversions with different values, including lower-value brand objectives like views, as well as higher-value conversions like sales.

So, we try to control for that and look at what we call value-weighted conversion rates. And when we do that, we're seeing very strong year-over-year growth with conversions growing faster than impressions. In fact, we saw meaningful acceleration in year-over-year conversions growth through Q4, and that was driven by our ad performance improvements that we made.

So, then getting to your specific one, we look at another metric, cost per action. This is also a complicated metric for the same reasons. There's many different objectives that advertisers optimize for, and they have different values. But when we try to control for that, we also saw healthy cost per action trends for advertisers in Q4 as well.

Operator: Your next question comes from the line of Shweta Khajuria with Wolfe Research. Please go ahead.

Shweta Khajuria: Okay, thanks a lot for taking my questions. Let me try two, please. First is on Meta Compute. So, is it fair to assume the level of capital intensity to be similar as we think about it going forward, given the commentary on tens of gigawatts of capacity coming online over the decade that translates to a very similar level of capacity growth? That's question one for at least foreseeable future.

And then question two is on impressions growth, a follow-up to the prior one. Could you please talk about the key drivers to pretty healthy impressions growth and the durability of that? Thank you.

Susan Li: Yes, so I think Meta Compute reflects the fact that we think based on our best understanding of the current scale at which models are scaling up, that we will want to have the option of having a lot more compute in the years to come. It is somewhat premature to say with certainty what those levels are.

This is a very dynamic area of planning, even just for 2026, much less the further-out years. But our focus is on giving ourselves the ability to scale compute up, and to do so in a way that will be at better prices than we're able to do today. And so, we're trying to build in option value at many points along the way.

In terms of where we are focused right now on driving down the cost of scaling compute, on silicon, that's obviously one of the big cost drivers. We're working to do that today through a variety of means, including diversifying our chip strategy so we can get the greatest cost efficiency for the workloads that we need to support.

For example, we run a number of workloads with different requirements for compute, memory, and networking. And we're focused on deploying the optimal chips for each of those workloads to deliver the best performance per watt and total cost of ownership. And we're also continuing to expand our MTIA custom silicon program to support our core ranking and recommendation inference and training workloads.

So, silicon is a big effort. I think over the longer term, we'll also be looking at other areas like reducing the cost of producing energy. That's not a near-term bottleneck, but could be depending on how compute scales over time.

So, Meta Compute is really intended to take a strategic view over a longer-run arc of what our compute needs could be, where the most likely supply

bottlenecks and cost efficiencies could be, and identify solutions to help give us better options in each of those areas.

Chad Heaton: And then to your question about impressions, so starting at the top, worldwide impressions grew 18% year-over-year in Q4. That was driven mainly by growth in users and engagement per user. Ad load increases were also a tailwind, although they contributed a minority of the year-over-year growth.

On engagement, video was the largest driver of engagement gains in Q4, with particular strength on Instagram, where Reels time grew more than 30% year-over-year globally. I'd also note that Facebook saw healthy double-digit growth in video time in Q4 as well.

On ad load, the tailwinds that we had in Q4 were driven by the higher levels of ad load on both feed and video surfaces on Facebook and Instagram. This was driven by some optimizations that we had made toward the end of Q3 and then continued into Q4. As Susan noted, we also had revenue gains from redistributing ad supply more optimally across users and sessions, though those show up in the reported pricing metrics rather than impressions since the optimizations don't grow absolute ad load.

The other thing I would note is we did see the three-point acceleration in worldwide impression growth. That was driven by ad load optimizations on Facebook and Instagram feed and video surfaces, in addition to the engagement strength on Instagram. The total year-over-year growth continues to be driven by growth in users and engagement, again, where ad load is a secondary contributor.

Operator: Your next question comes from the line of Barton Crockett with Rosenblatt. Please go ahead.

Barton Crockett: Okay, thanks for taking the question. I was wondering about the commentary that Mark made about your workforce getting more efficient, moving towards higher talent, flatter, fewer people doing work.

That sounds like a formula that could lead you to a position where your headcount growth would slow, flatten, reverse. I'm just wondering if I'm reading that correctly. And if so, with the big growth in OpEx that you're seeing this year, is this a situation where other things are offsetting headcount growth, like more comp per person or spending on other things?

Susan Li: Thanks for the question. So, we are currently expecting to continue growing headcount in 2026. That growth is happening in a targeted way across the company's priority areas, including infrastructure, monetization, AI, amongst others. That talent is also concentrated in technical roles that carry a higher compensation profile. So, it skews the -- it continues skewing the mix of talent over time, and we expect compensation expenses will continue to grow faster than headcount in 2026.

I should say on the broader point about the impact of AI on our hiring plans, we expect that there are going to be big improvements in productivity across our workforce over the next few years. But depending on how those improvements play out, we may want to add more people if, for example, each engineer can deliver greater impacts than they would have previously.

So, we don't know exactly how this will net out, but we're very focused on making sure that we lean into the productivity tools that exist, make them widely available, make sure that we are able to identify folks who are using them really well, and make sure that our organization structure is evolving over time to meet the needs of how people work in a new AI-enhanced landscape.

Operator: Your next question comes from the line of Deepak Mathivanan with Cantor Fitzgerald. Please go ahead.

Deepak Mathivanan: Susan, two questions from us. So, first, how should we think about your reliance on third-party cloud partners? Is this a temporary arrangement until Prometheus and others come live? Or do you expect third-party cloud partners to always be a significant part of compute strategy, maybe in similar magnitude of annual spending also beyond 2026?

And then second one on the CapEx guide, this is a broader range than your typical outlooks. What is bringing this uncertainty? Is it tied to some sort of like a cost inflation, either on data center side or memory or other components? Or is the uncertainty more due to the need to forecast compute early in the year?

Susan Li: Great questions. On cloud, we still anticipate investing significantly in our own owned-and-operated and leased data center capacity. The CapEx guide clearly points to that.

But a lot of that CapEx is for capacity that doesn't come online until 2027 or beyond. So, we have been signing cloud deals to enable us to bring capacity online more quickly this year to alleviate our current capacity constraints.

And I would say when we are evaluating more generally whether to build our own capacity or use public cloud, we are looking at a variety of factors, and that includes cost, timing, scalability, and other factors that we think about in terms of our ability to customize and configure, for example, the capacity to our needs.

In general, investing in our own data centers gives us greater customization and efficiency and, of course, secures long-term supply. But cloud has other advantages, including the ability to bring capacity online very rapidly, especially if cloud providers have pre-staged capacity available with shorter lead times.

In terms of the CapEx range, it is very early in the year. We are four weeks in, and the infra capacity process was very dynamic last year, as I'm sure you all remember, and it remains very dynamic this year. So, we are trying to reflect that in our wider range. We are really trying to scale up our capacity

significantly to support the range of capacity demand that we anticipate that we might have in '27 and beyond.

But there are a lot of variables that are fluctuating that will influence where we land in the range, including things like the availability and pricing of different components, servers obviously, but other components too, such as memory and storage, as well as the timing, and magnitude of leases that we sign this year for future capacity.

Operator: Your next question comes from the line of Michael Nathanson with MoffettNathanson. Please go ahead.

Michael Nathanson: Thanks, Susan. I have two following up on Deepak's question. I was going to the same place.

I know this is hard to get to, but when there's a steady state world of capacity that's kind of fully up and running, do you think it's margin enhancing at the end of the day? Because we're in this weird state now where you're basically leasing versus getting enough capacity on your own. So, I wonder how you think about steady state, moving more of your demand to your own data centers. Is that margin incremental at that point in time?

And then just on agentic commerce, this is a topic -- commerce is a topic that Meta has talked about for a while. Do you anticipate as people adopt agentic products that you will charge a fee from your commerce partners to use them or will they be free? So, I'm trying to figure out how you monetize all those great agentic tools you're going to get. Thanks.

Susan Li: Yes. On your first question, I chuckled a little bit when you said steady state world. I think we are, if nothing, not in a steady state world, we're in a remarkably dynamic one.

And so, right now we are really focused on trying to make sure that we have secured capacity to meet demand for the different time horizons that we're planning for. Again, for the nearest term, we're bringing cloud capacity, where available, online. Over the longer term, we believe that having owned and operated capacity gives us, again, more flexibility to configure the capacity exactly to our needs.

That's kind of complimented by the longer run compute strategy also that we alluded to earlier, which is how do we ensure that we are really -- we have a chip strategy that enables us to really get the maximum performance per watt and per dollar out of our chips so that we're running the right chip for the right workload. And those really go hand in hand as we think about them.

Chad Heaton: On the commerce question, I guess I'd first just note that commerce continues to be a big area of focus for our teams. Today, we already offer tools to help businesses connect with shoppers and boost sales through our ads products, including through our AI-powered Advantage Plus shopping campaigns.

We're investing in new product experiences to make it easier for people to discover and buy new products across our apps. But I think part of your question was also just around how we're going to go about monetizing these new AI revenue streams. And this is just an area we're going to be exploring.

We're going to be looking at subscriptions for high-value AI features. One feature of that will be Manus, as we've talked about. We intend to continue to scale its existing subscription offering to many more businesses.

We'll also integrate it into our products and bring a leading agent to billions of people. The other thread here is business messaging, where we're starting to see traction with the business AIs we've made available in Mexico and the Philippines. We've seen over 1 million weekly conversations between people and business AIs now happening.

So, these are just some of the areas that we're starting to explore, but I think there's a lot of opportunities both across commerce, agentic commerce, and then some of these new revenue streams that we're looking into.

Operator: Your next question comes from the line of Justin Patterson with KeyBanc. Please go ahead.

Justin Patterson: Great, thank you. Good afternoon. I wanted to go back to just the improved productivity per engineer.

How closely do you think you are toward reaching a ceiling around that? And how's that influencing the pace of product velocity versus what we've seen in the past? And then just on a separate topic, Reels did launch on TV in Q4. I'm curious to hear how you're thinking about both the time spent and ad opportunity around that. Thank you.

Susan Li: Yeah, on your first question about developer efficiency, so we talked about a 30% increase in output per engineer and the sort of power users, folks who are really leaning into and most effective at using AI coding tools, their output has increased more. That was the 80% year-over-year stat.

The big change that we saw in the second half of 2025, particularly in Q4, has been the increased effectiveness of agentic coding tools and accompanied by the growing internal adoption of them as the tools have become more capable.

And the percentages we cited are really talking about the number of code changes landed per engineer driven, again, by the continued adoption of agentic coding. I think there's a lot that we're going to learn over the course of this year and next in terms of how much more productive engineers can be. I would say that the tools themselves, I think, certainly will have room to grow and become more effective.

And I also think that our engineers will also, are probably early in the life cycle of kind of learning how to use them with maximum effectiveness. We are also starting to see functions outside of engineering, such as product managers and

designers be able to engage in more of the technical work because of coding tools.

So looking ahead this year, we're going to integrate these coding tools more deeply into our developer environment. We're going to give them access to more parts of our internal environment and other tools and context and try to deploy more custom models that are trained on our internal codebase.

Chad Heaton: And then in terms of the launch of Instagram on TV, so that's a fairly recent launch of IGTV with select partners. And we are in the early stages of evaluating the results. But overall, this is something we're excited about, excited to bring Reels from our favorite creators to the big screen so that our users can enjoy those with their friends.

This is something we've heard from our community - that watching Reels together is just more fun. This is a test that we've designed to learn which features make that experience work best on TV. Over time, we're going to explore new features here, including using your phone as a remote, find intuitive ways to channel surf, share feeds with friends, et cetera.

So, this is definitely something that we are early on but excited about. I think it'd probably be premature to talk about impacts on things like time spent or advertising opportunities.

Operator: Your last question comes from the line of Tom Champion with Piper Sandler. Please go ahead.

Thomas Champion: Thank you. Good evening. Susan, could you just talk about the impact of the OBBB in your CapEx guide for '26? Just curious if that had any impact.

And then, Chad, maybe for you, could you offer up a Manus for dummies or what do you see in the asset that is truly unique and strategically important going forward? Would love to just get your thoughts on that. Thank you.

Susan Li: Thank you, Tom. On the question about the OBBBA, we expect substantial cash tax savings from the new U.S. tax laws given the significant investments that we're making in infrastructure and R&D. So, we expect that that will benefit our cash taxes paid in 2026.

On the question about Manus, generally we are excited about Manus on multiple fronts. First and foremost, we believe that this acquisition is going to accelerate the development of new technologies, which we think will benefit businesses as well as our users across the world.

We are both going to leverage, we'll both keep sort of the existing Manus business running as is, but we're planning to leverage the technologies that Manus is building to help power the business agent work that we are building, which is a big driver of sort of new revenue opportunities for us.

And we expect that we're going to use it to make our AI tools broadly more effective. So, we're very excited about it. It's a team of exceptional talent who is joining us, and we think that general purpose agents will be a very important part of the overall consumer and business product experience in terms of what we're able to enable with AI.

Kenneth Dorell: Great. Thank you, everyone, for joining us today. We look forward to speaking with you again soon.

Operator: Ladies and gentlemen, this does conclude today's conference call. Thank you all for your participation, and you may now disconnect.