Meta Platforms, Inc. (META)
Second Quarter 2026 Results Follow Up Call
July 29th, 2026

Operator: Good afternoon. My name is Krista, and I'll be your conference operator today. At this time, I would like to welcome everyone to Meta's Second Quarter 2026 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 then the number 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. Chad Heaton, Meta's Vice President of Finance, you may begin.

Chad Heaton: Thank you. Good afternoon, and welcome to the follow-up Q&A call. With me on today's call is 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 back over to the operator for the first question.

Operator: Thank you. Your first question comes from the line of Ron Josey with Citigroup. Please go ahead.

Ron Josey: Great. Thanks for taking the question. Susan, I wanted to ask a little bit more about how the team balances just future investments, given all the conversations about '26 and '27 capacity relative to just managing free cash flow. I know there's a lot of answers here, but any insights on how you balance that?

Of course, there's a lot of questions around potential investments going forward. But then as a follow-up to that, just any insights or help us understand how the team goes about partnering with, call it a partner, like you mentioned, BlackRock on the call versus doing it on your own would be helpful. Thank you.

Susan Li: Hi, thanks. Thanks, Ron. So on your question about how we -- I'm going to -- I'm just going to make sure I got your question correctly. I think you're basically saying how do we think about what the right investment and CapEx levels are, relative to the way we think about managing free cash flow.

I think we were -- I -- as I mentioned on the call, we're obviously still working through what our capacity needs are going to be over the coming years. I think it is not new news that we are investing in CapEx in '26 and '27. And as I mentioned, we generally believe that near-term capacity is more valuable than long-term capacity.

We think that we have a lot of good opportunities to deploy this capacity toward, again, a combination of plenty of opportunities in our core business. We talked about the ranking and recommendations roadmap and how that's being enhanced by using LLMs to understand content more deeply, using AI-powered agents to actually do recommendations work is the next leg of work there.

And then outside of the core, a lot of new -- newer AI-driven experiences, for example, using Muse Image to help advertisers develop ad creative is something that we're quite optimistic about.

And then, of course, in the longer term, building more fundamentally new and agentic experiences. That's all mostly consumer-facing. We also talked about the enterprise-facing offerings that we're looking at also on the call, whether it's business agents that interact with the consumer agents that I mentioned, whether it's an API or compute.

So we really think that there are going to be a lot of opportunities for us to monetize that compute. And we are balancing that against, of course, making sure that we're delivering attractive financial results and we're bringing shareholders along with us.

I would say on your second question, which is about when we evaluate partnerships, we look at partnerships really to gain more flexibility for our compute plans. Like we said, there's a -- I think there are a lot of open questions in terms of '28 and beyond and how to think about what our capacity needs might look like there. And the partnerships really are structured in a way to give us a lot of flexibility in terms of making longer-term compute decisions.

Operator: Your next question comes from the line of Youssef Squali with Truist Securities. Please go ahead.

Youssef Squali: Thank you so much. Susan, just following up on Ron's question. So you guys have accelerated your compute capacity buildup through these JVs. I think you have one with Blue Owl, Hyperion, the one you announced last night with BlackRock, where you backstop some of these projects. Just how should we be thinking about these off-balance sheet obligations? And would you consider using equity to potentially fund some of these investments over time?

Susan Li: Yes, as I mentioned just now, again, these partnerships are really structured in a way that gives us flexibility to evaluate every four years what our ongoing compute needs are. And the value of the backstop also goes down over time. And I would point you to the disclosures in our queue to look at what the backstop arrangement looks like in more detail.

We're not at -- we don't have anything -- we don't have any plans to share in terms of the second half of your question. I think we will -- as we plan for our investments in 2027, I think we will certainly evaluate all of the financing options. But we will have more to share when we -- we'll have more to share as we go into '27. I think for now, again, our operating cash flow is really the primary pillar by which we are financing all of our ambitions.

We also again have -- we've been raising investment-grade debt. We've looked at these partnerships. I think we will continually be evaluating what is the right funding source and what is the right cost of capital for us as we look at our 2027 and beyond plans.

Operator: Your next question comes from the line of Mark Mahaney with Evercore. Please go ahead.

Mark Mahaney: Hey, thanks. I'll ask two questions, please. did you comment on why you tweaked up the CapEx? I know it was small, but was that just rising input costs, or was that an acceleration of capacity build plans?

And then just briefly talk about that FoA other revenue that hit $1 billion. The primary driver of that 73% year-over-year growth was what, and how sustainable is the growth like that? Thanks a lot.

Susan Li: Thank you. On the first question, we're just updating the guidance, frankly, based on the visibility we have three months further since the last time we issued guidance. And there's less time remaining in the calendar year forward, so we trimmed the guidance range or narrowed the guidance range to make it a little bit more precise based on our current visibility.

On the second question, which is about Family of Apps other revenue growing, there are a couple things I would call out here. So paid messaging year-over-year revenue growth, that continued to grow at a very healthy rate in Q2. That, along with the growth in subscriptions revenue, drove Family of Apps other revenue up to 73% year-over-year.

And really, that paid messaging growth is driven by the continued growth in number of active businesses, the scaling of paid messaging volume. Marketing messages remain the largest driver for us of paid messaging revenue growth, and we're seeing strong traction with optimized delivery.

The majority of marketing messages now flow through optimized delivery APIs, which take into account whether a person is likely to read or engage with a business's promotional messages when deciding who to deliver messages to.

We're also seeing rapid growth in the volume of utility and authentication messages sent too, and that is a meaningful opportunity, particularly in key markets like India and Brazil, and will remain a focus for us through the rest of the year.

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

Michael Nathanson: Hey, Susan. Can I just dig some more on the compute questions? Here's my series of questions for you. What are the [Indiscernible] building compute that you think will be cleared post-2027? If they're cleared for you, will they be clear for everyone else? And do you expect people to overbuild?

And then just longer term, we've always focused on consumer and excel at consumer. What do you think your strategic advantages would be on compute versus a more crowded field than maybe the social media markets you guys have competed in previously? So thanks.

Susan Li: Thank you for the question. There were a couple components to it, so I will hopefully have addressed them all. I'll do my best here. I think one of the things I mentioned in my remarks is, again, right now we are supply constrained as it pertains to compute, and we expect that to be the case for the foreseeable future, based on the use cases that we have in the core business, and as well as based on where we think our training and inference needs will go and our expectation for having -- building and scaling consumer products that leverage the AI technology that we are building.

So I -- what I would say is we feel pretty confident that there will be a lot of good use cases for the compute we're building in '26 and '27. Obviously, that visibility declines as you get to '28 and beyond. And that's why the '28 and beyond roadmap is more oriented around having the flexibility to scale up or down compute as we need it, when we get closer to that time.

So I think that's -- I think -- I know I think you asked the question, maybe do you need to overbuild? And our hope is that we are actually building, in fact, to absorb all of these good high ROI use cases that we expect to have.

I think there was an -- you had another part of your question, which is what are our strategic advantages? In part, given how central infrastructure has become to our cost structure, we're very focused on identifying ways to drive down the cost of scaling compute. We are 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.

We run a number of workloads with very different requirements for compute, memory, and networking. And we're focused on deploying the optimal chips for each of them to deliver the best performance per watt and total cost of ownership. And we've struck deals with different chip providers over the course of this year as part of the strategy. And beyond silicon, we're also looking at other areas like reducing the cost of producing energy.

So I would say that we have, in addition to our ability to build out our compute capacity at scale, we also believe that we are making strategic investments and have the leverage, frankly, to do so in a way that should bend the cost curve on compute over time.

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

Benjamin Black: Great. Thanks for taking my question. So I guess, what are the factors that you're considering when you're deciding between leveraging compute internally versus renting it out to third parties? And you spoke about talking to interesting -- interested parties. Are you discussing renting out the capacity in 2027 and beyond, or are these more near-term conversations?

And then secondly, you listed token costs as one of the drivers of the year-on-year increase in operating expenses. At what point should we expect owned infrastructure to maturely reduce the need for third-party cloud? Thank you.

Susan Li: As I mentioned, I think, in my first -- in the prior answer, we're really focused on using the compute for our own training and inference use cases first. Beyond that, we absolutely think that there will be many opportunities for us to sell intelligence, and a lot of the enterprise offerings that we -- that we're building are focused on this.

And right now, as Mark mentioned, the just market for pure compute is very, very strong. And we -- he talked about us getting offers for compute at multiples of what we paid for it. We are evaluating all of these opportunities. We think that we are going to have, again, multiple pathways to generating return on invested capital here, but we don't have a more definitive timeline to share right now.

And then I think on your second question, we still anticipate investing significantly in both owned and leased data center capacity. But that -- a lot of that is capacity that comes online in future years. And so the cloud deals that we have been signing are really in service of bringing on capacity more quickly this year and next year, getting back to the point we mentioned about focusing on compute in '26 and '27.

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

Shweta Khajuria: Thank you for taking my questions. If I could try two, please. Susan, would it be possible to give us a framework of how large the business AI opportunity is? How big of a revenue opportunity is it in the long term? And also is it one of those sources of revenue that can scale quickly in the near to midterm? Is that a fair way to think about it, given that you have a pretty good sizable SMB customer base?

And then the second question I have is just on regulatory risk. You did call that out. How should we be thinking about perhaps the downside scenario of what the regulatory risk looks like combined U.S. and international and EU when you think about your stress testing scenarios? What is the worst-case scenario? Thank you.

Susan Li: Thank you, Shweta. This is a great question. The -- I think broadly speaking, given the large advertiser base that we have, we think building agents for businesses is an extremely adjacent opportunity to that. And because, again, of the fact that so many advertisers touch our platform every day, we have the distribution platform in terms of making sure that they have access to this offering.

In the very near term, we announced this past quarter the global expansion of Meta Business Agent for SMBs. We launched the Meta Business Agent Platform that gives enterprise clients the infrastructure to build and deploy business agents at scale.

And as of June, Meta Business Agent now serves over a million active businesses, helping enterprises and SMBs automate sales and manage customer conversations while maintaining engagements that are personal and on-brand for them across WhatsApp and Messenger. And we're now beginning to roll it out to Instagram Direct.

And we also launched the Business Agent Platform earlier this month, which includes a suite of APIs and tools that allows enterprises to customize their business agent, integrate it through their existing systems, so they can connect to systems like catalog, CRMs, inventory management, giving their business agents the ability to take action on behalf of the business and to deploy those conversations on WhatsApp.

So we are just getting started here. We're seeing a lot of demand. The case study that I referenced in my prepared remarks about Movida, I think, showcases both just the power of WhatsApp globally and how that compounds to accrue value to business agents.

So we're very excited about this. We think the product is getting ready to scale. We're investing in driving greater awareness of it, while also continuing to improve the discoverability and onboarding and user experiences. And we'll be introducing more capabilities into the business agent over the rest of this year and into '27. And we'll also keep improving our models, which we think will drive increased performance and demand for Meta Business Agent going forward.

So overall, I think this is an opportunity that we think is pretty unique. We are building a turnkey solution that's going to leverage existing social media posts, ad campaigns, web presence, making it really easy for businesses to set up. And we want Meta Business Agents to work day one for each business. And we think there will be a big opportunity here again, given the size and scale of advertisers on our platform.

On your second question, which was about how to think about the downside of regulatory risk. This is -- it's early. I think it is a little premature to comment on any ultimate outcomes. Obviously, there's ongoing litigation. But right now, I would say that again, the litigation is ongoing. It's hard for us to comment. These are untested legal theories, and we plan to defend ourselves vigorously.

And in the broader regulatory landscape, we are working hard to partner with regulators and make sure that we are -- we're a collaborative force in the industry in terms of getting to the outcomes that we all want.

Operator: Your next question comes from the line of Rob Sanderson with Loop Capital Markets. Please go ahead.

Rob Sanderson: Yes, thank you. Thanks for taking my question. Susan, I want to go -- you were describing ways that you're using LLMs now to help with ranking and recommendations. And Mark was talking about building LLM-native systems. So that sounds like a pretty significant change to a transformer architecture at the core.

So I guess I have a couple of questions around this. Can you speak to some of the complexities of making such a significant change at the core algorithm level? Can you say anything about enthusiasm for potential performance on this?

These shifts have been more than incremental for other use cases. And is there similarly large potential here for your recommendation systems? And then any color on time frame, should we be thinking about these as multi-year, two, three year initiatives, or can some of these changes happen more quickly?

Susan Li: [Indiscernible] a lot of runway generally to drive gains from improving our existing systems. And we've actually -- we have built a team within our Meta Recommendation Systems organization that's really focused on researching and building the next generation of our recommendation system.

So the LLM native recommendation system that Mark alluded to, our MRS research teams are working to combine the world knowledge and reasoning ability of LLMs with the scale and speed of our existing systems.

Our current recommendation systems are very, very good at leveraging user interaction histories, but unlike LLMs, they don't have the ability to reason about content or user interest from first principles.

So part of this work involves developing an LLM native data infrastructure where we shift from using, for example, content IDs for ranking, which have a limited amount of information, to a more descriptive and meaningful semantic ID system that allows our models to just understand more about each piece of content and infer why it's interesting to someone.

And eventually, we hope to get to a place where we can collapse our current multi-stage recommendation model to a simpler unified model that, similar to how LLMs do next-token prediction, can directly produce a set of tokens to show users, which we can then match to either organic or ads content within our inventory. So we've already begun early validation work on LLM-native recommendation systems, and we'll look to ship components of that to complement parts of the current stack.

And the other thing that I would just highlight, I think I alluded to this earlier, but we're also seeing AI agents contribute meaningfully to recommendations research workflows, and that makes our machine learning developers more efficient and productive. And we're gradually building towards having AI agents help improve our recommendation stack autonomously by doing more independent and self-driven recommendations work over time.

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

Deepak Mathivanan: Great. Susan, can you talk about where Muse Video and rank -- Image rank in terms of priorities for MSL? Mark touched on this a little bit on the call, but it seems like there is obviously a large opportunity to follow Meta with frontier, image, and video model. Are current capabilities of Muse Video and Image sufficient to capture this opportunity? And how should we think about the roadmap to deploy these opportunities at scale?

And then second question on the third-party compute capacity. You've signed several of these over the past few months. Can you talk about how these are getting integrated with your own first-party capability, and how the teams at Meta are using it efficiently in terms of making the capacity available for a wide range of use cases internally? Thanks so much.

Susan Li: Yes, I mean, we're pretty excited about the release of Muse Image and Muse Video. Muse Image is our first media generation model from MSL, and it uses agentic visual reasoning and built-in self-refinement.

And we think that it's going to support both a lot of consumer engagement and creative advertising. And we also hope to unlock new experiences. But I think right now, Muse Image will lower the barrier for people to create high-quality content, similar to the way stories and reels unlocked new forms of creativity.

And we generally that expect AI assisted media tools are going to be able to enable new ways for people to express themselves, to empower SMBs, to create more personalized Ad Creative at a fraction of today's cost and Muse Image began rolling out in July within Advantage+ in the Creative Tool Suite. And so we're excited for that. And there are also new experiences we're building.

We're using it to help customers in their shopping experience. There's a new restyle experience in Meta AI Shopping that allows people to visualize catalog products in their own spaces. So, you could see maybe how a couch would look in your living room, for example.

So, we -- I think the progress here has been very exciting and I think we are still relatively early. I think we expect both Muse Image and Video to evolve towards becoming more capable over time. And I think it will expand the aperture of the opportunities that I described.

Your second question was about infrastructure deals, I guess that we signed over the last few months. I'm not sure if there is a specific -- if there's a specific deal that you're referencing. In general, we have been contracting capacity from a number of third party cloud providers to meet our infrastructure needs.

I think I alluded to this earlier, but the cloud deals are really the fastest way for us to bring '26 and '27 capacity online just given the lead times associated with building out our owned & operated footprint.

So, I think we, obviously, over the long term investing in our own data centers gives us sort of more efficiency and control in terms of what capacity delivery looks like. But cloud capacity gives us the ability to bring capacity online rapidly, given that they have pre-staged capacity available.

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

Barton Crockett: Okay, great. Thanks for taking the question. I guess there was one thing I was in particular interested in and that is the outlook for revenue.

In the next quarter, you're kind of decelerating a couple of percentage points from what you just put up and I'm just wondering if you could elaborate a little bit on what's behind that. Is it mainly tougher comps? Is there something in the macro that you're seeing and then, if I could throw something else in there, I'm just kind of curious on your legal strategy, which you've highlighted in terms of the expenses and everything.

Why -- can you talk a little bit about your -- how you guys see Section 230 kind of playing into your legal defense, particularly as we go through appeals court, if you think that might be more favorable kind of venue for you and maybe is driving some of the decision to go to trial and maybe set some precedent at appellate versus settle?

Susan Li: Yes, thanks for the question. On the first part of your question, which is the Q3 outlook, there are really a range of outcomes captured in the Q3 outlook. Overall, we expect another quarter of strong revenue growth.

We think it reflects healthy and broad-based advertiser demand and that we really think has benefited from our continued investments in ad performance as well as ongoing tailwinds to ad supply, primarily from engagement growth, to a lesser extent, ad load optimizations.

There are a few factors contributing to the decel and constant currency year-over-year growth that's implied in the outlook, so this is the constant currency piece.

First, we are lapping a quarter of accelerated impression growth, which benefited from engagement-related ranking improvements, mainly on Instagram feed and reels, as well as some ad load optimization on Instagram feed and stories.

Second, we're seeing the impact of less personalized ads offering in Europe now that it's fully rolled out, and that may be an additional headwind. And we're also expecting an impact from our continued integrity enforcement efforts.

But overall, we are pleased with the revenue trends we're seeing, even with some of the continued -- even with some of the continued macro uncertainty.

On the second part of your question, I think as I mentioned, I think it is premature to comment on the litigation cases that are underway. We believe that they are based on untested legal theories. We think any liability here is speculative and we plan to defend ourselves vigorously.

Operator: We have time for one more question. And that question comes from the line of Ross Walthall with Cleveland Research. Please go ahead. Ross, your line is open. Your next question will come from the line of Tom Champion with Piper Sandler. Please go ahead.

Tom Champion: Hi. Good evening, Susan. Could you talk about business agent revenue and will we expect to see that recognized in Family of Apps other revenue? And could you just talk about maybe the business model which you expect will be maybe the driver that gets a million businesses that are currently testing the product to pay for it, where we would maybe see the impact in the P&L? Thank you.

Susan Li: Yes. We want to -- we're very excited about this. We want to make paying for Meta Business Agent a seamless experience for the business regardless of size. And in the second half, we plan to begin charging businesses for Meta Business Agent usage through a mix of subscriptions and volume-based token pricing.

So, we're testing Meta Business Agent on Messenger and WhatsApp as part of our Meta One subscription offering. That is a freemium model so businesses can get started and send a limited number of messages for free, and then subscribe to Meta One for higher messaging limits. We'll make that more widely available in coming weeks.

Additionally, earlier this month, we announced a volume-based pricing model for messages sent using Meta Business Agent. And effective August 1st, we're planning to charge on a per token basis for Meta Business Agent messages with one token base charge that will encompass both AI agent processing and message delivery.

So, again, this is still, I would say, in the very, very early stages of its product life, but we are getting ready to scale it. We are getting ready to drive greater awareness of it, introduce more capabilities and improve our models in a way that we hope will make it a more valuable experience going forward and also grow demand for the product.

Tom Champion: Thank you.

Chad Heaton: Thank you for joining us today. And we look forward to speaking with you again.