

Meta Platforms, Inc. (META)
Third Quarter 2026 Results Follow Up Call
October 21st, 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 third 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. We ask that you limit yourself to one question, and this call will be recorded. Thank you very much.

Mr. Rob Bevegni, Meta's Director of Investor Relations, you may begin.

Rob Bevegni: Thank you. Good afternoon, and welcome to the follow-up Q&A call. With me on today's call is 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 as 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 Youssef Squali with Truist. Please go ahead.

Youssef Squali: Great. Thank you so much for taking the questions. Susan, on the Cheeky Pint podcast, Mark recently hosted a talk there, and in there, he talked about how he basically talked about subscription, token usage, and take rate models as ways of monetizing news.

There was no mention of advertising. So is that the right way for us to kind of think about monetization opportunities around news? Because certainly, if there is anybody that could make advertising work around something like this, it would be you guys.

So just wondering where that fits, if at all, in the monetization toolbox. And then can you talk a little bit about the key assumptions underpinning your guidance around revenue growth for Q4? Is it fair to assume that it's still heavier more on the monetization, on the pricing, than on engagement, just like what you guys showed in Q3?

Thank you so much.

Susan Li: Youssef, thank you for the questions. On your first question, we are, I mean, at present, we are, in fact, most focused on scaling Muse, both in terms of growth and adoption and also features. And there's a long list on the roadmap ahead of us to really get the flywheel going with Muse. So that's our first and foremost focus.

And I think Mark outlined the primary ways that we are thinking about news monetization today. I wouldn't rule out advertising. But, again, we are really focused on driving consumer adoption and getting the data flywheel going as our first priority.

Chad Heaton: And on your second question about Q4, I would say that there's a range of outcomes that are captured in our Q4 outlook. Overall, we expect another quarter of strong revenue growth. This reflects healthy advertiser demand broadly, which we believe is largely driven by our continued investments in ad performance, as well as ongoing tailwinds to ad supply, primarily from engagement growth.

So we obviously don't break out the specific contributions, but those are the general contributors.

Operator: Your next question comes from the line of Ron Josey with Citi. Please go ahead.

Ron Josey: Great. Thanks for taking the question, and welcome, Rob, and hey, Chad and Susan. So I want to ask just another one on Muse monetization.

We know the different types of potential monetization, but for transaction-related revenue, do we need to see more Connect partners before we really start turning that on now that the average might be \$200 per, and understood we need to get greater scale? And then I wanted to ask on the capital raise. I think the 8K that came out, Susan, it talks about the potential to raise hundreds of billions of dollars here over time. So just any insights on that comment relative to the Class C shares that were announced. Thank you.

Susan Li: Thanks, Ron. On the first question, I certainly think we need to see more scale in Muse broadly, both in terms of the number of users who use it, as well as the number of connectors/business partners before we get to a place where we are monetizing transactions.

But I think what makes us feel pretty good about our prospects there is the high degree of commercial intent that you see in the way that Muse is being used today, both with shopping being one of the leading activities at Muse is extraordinarily useful for and is -- in the revealed behaviors of the folks who use it is at the top activity.

And also I would say the fact that our -- many of our highest ARPU users in the family of apps are the earliest adopters of Muse. So those, I would say overall,

the commercial intent is high. And we were pretty optimistic that it's going to be pretty straightforward to find various, um, various avenues to monetizing that once things scale.

On your second question, which is about the Class C, we -- specific reference in the information statement about the hundreds of billions of dollars was really a theoretical maximum that was basically was computed to determine what level of equity financing would trigger a loss of founder control. It is not actually a description of our present equity financing needs or plans. And we don't have anything to announce on that front.

Operator: Your next question comes from the line of Ken Gawrelski with Wells Fargo. Please go ahead.

Ken Gawrelski: Thank you. Appreciate it two, if I may, please. First, maybe Susan, could you talk, could you characterize the environment for compute infrastructure right now?

There's been a lot of headlines and a lot of movement, including what we heard out of SpaceX last night. Could you talk about the mix of lease versus own infrastructure over the next several years? And then second question, please.

Maybe, could you detail a little bit, maybe talk about one or two KPIs that you're monitoring for Muse adoption and engagement? We understand it's still early, but anything you could share about how you assess the progress of Muse as an agentic product would be very helpful. Thank you.

Susan Li: Sure. Thanks, Ken. On your first question, overall the landscape for compute, remains pretty constrained and that's certainly something that we are experiencing ourselves.

We expect that our capacity needs will continue to grow until next year and our plans are built around bringing online the capacity that we think we will need, but we also know we've underestimated that in the past. And so we'll get more visibility to the compute needs to support our AI products as we learn more about how they get used. I would say we are focused certainly on building our own infrastructure for a number of reasons, including that we can really optimize it for our own use cases.

But obviously, that is a longer lead time, longer lead time effort. And so we have augmented that with third-party cloud agreements, primarily in service of bringing on capacity in nearer time frames, and I don't think we have specified the -- I don't think we've specified the mix, but I would say, we are facing the same kinds of capacity constraints as everyone else. We are building our AI fungibility -- infrastructure with fungibility in mind. So we can flex our capacity where we think that it will be put to best use.

And we are really trying to think ahead across different timescales. We are looking forward, we're trying to procure land and prepare large scale sites. So we can flex up capacity to meet demand in faster lead times, if needed, in outer

years while limiting how much we have to commit to today. Sorry, commit to that today because we can pace our server orders over time.

And I think we also think as we look forward, we believe our strong balance sheet is going to be a source of strength for us, I think, in a more uncertain environment as the market shakes out in terms of how this all gets funded.

On your second question, I think there are a lot of things we're looking at. Obviously, it's so new, we're really looking at some of the most bread and butter core usage and engagement numbers. We are looking at retention curves, very similar to the way we would have looked at retention curves for any of our traditional family of apps products. And we generally feel good about what we see and what the shapes of those curves look like by cohort.

We are keeping a very close eye on how people use Muse and really using that to help us sort of think about how to make subsequent features and iterations more useful to people. And what are the dimensions on which I think Muse is really, has been really helpful. The proactivity that I mentioned, I think, in my earlier comments on the main call is a good example of that.

And we have a lot more in the pipeline. I think Mark mentioned some of the next things in our roadmap, including rolling out to more countries, to more partners and connectors, launching voice mode, deeper integrations with the rest of the family of apps, deeper integration with ads and an ability to buy with Muse, expanding Muse to cover more business use cases. So, I would say there's a lot more for us to do, even in the very near term.

And even as we are learning about how people use it, how they engage with it, how much time they spend on it and what they use it for.

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

Shweta Khajuria: Okay. Yes. Welcome, Rob. And hi, Susan and Chad. Two questions, please.

One is on the overall market opportunity for consumer agency. You addressed the 500 million businesses and how that relates to business agents. How should we think about the monetization opportunity given the high commercial intent on the consumer agent front long term?

And then the second question I have is when we think about the 10x cost efficiency gains over the next year, how would you rank order the biggest drivers from smaller models to optimizing cost curves to optimizing capacity usage? How should we think about what drives to get you more cost efficient?

Thank you.

Susan Li: Thanks, Shweta.

On your first question, I think I might separate the components of the question. So, we think that the high commercial intent that is already shown up in the way people use Muse today is I think it's both a signal of where people find Muse the most useful and where Muse stands out doing proactive research for you, researching, providing recommendations, researching alternatives, helping you complete a purchase.

And so, while our family of apps has traditionally been strong at helping people for products and brands, this is a place where we are really excited that Muse can serve people who arrive with more explicit intent and allows us, we think, to kind of tap into that advertising market in a new way. And Muse simplifies the journey, we think, from consideration to transaction. An agent can understand what someone is looking for, it can compare relevant options, it can answer questions, it can help them complete the purchase, reducing the friction that exists versus a person moving across multiple websites and apps to research and buy products.

And we think the apps can also -- the family of apps, sorry to be clear, can continue to create moments of inspiration and discovery, and Muse gives us an opportunity to participate more deeply when a person has expressed commercial intent. So, we think the two experiences across family of apps and Muse we hope will be very complementary, and I think we hope that that will support, again, multiple revenue streams over time.

The figure I shared around the 500 million self-identified small businesses, I highlighted that for a couple of reasons. One is we're very excited about Muse for small business, which we launched in late September, and again, we think it's going to make it, we hope, much easier for businesses, again, hundreds of millions of which are already on our platform, manage their business, manage their marketing campaigns, engage with customers, recommend actions that they can take to improve, the performance of ads, to improve customer responses, that sort of thing.

So, we think that there's a very natural platform for small businesses who already use our products and tools to experience Muse. And then where I think the sort of, where I hope the real magic happens, I realize that's not a very CFO word to say, is eventually for consumer and business agents to be able to interact with each other. And again, with 3.6 billion consumers on our platform, with 500 million businesses on our platform, it feels like the most natural place for a agent-to-agent ecosystem to grow and flourish over time.

Chad Heaton: So, on the second part of your question, you asked about the opportunities to get cost efficiencies with Muse. We do see a significant opportunity to improve Muse's unit economics over time. A key part of that is going to be distilling the capabilities of Muse Wonder into smaller, more efficient, and highly capable models that can serve many everyday use cases at a substantially lower cost.

We also expect gains from inference and product efficiencies as well as more effective use of off-peak capacity. And then finally, having specialized models that we can optimize for specific tasks. So, there's kind of a range of ways that

we will sort of attack the cost efficiency problem, but we think there's a lot of opportunity there.

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

Benjamin Black: Great. Thank you for taking the questions. I guess it's a bit of a follow-up though.

If revenue can sustain strong growth next year and compute efficiency improves in some cases more than 10x, should we expect incremental dollars of revenue to require less incremental capital than is required over the last couple of years?

And then the second question is, you also spoke about AI-driven internal efficiencies in the press release. Maybe touch on what you're seeing there. Where exactly are you identifying those efficiencies? Are you leaning more on Muse models versus third-party models? Maybe just shed some more light there.

Thank you.

Susan Li: Yes. On your first question, I think ultimately getting to sustainable unit economics is really going to be a question of how both of those curves, just what the shape of both of them looks like over time.

So today, we are not monetizing Muse outside of subscriptions. That's kind of the currently the only vector. And it will take us time to build out those other monetization engines as we described. Although, again, we think a lot of them will be quite adjacent and straightforward. So, I think this is really a question of basically how both curves map out over the course of the year, how quickly we're able to ramp up Muse monetization and how quickly we're able to bend the curve on compute efficiency.

I think some of the lowest hanging fruit that on the efficiency side that we know are ahead of us, we are the most confident about, but also those won't be the 10x factors. So, I think 10x is still going to require more technological advances over the course of the year.

Chad Heaton: And on your second question about AI-driven internal efficiencies, a few thoughts there. So, first, we are seeing measurable productivity gains in engineering. Our user-facing engineering output is up meaningfully and growing faster versus a year ago.

A total software output per engineer has more than doubled year-on-year, which is demonstrating how significantly AI is increasing our development capacity. AI is also increasing our analytics output and broadening access to data. Employees are producing approximately 10% to 20% more analytics output year-over-year, and more than 70% of measured analytics output is now AI-assisted, which is up from effectively being zero at the beginning of the year.

We see AI agents also helping employees outside of traditional data roles, answer questions and generate insights more independently. AI is expanding into who can actually build the products. Nearly 30% of product managers, designers and data scientists and researchers shift code in the most recent measured period, that compares with about 5% a year ago.

Obviously, their total contribution remains smaller than engineering output, but I think it's an encouraging sign that AI can reduce some functional barriers and shorten the path from ideas to working product. We're seeing adoption of our internal AI tools being broad and sustained. Employees are regularly using our internal AI assistance for information retrieval, drafting and workflow automation.

As our models are becoming more capable, coding has really emerged as an increasingly important use case, helping both engineers and employees and other functions move quickly from just having ideas to actually building products. And we're beginning to realize operational and infrastructure savings. Since the second quarter, our AI-enabled data infrastructure work has generated meaningful savings in internal compute efficiency, while helping us remove more than 50,000 unnecessary data assets.

So, I think this is just an illustration that the opportunity extends beyond individual productivity to more efficient operation of our overall technical infrastructure.

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

Deepak Mathivanan: Great. Thanks for taking the question. Susan, in terms of scaling Muse, can you talk about the opportunities to leverage the almost 4 billion user base of Facebook and Instagram to drive adoption for Muse? The use cases obviously are different, but how do you think about potential integrations that can help scale Muse? And is there a timeline where we can expect a more aggressive push on it?

And then second question, now that Prometheus has come live, can you talk about where you stand in bringing some of the other larger clusters live for training potentially the next set of bigger models? How much of model advancements going forward is dependent on bigger clusters coming live versus some of the more fungible compute that you have procured?

Thank you so much.

Susan Li: Thanks, Deepak. On your first question, today Muse is only available in the U.S. and Canada. And Mark talked about how scaling to a wider range of geographies is going to be dependent on many things, but one of which is being able to do this sustainably from a business perspective, given the cost of serving Muse.

And so, I think we have to be thoughtful about that and make sure -- again, we care more than anything about making Muse and personal super intelligence available to people broadly, and that is part and parcel of our mission, but we have to do it in a way that's thoughtful and sustainable.

Obviously, as we are able to do that, I think our distribution platform will make that -- I think that there will be many very natural touch points to introduce folks around the world to Muse. I think deeper integrations across family of apps, I think using Muse in chat formats, I think there are going to be many ways to bring it to a larger audience. And I think we're excited about when we will be able to do that.

On your second question, which is how do we stand in bringing up other clusters live, I'm not sure that we have said that much about this publicly. We certainly -- we have a number of data center sites at different stages of their life cycle, and some of them are on track to be quite large, although those are, I think, largely still under construction.

But we have many post-training opportunities that we run basically on smaller clusters. And so, while pre-training base models is more dependent on large clusters, there's actually a lot of work we can do with the many other smaller clusters that we have access to. And we are also certainly working on larger data center projects that are in the pipeline.

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

Rob Sanderson: Yes, thank you. Good evening, and afternoon, everybody. Two questions for me, please.

First, a question on sort of go-to-market consideration on Muse. It's on one hand, getting -- there's a big time-to-market advantage, obviously, going -- getting people to use your agent before the competitors' products come is strategically important. But on the other hand, it's a new category, more sophisticated than chatbot. I think it takes some time for people to understand the value prop.

So, really, how do you balance the strategic importance of driving wide distribution as soon as possible versus making sure that you bring people on as they're ready and avoid having a negative first impression? That's one.

And then, second, on the business-as-usual ML, over the last year, we've talked a lot about expanding GEM, Andromeda, Lattice to additional surfaces. Like, is there any way to sort of dimensionalize, like, how much work's still ahead of us to have these models fully deployed?

And we saw a blog post from Engineering at Meta last week, I think, is detailing some very impressive cumulative gains from GEM. Is this -- getting these models fully optimized and deployed, is that really what we're talking about

with confidence on the healthy pipeline for rankings and recommendations?
Thank you.

Susan Li: Thanks, Rob. Your first question, I think, is a very good one. I think we are trying to be very thoughtful about it.

I will say, again, our experience with Muse so far has been just that it's -- the product market fit is quite good even for folks -- I mean, first of all, the agentic experience is somewhat new to a lot of people, but I think even for folks who are maybe less tech-savvy. And I think we have really learned a lot from kind of watching the way that our early adopters use it and what appeals to them about it, again, what they use it for, the degree to which the proactivity, I think, is a really standout feature.

I think we've been hearing a lot of feedback about Muse cases, as we call them, and things like the way in which -- I mean, people are using it in very creative and clever ways to, like, save money in very meaningful ways to negotiate deals.

So, I think it's actually -- it's been both -- I think it's been both a more intuitive experience, I think, than maybe a lot of people had imagined, and we feel very good about where we are with the product and where we are competitively. I think the challenge that you described is something that I think anyone launching an agentic experience will face. And I think that right now we are really leading in the market here, and we have, I think, an unparalleled ability to scale once we start rolling this out more broadly.

On your second question, there are, I mean, I alluded to this on the main call, I think we just have a lot of runway for continued ad improvements. We've made a number of improvements across our ads retrieval, and ranking stacks in the last many quarters. And I think we are continuing to find ways to either enhance the actual performance of the recommendation engine, such that it increases conversions, or to drive efficiencies across the ad delivery system.

And so, I think we feel very good about the opportunities here. I think that the expansion of our LLM-inspired modeling techniques across the ad system, which have allowed us to select better ad candidates using a generative approach rather than scoring every ad individually, that helps us predict not just whether someone will convert, but what the conversion is worth. So that's been the type of work that we've launched with Meta Generative Recommender.

And then, as is the case for us, we launch things in one service, and then we slowly kind of go through the exercise of scaling it to other services as we find that it works. So again, we feel like it's a healthy pipeline of ranking and recommendations gains that the team is working on.

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 curious about the commentary about hundreds of dollars, I think you said, in revenue attached to current U.S. Muse users across the Meta family of apps, and the idea that Muse could be even larger in time in success. I was wondering, first, if those Muse users are different or similar to the overall kind of average U.S.-Canada user base?

And secondarily, if you could give some description of what you think is the road to get to that end state of more revenues from Muse? And importantly, does that assume that there's got to be a transition, meaningful transition from Muse as a facilitator of transactions to Muse agentically completing transactions?

Susan Li: Thanks, Barton. Okay, so what I would say is, just on the sort of first part of your question, so Muse right now skews to our highest value users. I think over 20% of U.S. Muse users today are \$1,000-plus annual ARPU users. So, it gives you a sense of kind of what the early adopters look like, and so they are not totally representative of the overall U.S. user base, which is not representative of the overall worldwide user base. So that kind of gets to the first part of your question, and the second part, I think, is a very good question. I think we don't exactly know the answer to that.

Now, I think there are some analogs with our ads business, where obviously our ads business both facilitates a lot of transactions. Some of them are completed on our platform, but not all of them are completed on our platform. And I think we've obviously been able to grow a large and healthy ads business with -- despite the fact that not all of the transactions happen on our platform.

Ultimately, I think it will be about the value that we are able to drive to businesses and consumers and to businesses in particular relative to the other places that they would look to market and advertise their products.

Operator: We have time for one more question. And that question is going to come from the line of Michael Nathanson with Moffett Nathanson. Please go ahead.

Michael Morton: Hi, good evening. Thank you. This is Michael Morton on for Michael.

Two if I can, I wanted to ask Susan about what you're seeing from consumers with trust for Muse. If there's any data points you could share around Gmail shares or like kind of using the connector there or bank accounts to kind of give some perspective in the population's willingness to share this data with Meta and Muse going forward.

And then I hate to be a dead horse around commerce and partnerships, but longer term, how's Meta thinking about the willingness to allow transactions to occur offsite that might start within a Muse interaction?

Because we've seen some of the larger platforms more open to a kind of lead gen source. But if you look inside of Muse today, if you try to right click, you can't even do the traditional drop down of like search on Google for tech. So, anything around that would be great.

Thank you.

Susan Li:

Yes. On your first question, we understand that security and privacy are critical into the way that people think of and use Muse. And agent is acting on your behalf.

Muse books and buys and sends and manages accounts on someone's behalf. So, mistakes have real world consequences and trust is one of the most important factors for adoption. We have really thought hard about how we built Muse to address those types of concerns.

Each person's Muse runs on its own dedicated isolated computer in the cloud. It's contained so no one else's agent can reach it. Muse's memory, recurring tasks and artifacts all live inside that person's virtual machine.

There is a separate agent referred to as Sentinel that is kept apart from Muse at the system level and it checks actions and connected apps and anything that Muse sends to the internet. So, Muse proposes actions, but only Sentinel can grant permission to perform them. Muse has no visibility into your passwords, logins or payment methods.

So, credentials go into secure storage that Muse can use, but cannot read. And that includes passwords that you type into the browser yourself. Muse checks with the person before sensitive actions, such as sending an email on your behalf or making a purchase.

And it has a complete audit trail of everything it has done and plans to do. So, we have really built in a great deal of security features into Muse because we understand that that is so important to how it gets used.

On connectors, again, we talked about some of the dimensions of this on the main call.

A lot of businesses are eager to meet customers where they are. And that shows up in the number of applications that we've received for connectors. I think at Connect, we announced over 1,500 applications.

Today, that number has more than doubled. Almost 30% of users have connected at least one service to Muse. So again, obviously, it is not the case that every single business has connected to Muse, and there are businesses who clearly have concerns.

But at the same time, we think this will be a natural evolution of the ecosystem. And I think there will be many businesses who are excited to, again, find consumers and their agents where they are.

In terms of willingness to allow transactions off-site, we'll see how that evolves. As I said earlier, what we have found, at least with the ads businesses, the primary -- the most important thing is the ability of a business to measure

the value that we've actually accrued to them, the return on the ad spend that they generate with us, regardless of where the transaction completion is.

And so I think again, demonstrating value is going to be the most -- is going to be the critical axis on which monetizing transactions on Muse will revolve around. And we'll have to figure out exactly what that looks like over time.

Rob Bevegni: Thank you for joining us today. And we look forward to speaking with you soon.