

Meta Platforms, Inc. (META)
Third Quarter 2026 Results Conference Call
October 21st, 2026

Rob Bevegni, Director, Investor Relations

Thank you. Good afternoon and welcome to Meta Platforms' third 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.

This has been a busy and exciting few months. Our family of apps is now used by 3.64 billion people each day, and we're seeing all-time highs in engagement on both Facebook and Instagram driven by strong growth in video. Threads has now overtaken X in monthly actives and continues to grow quickly. WhatsApp is hitting all-time highs in messaging and is becoming an increasingly important platform for agents as well.

In the last few months we have had several major milestones towards our goal of delivering personal superintelligence to everyone. We had one of our most successful launches ever with the Muse personal agent; we shipped Muse Wonder as our first frontier model; and we shipped several new glasses and gadgets with Muse integration, including coming soon to our new audio glasses.

We have been discussing our vision for personal superintelligence for more than a year now. We want to help every person have an extremely capable personal agent that understands your goals and can help you achieve them and improve your life in whatever ways you want -- whether that's improving your relationships, helping you save time, save money, build a business, create new things, or anything else that you can dream of. Muse is the product form of this vision. So now we have the foundational technology in place to deliver this. There is still a lot of work ahead, but our initial position is very strong.

In just a month and a half, more than 10 million people in the US and Canada have set up Muse personal agents. Growth continues to be strong and we have been #1 in the App Store almost every day. Engagement and retention are also off to a strong start. About 25% of people who have set up Muse use it on any given day. For our earliest adopters, the longer that they use Muse,

the better that retention seems to get. To put this in perspective, this is more than twice the level that Threads started at or where we believe that other AI apps started at as well.

The most common use case for Muse is shopping, with more than a quarter of people who are using Muse engaging in this with an average purchase size over \$200. Plaid told us they hadn't seen activity like they're seeing with their connector on any other service before. And in the first week we had to work with Stripe to raise credit limits since the volumes were above what any of us expected.

Beyond shopping, the most common connectors are for personal productivity and connecting Muse to your other services like email and calendar. This shows the importance of the Muse Secure VM that we've built to provide state-of-the-art privacy and security, and we're looking forward to following up soon with the Muse Confidential VM, which will make it so that even Meta cannot see your content.

We have also seen that Muse Wonder improves quality, engagement, and retention further, which suggests that as we continue to improve intelligence, Muse will perform better and better over time.

Now that Muse has clear product-market fit, we're entering a new phase and a more operational phase that is focused on scaling it. The goals of this next phase play to many of Meta's strengths. We need to do three things: scale the product, grow revenue, and optimize compute costs.

We expect to grow Muse to serve billions of people. This is both our mission and our strategy. It's our mission in that we believe in putting the power of personal superintelligence in everyone's hands. We believe that that is the path to ensuring a positive future for everyone with individual empowerment as the source of prosperity and balance of power as the foundation for safety. On a values level, we believe that delivering this capability widely is the best thing that we can be doing. This is also our strategy, because we think that billions of Muses in a network at scale will be able to get more done for people and businesses by working together. So we'd like to create a large scale Muse network as quickly as possible.

From a business perspective, the average revenue per American user for our existing Family of Apps users is several hundreds of dollars annually, and we expect that Muse will reach an even higher level over time due to the amount of commerce that we already see taking place. This will take some time, but we're starting to work on this now. Our business model will have a few components:

For most people, the product should be free and we will build commerce tools so we get paid a fee based on the transactions that we help facilitate. We expect this will be paid for by businesses.

For people who want the most advanced models or to use extra tokens, they subscribe directly and pay.

For business owners who want additional features to use Muse to grow their businesses or for companies that want people to use Muse for work, they will be able to subscribe directly and pay as well.

And for those who want Muse to help improve the personalization and quality of their content and ads across Meta, they will be able to connect to those experiences as well.

Just to clarify, the average revenue that I was talking about before is the Family of Apps revenue for our existing Muse users.

On the compute side, we are building out the capacity to serve more than a billion people over the next few years. In the near term, Muse is very compute-intensive to run. But we expect to improve its cost efficiency by more than 10x in the next year, which will enable it to be profitable once the revenue lines are at scale. If compute efficiency continues improving quickly for another few years, which we think is highly likely, then Muse will be highly profitable over multiple years. In the near term, the strength of our core business enables us to scale the Muse network before then. And we believe that we are the best positioned company to pursue this strategy.

One other note on compute economics is that since training is a large fixed cost, now that we have leading models, we want to serve as much high-margin inference demand as possible. That is why we are building the Meta Enterprise Platform, led by CJ Desai. Muse Wonder and even more advanced models that we are training can serve additional high-margin services beyond the Muse network, including our API, Muse Code, and other enterprise services.

A fundamental feature of our business is that we have many good uses for this compute. It is difficult to forecast the exact growth rates for the Muse network, for our enterprise platform, and for how Muse Wonder and beyond will accelerate our core business. But through building out multiple ways to deploy these advanced models, we increase the chances that one or more of them significantly outperforms our expectations. Beyond serving other businesses, we have not yet found the limit of how effectively we can apply compute to accelerating our own core business. We have implemented an eval with Muse Wonder to improve how effectively it can autonomously improve our recommendations and ads systems. And as our upcoming models get even more advanced, I expect that we're going to see an increasing rate of improvement in these systems.

Alright, let's talk about the models now. Muse Wonder is state-of-the-art across many dimensions that matter for personal superintelligence. This validates the importance of training models ourselves to deliver the best experience. In addition to Muse, we've also seen that engagement per person in Meta AI has more than doubled since April when we started rolling out our new models.

While Muse Wonder has been our most important release, I've been very pleased with our pace of execution here. In the last few months we've shipped Muse Spark 1.2 and 1.3, Muse Code, two open weights models with Muse Glimmer and Muse Spark Open, and also Muse Voice.

Going forward, with the field advancing quickly, we plan to focus on both pushing the frontier of intelligence and distilling those models into much smaller and more efficient product models that will be optimized for personal agents. My view is that we will be able to build superintelligence through scaling, so we are methodically working our way up the ladder there. We are already working on larger models after Muse Wonder, and they're going to be very impressive and capable. We're also working on new smaller models that are coming out soon that will deliver Muse Wonder-level intelligence more efficiently.

Moving on, our work on personal superintelligence is now also the focus of Reality Labs. Muse is coming to Meta Glasses starting next month, and that will be an important step towards this. People already love many features of the glasses already -- music, calling, capturing moments,

hearing assistance, and more -- but over time we expect Muse and having your personal superintelligence with you throughout the day will be the most important part of the experience.

Muse is also central to the business model for the glasses. Again, our mission and strategy are to serve everyone here. We believe in putting the power of this technology in everyone's hands. So we have priced our devices not to make a large profit on the hardware, but to build a large software and superintelligence business on top of them over time. Glasses sales continue to grow rapidly, and now with Muse, we expect that paid subscriptions will increase and that this business will begin improving meaningfully from here.

I also shared at Connect that we are changing our approach to virtual reality to only ship new devices when we have major breakthroughs like we had with the Meta VR Glasses this year -- which of course are the first VR experience ever in a pair of glasses. I continue to believe that the path to delivering the holographic experience of bringing together the physical and digital worlds in everyday glasses will come through VR first, but we will pace this aspect of our investment more to focus on delivering personal superintelligence first.

So with these factors combined, we expect that the P&L of Reality Labs to improve from here, with operating loss coming down to around \$15B next year and trending towards profitability from there.

Before we conclude, I want to spend a moment on the sheer velocity that we're seeing. An important part of our approach has always been to move fast so we can learn quickly and serve people better at scale. But recently things have accelerated. In the last year, we shipped Edits and it has already passed 50 million actives. New apps like Forum and Seller were built by small teams and are scaling quickly to serve power users. Muse went from an idea on a whiteboard to a successful product in just about eight months. And we're about to ship Muse Charm, which will have gone from concept to a hardware device in people's pockets in less than a year.

I'm proud of what we're building. People have all kinds of different opinions about the future, but at Meta we believe that the future is for everyone and that putting power in people's hands is how we get there. I think that this matters. There needs to be a company out there that represents these ideas. But this isn't just about ideas; it's also about execution. So I'm grateful for everyone who is working so hard to deliver all of this to billions of people around the world. And as always, thank you to everyone who is 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.

Q3 Total Family of Apps revenue was \$64.2 billion, up 26%.

Q3 Family of Apps ad revenue was \$63.1 billion, up 26% on both a reported and constant currency basis.

In Q3, the total number of ad impressions served across our services increased 10%. 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 14% year-over-year, driven by ad performance gains and improvements in macro conditions relative to Q3 of last year. This was partially offset by strong impression growth, particularly from lower-monetizing surfaces and regions.

Family of Apps other revenue was \$1.1 billion and grew 59%, driven primarily by WhatsApp paid messaging and subscriptions revenue.

Within our Reality Labs segment, Q3 revenue was \$443 million, down 6% due to lower Quest headset sales, which were partially offset by continued strong growth in AI glasses revenue.

Moving now to our consolidated results.

Q3 total revenue was \$64.6 billion, up 26% on both a reported and constant currency basis.

Q3 total expenses were \$51.6 billion, up 68% compared to last year, and included a \$9.9 billion charge related to the August 2026 multistate attorneys general agreement. Year-over-year growth was primarily driven by legal-related costs, infrastructure costs, and third-party AI token costs.

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

We ended Q3 with over 66,000 employees, down 12% from Q2 following the May 2026 headcount reduction.

Third quarter GAAP operating income was \$13.0 billion, representing a 37% decline year-over-year and a 20% operating margin. Excluding the Q3 legal charges related to the aforementioned legal settlement, our third quarter operating income would have increased 11% year-over-year.

Our tax rate for the quarter was 24%, above our outlook range primarily driven by the non-deductible portion of legal expenses.

Net income was \$9.9 billion or \$3.83 per share.

Capital expenditures, including principal payments on finance leases, were \$34.9 billion, driven by investments in servers, data centers, and network infrastructure.

Free cash flow was negative \$2.1 billion. We ended the quarter with \$82.4 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 per user reached an all-time high and grew double-digits year-over-year in Q3, with strength across both Reels and Feed.

On Facebook, Reels time spent reached all-time highs both globally and in the US, with both growing double-digits year-over-year.

Threads also continued to scale, with time spent per user growing over 20% year-over-year.

Engagement gains across our services continue to be driven by improvements we're making to the relevance, personalization and timeliness of the content we recommend. Investments into our recommendations engine across retrieval and ranking have continued to yield positive results on Instagram and Facebook.

For instance, on Instagram Reels, our relevance launches in Q3 increased Reels time spent by over 4%, driven in part by our upgraded model architecture that ranks sequences of Reels at a time instead of evaluating each Reel independently, as well as improvements we made to add richer watch time signals.

On Facebook Feed and Reels, we completed the global launch of a new generative retrieval architecture which, inspired by LLMs, labels pieces of content with semantic representations and generates potential recommendations, helping our systems surface relevant new posts before they've accumulated significant engagement. This new architecture is now rolled out to roughly one-quarter of unconnected Feed views and half of Reels views and increased Reels watch time by 1%.

Our recommendations work is also helping us deliver fresher and more personalized content. Nearly 45% of recommended Instagram Reels are now less than one day old, up more than 11 percentage points from Q2. On Facebook Reels, we doubled the amount of history our ranking model can consider, driving a 1% increase in Reels watch time in the US.

LLMs are also helping us understand content more deeply. We continue to process public Instagram posts each day through LLM-based content-understanding systems. These signals have improved the engagement with Instagram search results by 1% through the end of this quarter.

This quarter, we also began moving our content-understanding models onto our first-party Muse model family. In offline evaluations, these models are outperforming our prior production models and are reducing per-unit content processing costs by more than 80% across Instagram, Threads, Facebook, and ads. Beyond improving quality, owning more of this stack gives us greater flexibility to develop models around the signals that matter the most for personalization.

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

We see significant opportunity to continue improving performance for the businesses who use our services.

Within our ads systems, we are continuing to drive performance gains as we deploy more predictive models that understand users more deeply and deliver ad experiences better aligned with their unique preferences, while also making our infrastructure more efficient.

Last quarter, we introduced Meta Generative Recommender, which uses LLMs to reason across ad content and user preferences to predict the best ad for each person. In Q3, we launched this

approach to generative retrieval on Facebook Feed, driving a nearly 2% increase in conversions across targeted segments.

In addition, we expanded our GEM model for ads ranking to additional surfaces, driving a 1.7% increase in conversions for offsite conversion campaigns and a 9.6% increase in app ad conversions.

And we continue to scale our retrieval systems efficiently. In Q3, improvements to Andromeda such as scaling model complexity and throughput increased conversions by 1.4% on Facebook.

We're also leveraging AI to make it easier for businesses to manage campaigns, develop creative and act on performance opportunities.

Meta AI business assistant is becoming a more proactive and capable partner for advertisers. We estimate that advertisers who adopted its recommendations drove almost 5% more conversions at the same cost per conversion and increased their advertising investment by 9% compared with similar non-adopters. We are rolling out additional capabilities, including creative insights, stronger business context and more agentic actions. These features will allow marketers to plan, create and optimize campaigns, as well as schedule tasks and alerts, through a conversational experience with approvals and previews that keep them in control.

I'll share one example of how Meta AI business assistant is empowering our advertisers. Simply Organic Maple, a family-run maple syrup producer in Vermont, used Meta AI business assistant to improve its campaign messaging. After the business assistant recommended testing longer, more descriptive advertising copy, the business saw a 57% increase in purchases and a 36% reduction in cost per purchase. Its strongest revised ad delivered a 3x return on ad spend, a 52% improvement from the original.

Adoption of our generative AI ad creative tools also continues to scale. In Q3, nearly 10 million advertisers used at least one of these tools. Advertiser adoption of our video generation features grew 25% quarter-over-quarter, driven in part by ready-made videos that reduce generation time to just a few seconds from a single image upload.

We're also expanding the opportunity at the intersection of creators and advertising.

Partnership Ads, which allow advertisers to run ads with creators, brands and other businesses, continued to scale rapidly in Q3. Our partnership ads product grew over 97% year-over-year and surpassed \$15 billion in annual revenue run-rate.

In Q3, we also completed the global rollout of Creator Marketing Hub to all advertisers, bringing creator discovery, content organization and partnership ad activation into one place, with more than 10 million creators now discoverable. We are integrating Meta AI Business Assistant into the Hub so advertisers can conversationally search for creators, identify content to use in partnership ads and get guidance across their creator marketing workflows.

Outside of ads, we are also building out additional ways to monetize engagement across our ecosystem.

In September, we announced the global expansion of Meta One, our unified subscription service for consumers, creators and businesses. Meta One is rolling out gradually with more than 50

features across Instagram, Facebook, and WhatsApp, including Meta Business Agent and Meta AI, and has surpassed 15 million subscriptions and trials to date. The core experience across our apps and Meta AI will remain free, while subscriptions offer more specialized capabilities and expanded AI usage.

Beyond monetizing engagement across our existing services, we are building new experiences and revenue streams around consumer agents.

Mark highlighted the impressive early momentum with Muse. Users have found many reasons to use Muse with the top 3 categories being shopping, personal productivity and creativity. We are encouraged by the breadth of use cases, with shopping in particular being used by over 25% of Muse weekly actives.

While our near-term priority is to grow adoption and expand Muse's reach, we see multiple paths to monetization, beginning with subscriptions.

We also see a significant opportunity to extend Muse to small businesses. Based on a recent user survey, we estimate that at least 500 million self-identified small businesses use Meta's family of apps to reach customers and generate revenue. Muse could help these businesses acquire customers and operate more efficiently, while helping people complete transactions. Over time, consumer and business agents could connect and transact directly, creating opportunities around paid usage, commerce, and business services.

Alongside growth, we are focused on building sustainable unit economics. Muse's compute intensity means that expanding its user base will require significant investment, and we expect total serving costs to rise. At the same time, we expect to improve compute cost efficiency by more than tenfold over the next year, driven by better utilization of off-peak capacity, more efficient models and inference optimization.

As revenue scales, these efficiency gains give us confidence in our ability to build a highly profitable business over time.

Moving now to our financial outlook.

We expect fourth quarter 2026 total revenue to be in the range of \$73-76 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 our expense outlook to incorporate the \$9.9 billion charge related to legal proceedings recognized in Q3. We now expect full year 2026 total expenses to be in the range of \$175-179 billion.

Excluding the Q3 legal charge, we continue to expect 2026 operating income to be above 2025 operating income.

At a segment level, we now expect Reality Labs losses to decline compared to 2025 to a range of \$18-19 billion.

We anticipate 2026 capital expenditures, including principal payments on finance leases, to be in the range of \$135-145 billion, narrowed from our prior outlook of \$130-145 billion.

Absent any changes to our tax landscape, we expect our fourth quarter 2026 tax rate to be between 15-17%.

Additionally, 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 in the coming months in the United States.

Looking beyond 2026, we are still developing our budget and capacity plans for next year, but we expect to continue investing aggressively to support both our core business and emerging AI initiatives. Our current capacity and what is scheduled to come online gives us the flexibility to allocate resources across model training, core business growth, and the development and support of new products.

Our AI investments are already delivering substantial incremental advertising revenue through stronger engagement, better recommendations, and more efficient monetization. Building on this progress, our MSL investments are creating the technology and infrastructure needed to unlock significant new revenue opportunities and strengthen capabilities across the business. While these investments will take longer to generate returns than those in our core business, the breadth of monetization opportunities and internal efficiencies gives us confidence in their long-term potential.

Looking ahead, although we are not providing a full-year 2027 revenue outlook, the benefits our AI investments are delivering to the core business, combined with emerging opportunities from model advances and product launches, position us to sustain strong overall revenue growth throughout 2027.

Separately, we remain focused on improving Reality Labs profitability. As Mark mentioned, we expect losses to narrow next year to around \$15 billion as we reallocate investment toward other priority areas, primarily AI.

Finally, today we announced plans to create a new class of non-voting capital stock, which will be distributed to all existing stockholders through a stock dividend. For each existing share held, stockholders will receive one share of the new non-voting stock, effectively resulting in a two-for-one stock split. We expect trading to begin in late January 2027.

In closing, our business momentum continued in Q3, with strong execution across our core ads and engagement initiatives. And we're also excited by our progress towards bringing personal superintelligence to everyone.

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

Operator: And your first question comes from the line of Brian Nowak with Morgan Stanley. Please go ahead.

Brian Nowak: Thanks for taking my questions. I have two, one on Muse, one on enterprise. Maybe, first, congrats on the success of Muse.

But as you analyze the user behavior and the product, can you talk with us a little bit about any of the constraints to the slope of forward user growth that you're most focused on addressing, whether it's capabilities, model efficiency, capacity, et cetera, just so we understand what are the limiting factors to forward user growth that you want to address? And the second one is on enterprise. Mark, will you just sort of walk us through your vision for the Meta enterprise product offering and the key value drivers you hope to provide for the enterprises adopting the service?

Mark Zuckerberg: Yes. I think I can talk about both of those, and then Susan can jump in with anything else. All right.

So right now for Muse, we are -- and one of the things I pointed out is we have more than 10 million people in the U.S. and Canada who have set it up, but I think one of the biggest constraints so far is that because this is a compute-intensive service to run, we have not launched it globally yet. That's partially because of wanting to optimize it for each place, and there's some work around connectors for that. But it's partially that we want to make sure that we do this efficiently.

So I mean there's pros and cons there. The pro, of course is that, that just makes it even more impressive that you have 10 million people using it and signing up to use it in the first six weeks because those are all kind of Americans and Canadians in a small number of markets.

So if you actually did roll it out more broadly, then you would expect that number to be significantly larger than what we're seeing today.

But then on the flip side, we actually have to go do the work to bring the efficiency in line to make that make sense.

And so that's a big part of it that we need to do. We -- the compute efficiency work is going to be important on this also because we want to make sure that we've built enough compute to be able to serve what we think is going to be 1 billion plus people if we can achieve certain compute efficiency targets, but we obviously have a lot of work to do there as well. And that will be one of the limiting factors over the next year.

We want to build out the business around it. So I talked about that. And I think doing that in tandem with bringing down the compute -- or increasing the compute efficiency is going to be important for getting this to work from a business perspective which I'm optimistic about, but I think that that's one of the things that we need to do.

And then there's just continuing to increase the value in the product which I think we will largely do by continuing to develop new models. That's why I mentioned that we launched Muse first on Muse Spark, and that's what is currently the default across the free tier. And we rolled out Muse Wonder, and people can access that in the paid tier today. And we know that it performs a

lot better, so we're working on getting a more efficient version of Muse Wonder that we'll be able to roll out more broadly.

But you can imagine just as we iterate forward in two, three models into the future, when you have significantly higher intelligence and capability to solve many more problems for people and the ability to deliver that at a much more affordable price point, I think that that's going to just unlock a lot more value to people. So we're very excited about that.

Transitioning to talking about the enterprise piece because I think these two pieces, there is some connection here in that the main enterprise offering in the near term is going to be the model API and Muse Code. And we have business agents that we've been scaling for a while, and each of those three things is growing quite quickly. So we're excited about that.

Over time we would like to make it so that we have a whole kind of business network of Muse agents as well.

There's the business agent which we've talked about on previous earnings calls for businesses to help out with sales and customer support.

But what we see is a lot of people are actually using Muse today to help run especially small businesses. And we think that there's both an opportunity to help create an agent that people use to really enable them to automate and do a lot of the parts of spinning up a small business that are just kind of difficult and time consuming but aren't the core thing that a person is trying to do.

So we think that's a big deal. And then making a version of Muse that people can bring into different corporate environments and work there, we think is going to be a big deal.

And then beyond that, building what I refer to as the Muse network which is making it so that Muse is not just a kind of single player game where you have your Muse but making it so that your Muses can work with your friends and your families. And if you want to buy something, it's not just navigating to that business' website. It's talking to their agent. I think that getting that all to work well together is going to unlock tremendous value at scale.

So I'm very excited about all of that. I think the enterprise use is going to be a big part of it, but it's also going to kind of have this very valuable boost in the utility for consumers, too. So I'm very excited about both of those.

Susan Li:

I would add on to Mark's comments on enterprise. AI is creating huge amounts of disruption in enterprise today. And we don't -- from what we can see, we really don't believe that traditional enterprise companies have this locked down. And in fact, some of the leading companies in this space today didn't exist a few years ago. Those companies made large investments in model development, and we've done the same, and we have increasingly competitive models to bring to bear for the benefit of all of our users including small businesses and including enterprises.

And already today we have 500 million businesses using our products. So we really believe that AI is going to reshape how businesses operate.

We expect that businesses of all sizes are going to use AI agents across software development and productivity, sales, marketing, consumer -- customer service, et cetera.

We have differentiated assets to serve this market, bringing together models, infrastructure, global distribution and long-standing relationships with businesses. And it's our hope that this will be able to be a new diversified recurring revenue stream for us that will be able to support multiple business models including subscriptions and usage-based API and messaging and other revenue lines.

So we expect that this will be an opportunity that will develop over multiple years, and we intend to scale our investment behind demonstrated customer demand and product market fit in this area.

Operator: Your next question comes from the line of Eric Sheridan with Goldman Sachs. Please go ahead.

Eric Sheridan: Thanks so much for taking the questions. Maybe just the first one, building on Brian's question. When you think about the friction points around agentic commerce, how do you guys as a company think about solving for friction that's rising where there are aggregators or suppliers that don't want agents to interact with their platforms or their services and how you think about lowering that friction or creating the right balance between the existing services that consumers access today relative to how this might all be changed or altered by agentic behavior over the medium to long term?

And the second one would be, with the announcement of the new class of shares tonight with earnings, any updated thoughts broadly about how you think about sources of capital and optimizing your capital structure to make sure you invest the proper amounts to capitalize on the opportunities that AI are presenting in the next couple of years? Thanks so much.

Mark Zuckerberg: I can answer the first one. So the question -- sorry, what was the first question there?

Susan Li: Friction points around agentic commerce.

Mark Zuckerberg: Got it. All right. Yes. So the -- right now what we're seeing is that, overwhelmingly, businesses are very excited about enabling this because they recognize that it's going to be a new opportunity for them not only to increase the sales and the quality of the experience for existing customers but also potentially to bring in new customers.

So now I mean across the whole business ecosystem, different people will kind of adopt this at different paces and how our strategy evolves on that to develop the ecosystem will depend on how that goes.

But overall, I would say we've been very positively surprised by how quickly, I think a huge amount of businesses have wanted to not only figure out how to make it so that Muse can use their services more effectively but have reached out to us to integrate around it.

So I view that as a very good sign. And obviously not everyone wants to be an early adopter. That's fine. We'll work with folks on the terms that makes sense for them.

But I do think that, over time this is going to be quite valuable for everyone to make this work well and I think we're on a good trajectory on that so far.

Susan Li:

Yes. I would add to that, and I will also answer your second question and just say I think agents are going to be an important part of how reality evolves, and I think everyone is going to have to adapt to what that means, the new age of agentic commerce.

We obviously have had a lot of businesses we're very excited to partner with us. There are also businesses that are understandably focused on protecting customer relationships. And we expect the market, I think ultimately to evolve away from agents attempting to navigate every digital experience like a human.

And I think over the long run, this will move towards more permissioned, structured interactions, and that will include authenticated APIs, approved agent partnerships or direct communications between the consumer's personal agent and a business' own agent.

And we're helping establish the standards needed to support that model. Earlier this month, we previewed the Personal Agent Protocol, an open standard across Muse and Meta business agents that we are developing with industry partners like Sierra, Walmart, Stripe, Shopify, et cetera. The protocol is intended to define how a consumer's personal agent can interact securely and reliably with the business' agent and preserve appropriate controls for both the business and customers.

So again I think adoption will likely be uneven and will take time and in the near term, we expect any experience -- many experiences will remain AI assisted rather than fully autonomous.

But as trusted integrations develop and agents demonstrate that they can generate incremental value and demand and operate safely, we expect that more and more businesses will be excited to participate.

Your second question was about the announcement of the new class of shares.

So I should say we created this new class of shares with the purpose of creating a publicly traded non-voting stock usable for capital raising, employee compensation, acquisitions, other general corporate purposes while preserving the voting power of existing Class A and B shareholders.

We believe this share class will maintain our founder-led structure. It will expand our funding tools and give us flexibility to invest behind the unprecedented momentum we're seeing across every part of our business driven by AI.

We don't have any specific financing plans to announce at this time but we thought it was prudent to establish the appropriate capital structure framework to ensure that we have access to a diverse pool of funding sources.

Operator: Your next question comes from the line of Doug Anmuth with JPMorgan. Please go ahead.

Douglas Anmuth: Great, thanks for taking the questions. I have two.

First, I just wanted to follow up on the -- your last comment, Susan, just around as Muse evolves and more businesses adopt agents. Just wanted to get your sense of how things kind of change and shift as you go to more permission-based and agent-to-agent transactions. And like what does that really unlock over time?

And then secondly just on kind of expenses and CapEx, thinking about '27. I know you're not really there yet. But if I go back to some previous comments, you've talked about maximizing compute in '26 and '27 and then having more flexibility in '28. A year ago, you talked about kind of notably larger dollar increase in '26 than '25, and then you've also used kind of growing operating income on an annual basis as well.

So I'm just curious if there's any of those that you can kind of tie or to help us think about '27. Thanks.

Susan Li: Thanks, Doug.

So on your first question about how the agentic ecosystem evolves, I wish I could sort of predict that more perfectly. And I think it will take some time. And again I think it will be a little uneven.

But I definitely think that when I use Muse today it's clear that there are a lot of parts of the experience that would be even better if Muse was able to interact more seamlessly with business agents on the other side.

For example, when -- if Muse has to call a business on my behalf, I think it would be easier for the business if that call was handled by business agents that were representing it. And I think also one of the things that is really distinctive about Muse and one of the things that makes the product

experience so great is the degree of proactivity that it exhibits on my behalf, right?

I can give it a task and Muse is going to go execute against that task for me even as it encounters some obstacles or can't quite find the right thing and needs to chart a new path to helping me look for the thing I want.

And so I think the -- that sort of dimension of proactivity exhibited by either consumer agents or business agents, I think are going to enable more transactions to happen or more activity to happen broadly than happens today where -- when things are -- when more actions are taken, primarily -- when like the latent demand expresses itself versus things where I can have Muse searching for things for me on an ongoing basis, giving me regular updates on things I'm looking for, sending me 10 new books a month.

And agents should be doing similar things for businesses, and I think that creates sort of opportunities, again for a new more proactive sort of framework for what agent-to-agent interactions look like over time.

On your second question which is about 2027, we are quite early in developing our budget and capacity plans for next year.

So we are, I would say at the point where we are both investing -- sorry, evaluating a lot of investment proposals that have attractive return profiles over different time horizons, and we want to figure out how we can fund as many of those as possible within the construct of our overall budget and financial goals.

But it's very dynamic and it's still pretty early. And so we don't have anything more specific or quantitative to share until we get to the Q4 call which is typically when we share guidance for the next year.

Operator: Your next question comes from the line of Justin Post with Bank of America. Please go ahead.

Justin Post: Great. Thanks. One for Mark. One for Susan. Mark, you've got this powerful new model on your hands, and it seems like there's a lot you could do internally besides externally. But specifically to the ad business, you rearchitected your systems around AI.

How much room is there for improvement as you deploy that new model? How are you thinking about that?

And then for Susan, you mentioned strong revenue growth next year. Just wondering if that's advertising. Or does that incorporate some incremental AI revenues hitting next year? Thank you.

Susan Li: I'm happy to take both of those. Great.

On your first question, I think our ads business has been benefiting for now many years from our ongoing investments in both kind of what we would have called traditional AI machine learning models and then, of course from the sort of more recent generation of models.

I would say that there are three themes there that I think are worth outlining. The first is just the kind of traditional roadmap of performance improvements that we have in the pipeline based -- that is what I would describe more as business-as-usual machine learning work.

That continues to be a robust and healthy pipeline. And as we are looking into 2027, again we're not providing an outlook, but I would say that our investments and product improvements across both engagement and ads initiatives give us confidence that we have meaningful runway to grow our core business throughout 2027. That also gets to your second question also.

But I should say basically, I think the ranking and recommendations pipeline is healthy. The second theme I would say is that we are excited about some of the new research bets that we have as we move into sort of a more LLM-oriented AI infrastructure to help us more deeply reason about people's preferences and what they're interested in both across organic content and recommendations, and we expect that we will be able to unlock benefits to our ranking and recommendations engine from that research, too.

And then the third thing is that we expect that generative AI tools which have already very valuable for advertisers will continue to grow in terms of the features that we make available to advertisers primarily through the Advantage+ suite but features that make it easier for advertisers to quickly and seamlessly generate, whether it is creative or take an image and generate video or generate sound, all, I think in service of making the ads increasingly personalized for the person that they're being served to.

So I would say we feel like there is meaningful runway across all three of those things.

Your second question was about revenue growth next year. So as I mentioned, I think we're both feeling good about the budget investments that we expect we will make in both core ads and engagement initiatives next year.

We're also building several new AI businesses, and we expect that those can begin contributing incremental revenue in 2027 but more importantly, can scale meaningfully over time.

And those will include monetizing our AI products through subscriptions, transactions, advertising, developing enterprise offerings, as we talked about in one of the earlier questions.

So I would say while many of these opportunities remain early and our near-term focus is really on building high-quality products and achieving product market fit and scaled product market fit, I think our track record of reaching

billions of people and building businesses at scale makes us feel good that AI is going to become a meaningful source of revenue diversification over time.

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

Mark Shmulik: Yes, thanks for taking my questions. Mark, coming out of Connect, I kind of felt like we were ushering in the next era of the company and the business reorienting around kind of Muse at the center of everything.

Is that a fair assessment of kind of how things are going internally? And if so, can you just talk a little bit more about how some of the business mandates and outcomes may have changed? And then, Susan, it was interesting to kind of see the launch cadence where enterprise announcements kind of preceded some of the small business ones. And I appreciate your comments on some of what's going on there around the SMB agent opportunity.

But can you just share more details on kind of the work required to serve and scale the offerings to SMBs?

Mark Zuckerberg: Yes. I mean I can talk briefly about the first one.

I mean look, it was more than a year ago that I think I wrote the letter about personal superintelligence as the North Star for the work that we're doing here. And it basically -- it fits with this line of work that we've done for more than 20 years at the company that is focused on empowering individuals, giving people a voice, giving people the tools to connect in the ways that they want. And when you extend that to the work around AI now we think that the most important thing that we can be building for this next generation of technology is giving -- is basically putting the power of this technology in everyone's hands which is the personal superintelligence vision.

So we've been working on this for a while, and this is when we did the reboot around MSL and building the lab. A lot of it was this idea that we need to build models specifically to be good as personal agents, that we'd be able to build the best experience by owning and building and designing ourselves the end-to-end technology. That includes the models and then the agents. And this has been -- this has kind of been in the works for a while.

I mean the vision around the glasses has also been this. I mean the -- we've wanted to build the next generation of hardware platforms so that we can deliver the best end-to-end experiences to people. And I think that that's starting to come together as well and I think will even more over the coming years.

I mean these are long-term bets. I feel like this is what we're here to do, is make long-term technology bets and design the best experiences end to end that we can for billions of people around the world. And I think as long as we continue to do that, then over time the business takes care of itself as long as we

continue executing well which I'm really proud of the team for how we're doing on that.

Susan Li: Thanks, Mark. Your second question was around Muse for small business.

So we launched that which you can think of as an expansion for Muse in late September and rather than a separate product, it's kind of a new -- a set of new skills and connectors inside the existing Muse app aimed at helping SMB owners run operations, find new customers.

It connects to a business', for example, Instagram account, analytics or Facebook pages or Meta ad accounts in a few clicks, along with dozens of third-party connectors that businesses already run on like Shopify, QuickBooks, Canva, et cetera. It already understands what businesses sell. Its brand voice, what customers are asking for.

And Muse can take actions including analyze ad and content performance, draft campaigns, customer responses, flag issues, and we are actively listening to our SMB users to understand what their top use cases are.

We are seeing them use it for ads optimization, for direct message organizing, use it to expand their sales footprint to keep their website current based on things like inventory and new products to deal with employee load balancing.

So we're pretty closely following how businesses use Muse to make sure that we are being thoughtful about new features that we release in Muse to make it maximally useful and productive for SMBs.

I'd also say we have unique touch points to, again I mentioned in my opening remarks, 500 million self-identified small businesses. And we think that many of the investments that we make in consumer products are going to carry over to help power those businesses also.

We have multiple excellent channels to surface those opportunities to small businesses globally.

So this is a place where we feel like we have an existing extremely strong platform, and we're very excited about it. And we also believe the transition to AI has broadly unlocked enterprise as a new opportunity for us as well.

Operator: Your next question comes from the line of Ross Sandler with Barclays. Please go ahead.

Ross Sandler: Great. Mark, you touched on this earlier, but could you elaborate a little bit on maybe some high-level metrics around the per-user cost structure of Muse today and then how you kind of see that cost curve coming down as you go forward? And what are some of the things you're doing to bring that cost curve down? Thank you.

Susan Li: I'm happy to answer that. And Mark, you should add anything that you feel like I'm missing.

We are -- we don't have sort of specific numbers to share around the cost structure of Muse today. But of course it is, as Mark said, pretty compute expensive, and making sure that we are focused on compute efficiency is going to be one of the key vectors towards being able to both scale Muse and grow it into the business that we want.

I think there are a couple of dimensions of achieving that efficiency.

I think some of the more nearer-term pieces include things like leveraging off-peak or spare computing capacity to run some sets of jobs, using smaller models. There's a lot of work to be done across different inference efficiency initiatives, like all in service effectively of higher GPU utilization and other product efficiency work we can do also. And I think over the longer term, I think this is a place where we expect that we're going to be able to, over time as LLMs scale, get more tasks done with fewer tokens and also rely on AI models themselves to help us improve the efficiency of the software that we use to run the model.

So I think it's a place where we both feel like there is a lot more work on our efficiency roadmap that is going to help us bend the curve here. And we also expect that the sort of arc of how the sort of AI-driven technological advances are going to apply here will also help us achieve larger savings over time.

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

Mark Mahaney: Okay, thanks. Two questions. There's a bit of a playbook or there has been at Meta over the years when you surface new surfaces and you build up the user growth, engagement. Then you layer in monetization. Do you think it's sort of a similar playbook here? Or does it have to be a little bit faster with Muse just given that the costs associated with it are materially greater arguably than launching those surfaces in the past?

And then just a clarification on that commerce data point that you provided, about 25% of users and the \$200 or greater ASP. Are you including in commerce travel purchases? Or is that just a retail data point?

Mark Zuckerberg: My guess I think that we are going to work on monetization somewhat sooner on this than our normal playbook for a couple of reasons.

One is just the size of the numbers that we're talking about in terms of both the costs and the opportunity, I think make it kind of important to make progress on that in order to be able to scale it. The other is just that such a large part of what people are doing is commerce. And to your question, that does include travel, too, I think we might consider travel a different case, but travel is actually one of the other top use cases for Muse, is basically people planning

that out and planning different kinds of things that they would get than they might get through our Family of Apps.

So I think that this is going to be an important thing for us to work on to deliver the consumer value that people want, and in doing so, I think we will be able to build a model where we get a percent of the value that we are helping to create for people.

So I think that that's going to be -- it's going to be very important to kind of get the math on all of this to work out so that way we can continue the compute build-out that we want to build in order to get this to reach everyone at scale.

And in addition to everything that Susan said in response to the last answer, I would also just add that like we just see this trend that year-over-year, I just expect that we will be able to deliver the same intelligence that we have now in Muse Wonder or, frankly, any of the current frontier models at a fraction of the cost in a smaller model a year from today.

So to the extent that that represents some kind of baseline level that people seem to -- that kind of is sufficient for product market fit for Muse, I think that we will just continue through continued model development, make progress on that efficiency. The thing that is -- and of course less clear is that I don't think in a year from now most people are going to want to use the current model level of intelligence because that's also going to make progress.

So I think we're going to need to work on all of these things, both the revenue pieces of it and scaling the different aspects of that as well as progressing on the model development and all the efficiency in order to be able to both fund and support the compute build-out and reach billions of people which, again we want to get to that as soon as possible.

Part of the reason why we're not just charging everyone upfront to use this is we want to establish this Muse network, and we think that it's going to be more valuable for people at scale when there are a lot of people and businesses using Muse so that way they can collaborate and work together.

I think that that's really a very important state to get to and one that we're quite excited about. And again I don't think that there are many companies in the world that are well situated to be able to do that in the way that we are. I think we are very well situated to do that, and we plan to push on that quite aggressively.

Operator: And that does conclude our question and answer session. Mr. Bevegni, I would like to turn the call back over to you.

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

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