

**Meta Platforms, Inc. (META)  
Third Quarter 2024 Results Conference Call  
October 30<sup>th</sup>, 2024**

**Kenneth Dorell, Director, Investor Relations**

Thank you. Good afternoon and welcome to Meta Platforms third quarter 2024 earnings conference call. Joining me today to discuss our results are Mark Zuckerberg, CEO and Susan Li, CFO.

Before we get started, I would like to take this opportunity to remind you that our remarks today will include forward-looking statements. Actual results may differ materially from those contemplated by these forward-looking statements.

Factors that could cause these results to differ materially are set forth in today's earnings press release, and in our quarterly report on form 10-Q filed with the SEC. Any forward-looking statements that we make on this call are based on assumptions as of today and we undertake no obligation to update these statements as a result of new information or future events.

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.fb.com](http://investor.fb.com).

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

**Mark Zuckerberg, CEO**

Thanks, Ken. This was a good quarter with strong product and business momentum, and with parts of our long-term vision around AI and the future of computing coming into sharper focus. We estimate that there are now more than 3.2 billion people using at least one of our apps each day -- and we're seeing rapid adoption of Meta AI and Llama, which is quickly becoming a standard across the industry.

So let's start with some highlights from the apps. For WhatsApp, the US remains one of our fastest growing countries. And we just passed a milestone of 2 billion calls made globally every day. On Facebook, we continue to see positive trends with young adults, especially in the US. On Instagram, global growth remains strong. We also launched Teen Accounts this quarter, which add built-in protections that limit who teens are messaging and what content they can see. On Threads, the community now has almost 275 million monthly actives. It has been growing more than 1 million sign-ups per day. Engagement is growing too. So we continue to be on track towards this becoming our next major social app.

We are making a lot of progress with our AI efforts too, and we're seeing AI have a positive impact on nearly all aspects of our work -- from our core business engagement and monetization to our long-term roadmaps for new services and computing platforms. And I think that this partially comes from having a vision and roadmap that is aligned with the direction that technology is heading, but even more importantly from our teams doing some really excellent work on execution on so many fronts.

Meta AI now has more than 500 monthly actives. Improvements to our AI-driven feed and video recommendations have led to an 8% increase in time spent on Facebook and a 6% increase on Instagram this year alone. More than a million advertisers used our GenAI tools to create more than 15 million ads in the last month, and we estimate that businesses using Image Generation are seeing a 7% increase in conversions -- and we believe that there is a lot more upside here.

We're also seeing great momentum with Llama. Llama token usage has grown exponentially this year, and the more widely that Llama gets adopted and becomes the industry standard, the more that the improvements to its quality and efficiency will flow back to all of our products. This quarter, we released Llama 3.2, including the leading small models that run on-device and open source multi-modal models. We're working with enterprises to make it easier to use, and now we're also working with the public sector to adopt Llama across the US government.

The Llama 3 models have been something of an inflection point in the industry, but I'm even more excited about Llama 4, which is now well into its development. We're training the Llama 4 models on a cluster that is bigger than 100k H100s or bigger than anything that I've seen reported for what others are doing. I expect that the smaller Llama 4 models will be ready first, and they'll be ready, we expect sometime early next year, and I think that they're going to be a big deal on several fronts -- new modalities, capabilities, stronger reasoning, and much faster. It seems pretty clear to me that open source will be the most cost-effective, customizable, trustworthy, performant, and easiest-to-use option that is available to developers, and I'm proud that Llama is leading the way on this.

Now it's the time of the year at Meta when we plan our budget for the next year. That's still in progress, but I wanted to share a few things that have stood out to me as we've gone through this process so far. First, it's clear that there are a lot of new opportunities to use new AI advances to accelerate our core business that should have strong ROI over the next few years, so I think we should invest more there. Second, our AI investments continue to require serious infrastructure, and I expect to continue investing significantly there too. We haven't decided on a final budget yet, but those are some of the directional trends I'm seeing.

Moving on, this quarter we also had several milestones around Reality Labs and the integration of AI and wearables.

Ray-Ban Meta glasses are the prime example here. They're great-looking glasses that let you take photos and videos, listen to music and take calls. But what makes them really special is the Meta AI integration. With our new updates it'll be able to not only answer your questions throughout the day, but also help you remember things, give you suggestions as you're doing things using real-time multimodal AI, and even translate other languages right in your ear for you. I continue to think that glasses are the ideal form factor for AI because you can let your AI see what you see, hear what you hear, and talk to you.

Demand for the glasses continues to be very strong. The new clear edition that we released at Connect sold out almost immediately and has been trading online for over a thousand dollars. We've deepened our partnership with EssilorLuxottica to build future generations of smart eyewear that deliver both cutting-edge technology and style.

At Connect we also showed Orion, our first full holographic AR glasses. We've been working on this one for about a decade, and it gives you a sense of where this is all going. We're not too far off from being able to deliver great-looking glasses that let you seamlessly blend the physical and

digital worlds so you can feel present with anyone no matter where they are. We're starting to see the next computing platform come together, and it's pretty exciting.

We also released our newest mixed reality headset: Quest 3S. It brings the best capabilities of Quest 3 -- high quality color passthrough, a new chipset, and more -- at the much more accessible price point of \$300. Reviews are great so far, and I'm looking forward to seeing how well it does this holiday season as more people get their hands on it.

Overall, this has been a good quarter. I'm pretty amped about all the work that we're doing right now. This may be the most dynamic moment that I've seen in our industry, and I'm focused on making sure that we build some awesome things and make the most of the opportunities ahead. And if we do this well, then the potential for Meta and everyone building with us will be massive. As always, I'm grateful for everyone who is on this journey with us -- our teams, our partners, and our investors.

And now here's Susan.

**Susan Li, CFO**

Thanks Mark and good afternoon everyone.

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

Q3 total revenue was \$40.6 billion, up 19% or 20% on a constant currency basis.

Q3 total expenses were \$23.2 billion, up 14% compared to last year.

In terms of the specific line items:

Cost of revenue increased 19%, driven primarily by higher infrastructure costs.

R&D increased 21%, mostly driven by higher headcount-related expenses and infrastructure costs.

Marketing & Sales decreased 2% driven primarily by lower restructuring costs.

G&A decreased 10% driven primarily by lower legal-related expenses.

We ended the third quarter with over 72,400 employees, up 9% year-over-year, with growth primarily driven by hiring in our priority areas of monetization, infrastructure, Reality Labs, generative AI, as well as regulation and compliance.

Third quarter operating income was \$17.4 billion, representing a 43% operating margin.

Our tax rate for the quarter was 12%.

Net income was \$15.7 billion or \$6.03 per share.

Capital expenditures, including principal payments on finance leases, were \$9.2 billion, driven by investments in servers, data centers and network infrastructure. Our capital expenditures were

impacted in part by the timing of third quarter server deliveries, which will be paid for in the fourth quarter.

Free cash flow was \$15.5 billion. In Q3, we completed a debt offering of \$10.5 billion, repurchased \$8.9 billion of our Class A common stock and paid \$1.3 billion in dividends to shareholders, ending the quarter with \$70.9 billion in cash and marketable securities and \$28.8 billion in debt.

Moving now to our segment results.

I'll begin with our Family of Apps segment.

Our community across the Family of Apps continues to grow, with more than 3.2 billion people using at least one of our Family of Apps on a daily basis in September.

Q3 Total Family of Apps revenue was \$40.3 billion, up 19% year over year.

Q3 Family of Apps ad revenue was \$39.9 billion, up 19% or 20% on a constant currency basis.

Within ad revenue, the online commerce vertical was the largest contributor to year-over-year growth, followed by healthcare and entertainment and media.

On a user geography basis, ad revenue growth was strongest in Rest of World and Europe at 23% and 21%, respectively. Asia-Pacific grew 18% and North America grew 16%. On an advertiser geography basis, total revenue growth was strongest in North America and Europe at 21%. Rest of World was up 17%, while Asia Pacific was the slowest growing region at 15%, decelerating from our second quarter growth rate of 28% due mainly to lapping a period of stronger demand from China-based advertisers.

In Q3, the total number of ad impressions served across our services increased 7% and the average price per ad increased 11%. Impression growth was mainly driven by Asia-Pacific and Rest of World. Pricing growth was driven by increased advertiser demand, in part due to improved ad performance. This was partially offset by impression growth, particularly from lower-monetizing regions and surfaces.

Family of Apps other revenue was \$434 million, up 48%, driven primarily by business messaging revenue growth from our WhatsApp Business Platform.

We continue to direct the majority of our investments toward the development and operation of our Family of Apps. In Q3, Family of Apps expenses were \$18.5 billion, representing approximately 80% of our overall expenses. Family of Apps expenses were up 13%, primarily due to higher infrastructure and headcount-related expenses, partially offset by lower legal-related expenses. Family of Apps operating income was \$21.8 billion, representing a 54% operating margin.

Within our Reality Labs segment, Q3 revenue was \$270 million, up 29% driven by hardware sales. Reality Labs expenses were \$4.7 billion, up 19% year-over-year driven primarily by higher headcount-related expenses and infrastructure costs.

Reality Labs operating loss was \$4.4 billion.

Turning now to the business outlook. 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 are focused on both improving people's experiences within our apps today and investing in longer-term initiatives that have the potential to contribute to engagement in the years ahead.

We expect our content recommendations roadmap will span both of these timeframes, as we have nearer-term workstreams focused on improving recommendations as well as multi-year initiatives to develop innovative new approaches.

I'll focus first on the near-term. In the third quarter, we continued to see daily usage grow year-over-year across Facebook and Instagram, both globally and in the U.S.

On Facebook, we're seeing strong results from the global rollout of our unified video player in June. Since introducing the new experience and prediction systems that power it, we've seen a 10% increase in time spent within the Facebook video player. This month, we've entered the next phase of Facebook's video product evolution. Starting in the US & Canada, we are updating the standalone video tab to a full screen viewing experience, which will allow people to seamlessly watch videos in a more immersive experience. We expect to complete this global roll out in early 2025.

On Instagram, Reels continues to see good traction and we're making ongoing progress with our focus on promoting original content, with more than 60% of recommendations now coming from original posts in the U.S. This is helping people find unique and differentiated content on Instagram while also helping earlier stage creators get discovered.

Next, let me talk more about our multi-year roadmap for recommendations. Previously, we operated separate ranking and recommendation systems for each of our products because we found that performance did not scale if we expanded the model size and compute power beyond a certain point. However, inspired by the scaling laws we were observing with large language models, last year we developed new ranking model architectures capable of learning more effectively from significantly larger data sets. To start, we have been deploying these new architectures to our Facebook video ranking models, which has enabled us to deliver more relevant recommendations and unlock meaningful gains in watch time.

Now, we're exploring whether these new models can unlock similar improvements to recommendations on other surfaces. After that, we will look to introduce cross-surface data to these models so our systems can learn from what is interesting to someone on one surface of our apps and use it to improve their recommendations on another. This will take time to execute, and there are other explorations that we will pursue in parallel. However, over time we are optimistic that this will unlock more relevant recommendations while also leading to higher engineering efficiency as we operate a smaller number of recommendation models.

Beyond recommendations, we're making progress with our other longer-term engagement priorities, including Generative AI and Threads.

Meta AI usage continues to scale as we make it available in more countries and languages. We're seeing lifts in usage as we improve our models and have introduced a number of enhancements in

recent months to make Meta AI more helpful and engaging. Last month, we began introducing Voice, so you can speak with Meta AI more naturally, and it's now fully available in English to people in the US, Australia, Canada and New Zealand. In the US, people can now also upload photos to Meta AI to learn more about them, write captions for posts, and add, remove, or change things about their images with a simple text prompt. These are all built with our first multi-modal foundation model, Llama 3.2.

Threads remains another area where we see exciting potential. We are bringing on an increasing number of new users each quarter while depth of engagement also continues to grow. Looking ahead, we plan to introduce more features to make it even easier for people to stay up to date on topics they care about.

Now to the second driver of our revenue performance: increasing monetization efficiency. There are two parts to this work.

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

We continue to see opportunities to grow ad supply on lower monetizing surfaces, like video. Within Facebook, video engagement continues to shift to short form following the unification of our video player, and we expect this to continue with the transition of the video tab to a full screen format. This is resulting in organic video impressions growing more quickly than overall video time on Facebook, which provides more opportunities to serve ads.

Across both Facebook and Instagram, we are also continuing our broader work to optimize when and where we should show ads within a person's session. This is enabling us to drive revenue and conversion growth without increasing the number of ads.

The second part of improving monetization efficiency is enhancing marketing performance.

Similar to organic content ranking, we are finding opportunities to achieve meaningful ads performance gains by adopting new approaches to modeling. For example, we recently deployed new learning and modeling techniques that enable our ads systems to consider the sequence of actions a person takes before and after seeing an ad. Previously, our ads systems could only aggregate those actions together, without mapping the sequence. This new approach allows our systems to better anticipate how audiences will respond to specific ads. Since we adopted the new models in the first half of this year we've already seen a 2-4% increase in conversions based on testing within selected segments.

We are also evolving our ads platform to ensure that the results we drive are customized to each business' objectives and to the way they measure value. In Q3, we introduced changes to our ads ranking and optimization models to take more of the cross-publisher journey into account, which we expect to increase the Meta-attributed conversions that advertisers see in their third party analytics tools. We're also testing new features and settings for advertisers that will allow them to optimize their campaigns for what they value most, such as driving incremental conversions rather than absolute conversions.

Finally, there is continued momentum with our Advantage+ solutions, including our ad creative tools. We're seeing strong retention with advertisers using our generative AI-powered image expansion, background generation and text generation tools, and they're already driving improved performance for advertisers even at this early stage. Earlier this month we began testing our first

video generation features - video expansion and image animation. We expect to make them more broadly available by early next year.

Next, I would like to discuss our approach to capital allocation. We continue to take a long-term view in running the business, which involves investing in a portfolio of opportunities that we expect will generate returns over different time periods. We are very optimistic about the set of opportunities in front of us, and believe that investing now in both infrastructure and talent will not only accelerate our progress, but increase the likelihood of maximizing returns within each area. This includes investing in both near-term initiatives to deliver continued healthy revenue growth within our core business, as well as longer-term opportunities that have the scale to deliver compelling returns over time. Given the lead time of our longer-term investments, we also continue to maximize our flexibility so that we can react to market developments. Within Reality Labs, this has benefited us as we've evolved our roadmaps to respond to the earlier-than-expected success of smart glasses. Within generative AI, we expect significantly scaling up our infrastructure capacity now, while also prioritizing its fungibility, will similarly position us well to respond to how the technology and market develop in the years ahead.

Moving now to our financial outlook.

We expect fourth quarter 2024 total revenue to be in the range of \$45-48 billion. Our guidance assumes foreign currency is approximately neutral to year-over-year total revenue growth, based on current exchange rates.

Turning now to the expense outlook. We expect full year 2024 total expenses to be in the range of \$96-98 billion, updated from our prior range of \$96-99 billion. For Reality Labs, we continue to expect 2024 operating losses to increase meaningfully year-over-year due to our ongoing product development efforts and investments to further scale our ecosystem.

Turning now to the capex outlook. We anticipate our full year 2024 capital expenditures will be in the range of \$38-40 billion, updated from our prior range of \$37-40 billion. We continue to expect significant capital expenditures growth in 2025. Given this, along with the back-end weighted nature of our 2024 capex, we expect a significant acceleration in infrastructure expense growth next year as we recognize higher growth in depreciation and operating expenses of our expanded infrastructure fleet.

On to tax. Absent any changes to our tax landscape, we expect our fourth quarter 2024 tax rate to be in the low-teens.

In addition, we continue to monitor an active regulatory landscape, including the increasing legal and regulatory headwinds in the EU and the US that could significantly impact our business and our financial results.

In closing, this was another good quarter for our business. Our global community continues to grow, we're seeing ongoing momentum across our core priorities, and we have exciting opportunities ahead of us to drive further growth in our core business in 2025 and capitalize on the longer-term opportunities ahead.

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

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

Brian Nowak: Thanks for taking my questions. I have two, one for Mark and one for Susan. Mark, I wanted to sort of ask you about Meta AI a little bit.

Can you help us understand some of the more recurring types of interactions or query types you're seeing with this product and whether they have commercial intent? And then just over time, how do you think about building your own in-house search offering as opposed to partnering and having another player partner those queries?

And then Susan, I wanted to ask you about sort of head count because you talked a lot about sort of infrastructure investment in '25. How do we sort of think about relative headcount investments into '25 to sort of support all that infrastructure versus what you've been doing in 2024? Thanks.

Susan Li: Brian, thanks for the question. This is Susan. So your first question was around what kinds of recurring interactions that we see between people and to their usage of Meta AI. And we're seeing -- first of all, I think, as Mark mentioned, just we're excited about the progress of Meta AI.

It's obviously very early in its journey, but it continues to be on track to be the most used AI assistant in the world by end of year and it has over 500 million monthly actives.

And people are using it for many things. A number of the frequent use cases we're seeing include information gathering, help with how to tasks, which is the largest use case. But we also see people using it to go deeper on interests, to look for content on our services, for image generation, that's also been another pretty popular use case so far. And I would say, that in the near term our focus is really on making Meta AI increasingly valuable for people and if we're successful, we think there will be a broadening set of queries that people use it for, including more monetizable queries over time.

The second part of your question, Meta AI draws from content across the web to address timely questions from users, and it provides sources for those results from our search engine partners. We've integrated with Bing and Google, both of whom offer great search experiences.

Like other companies, we also train our Gen AI models on content that is publicly available online, and we crawl the web for a variety of purposes. Your second question was really, I think, around maybe how we're thinking about headcount in 2025. And we are still working through our budgeting processes for '25. That's in part why we changed our forward-looking guidance approach to give guidance in the next call.



But as we're working through this, we are looking at where there are opportunities for us to invest in our strategic priorities and that includes monetization, infrastructure, Reality Labs, Gen AI, our ongoing investments in regulation and compliance.

And we're really evaluating each of those opportunities with an eye towards what either the measurable ROI looks like or what the strategic opportunity looks like, depending on what the area is. And we're supporting that by continuing to really focus on streamlining our operations elsewhere.

So we don't have -- we don't have specifics to share about headcount growth in 2025, but that gives you a little bit of the flavor of where we are in the budgeting process.

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 question. Maybe just one building on that question from Brian and going back to Mark's comments about the learnings as you do go through the business planning process. Mark, I wanted to understand better what you continue to learn about what the biggest opportunity sets are to apply AI to when you think about your platform, your product portfolio, and your internal processes because you sound quite optimistic about key learnings and how they continue to ramp and maybe even accelerate in terms of the potential for return profile? I just want to go a little bit deeper into what your key learnings are as you go through that process? Thank you.

Mark Zuckerberg: I think the main point here is just that it seems broadly applicable to a very wide variety of products. So there are areas that are more part of our core business from making Feed more relevant and Reels more relevant to making ads more relevant to helping advertisers generate better ads, to helping people create the content that they want, helping our integrity operations and compliance and the work that we do there, that's important.

It's very valuable across all these aspects of the core business, but then it also is going to enable completely new types of services, like we didn't have something like Meta AI before.

We didn't have something like the Ray-Ban Meta glasses before and AI is going to be a really important ingredient of all of these things. There are also other new products like that, things around AI Studio.

This year, we really focused on rolling out Meta AI as kind of our -- our kind of single assistant that people can ask any question to, but I think there's a lot of opportunities that I think we'll see ramp more over the next year in terms of both consumer and business use cases, for people interacting with a wide variety of different AI agents, consumer ones with AI Studio around whether

it's different creators or kind of different agents that people create for entertainment.

Or on the business side, we do want to continue making progress on this vision of making it so that any small business or any business over time can, with a few clicks, stand up an AI agent that can help do customer service and sell things to all of their customers around the world, and I think that's a huge opportunity.

So it's very broad. And I think part of what we're seeing is that there are a lot of opportunities. Some of the longer-term ones are on Meta AI or AI Studio; those aren't necessarily a next few years, massive profit opportunity. But there are a lot of things in the core business around engagement and monetization which I think will be, over the next few years.

So I think we're trying to make sure that we get the right people working on this and that we have the right amount of investment that's just going towards what we view as a very, very large opportunity.

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

Douglas Anmuth: Great, thanks for taking the questions. Maybe just a follow-up first on Meta AI, Mark. I mean helpful context certainly to understand how people are using the platform today. Maybe you can just talk more about some of that functionality over time as agents are introduced, just how you really expect use cases to expand beyond just longer and more complex queries?

And then, Susan, on CapEx, just trying to understand your comment on 4Q a little bit more. It sounds like some of the payments pushed into 4Q with the guidance suggesting \$15 billion to \$17 billion in CapEx in the quarter? And is that something we should think about as run rate into 2025? Thank you.

Mark Zuckerberg: Yes. I mean I can take the Meta AI question, although I'm sort of intentionally now not saying too much about the new capabilities and modalities that we're launching with Llama 4 that are coming to Meta AI.

I mean I noted in the comments upfront that there -- with each major generational update, I expect that there will be large new capacities that get added.

But I think that, that's just going to be -- that's partially what I'm excited about, and we'll talk more about that next year when we're ready to. One of the trends that I do think we're going to see though is having the models not just power Meta AI or a single assistant, but across AI Studio and business agents, have that grow.

I mean this year, if you look back to where we were about a year ago, we were starting to roll out Meta AI. This year we have really so far succeeded in having

that grow and having a lot of people use that. There's obviously a lot more depth of engagement and new use cases that we want to add over time.

But I'd say that we're -- today, with AI Studio and business AIs about where we were with Meta AI about a year ago.

So I think in the next year, our goal around that is going to be to try to make those pretty widespread use cases, even though there's going to be a multiyear path to getting kind of the depth of usage and the business results around that that we want. So there's a lot to do here, though, and I'm excited to talk about that starting earlier next year.

Susan Li: Thanks, Doug. So your second question was about Q4 CapEx. The expected step-up in Q4 CapEx from Q3, part of that comes from increases in server spend and to a lesser extent, data center CapEx.

But with servers there are these timing dynamics at play that you referred to because we had these server deliveries that landed late in Q3. And so the cash doesn't go out the door basically until Q4, and that's when you'll see the CapEx show up.

And given the nature of capital expenditures, generally, there is some -- actually quite a bit of lumpiness quarter-to-quarter. So it's a little bit hard to sort of extrapolate from any particular quarter. Overall, I'd say we're growing our infrastructure investments significantly this year, and we expect significant growth again in 2025.

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

Justin Post: Great. I think I'll ask a cost question this time. Just thinking about use of AI and employee productivity, how are you able to utilize AI internally? And are you seeing big productivity gains in your R&D group? And second, I know I got to the headcount one more time but, Susan, how flexible is your headcount as you think about cost growth in other areas? Thanks.

Susan Li: Justin, so I'll take a crack at both of those. On the use of AI and employee productivity, it's certainly something that we -- that we're very excited about. I don't know that we have anything particularly quantitative that we're sharing right now.

I think there are different efficiency opportunities with AI that we've been focused on in terms of where we can reduce costs over time and generate savings through increasing internal productivity in areas like coding.

For example, it's early, but we're seeing a lot of adoption internally of our internal assistant and coding agent, and we continue to make Llama more effective at coding, which should also make this use case increasingly valuable to developers over time.

There are also places where we hope over time that we'll be able to deploy these tools against a lot of our content moderation efforts to help make the big body of content moderation work that we undertake, to help make it more efficient and effective for us to do so.

And there are lots of other places around the company where I would say we're relatively early in exploring the way that we can use LLM based tools to make different types of work streams more efficient. So all that is to say, it's something we're pretty excited about.

We have lots of teams focused on it. There are sort of small opportunities and G&A functions to what we hope will be big opportunities in areas like content moderation and coding productivity over time.

On your second question about headcount, we're really -- again, we're still mid-budget. So there's -- we don't have very much that is definitive to share about this at the time. But as we're evaluating where there are opportunities for us to make good investments, we really think about there is a bucket of very ROI-driven headcount opportunities.

We're very rigorous about the way we think about returns there and what the return opportunity is and how -- what we think is the likelihood of those returns and what is the aggregate incrementality of those investments, and those are all things that sort of we're evaluating when we think about where to invest in the core business and where we think we can -- we can deliver sort of ROI on a nearer-term basis.

And then at the same time, we're also assessing what the opportunities look like in some of the more medium- and long-term strategic areas of investment and that includes our efforts in Gen AI and the infrastructure needed to support it.

It includes our investments in Reality Labs. And so those are all things that we're kind of assessing in kind of a portfolio of what we could -- what we think we would do in 2025. With a couple of thoughts, one is, again, where can we sort of build the most flexibility into the way that we're thinking about either infrastructure or headcount plans.

And the second is we're really focused across the company on our efficiency efforts broadly and making sure that we feel like we're continuing to push the whole company including areas in which we expect that we will be making additional headcount investments to think about how they can be more efficient in '25 than they were in '24.

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

Ross Sandler: Great. Just two quick ones, Mark. You said something along the lines of the more standardized Llama becomes the more improvements will flow back to the core Meta business. And I guess, could you just dig in a little bit more on

that? So the series of Llama models are being used by lots of developers building different things in AI.

I guess, how are you using that vantage point to incubate new ideas inside Meta? And then second question is, you mentioned on one of the podcasts after the Meta Connect that assuming scaling laws hold up, we may need hundreds of billions of compute CapEx to kind of reach our goals around Gen AI.

So I guess, how quickly could you conceivably stand up that much infrastructure given some of the constraints around energy or custom ASICs or other factors? Just any more color on the speed by which we could get that amount of compute online at Meta? Thank you.

Mark Zuckerberg: Yes. I can try to give some more color on this. I mean the improvements to Llama, I'd say come in a couple of flavors. There's sort of the quality flavor and the efficiency flavor.

There are a lot of researchers and independent developers who do work and because Llama is available, they do the work on Llama and they make improvements and then they publish it and it becomes -- it's very easy for us to then incorporate that both back into Llama and into our Meta products like Meta AI or AI Studio or Business AIs because the work -- the examples that are being shown are people doing it on our stack.

Perhaps more importantly is just the efficiency and cost. I mean this stuff is obviously very expensive. When someone figures out a way to run this better, if that -- if they can run it 20% more effectively, then that will save us a huge amount of money.

And that was sort of the experience that we had with Open Compute and why -- part of why we are leaning so much into Open Source here in the first place, is that we found counterintuitively with Open Compute that by publishing and sharing the architectures and designs that we had for our compute, the industry standardized around it a bit more.

We got some suggestions also that helped us save costs and that just ended up being really valuable for us. Here, one of the big costs is chips -- a lot of the infrastructure there.

What we're seeing is that as Llama gets adopted more, you're seeing folks like NVIDIA and AMD optimize their chips more to run Llama specifically well which clearly benefits us.

So it benefits everyone who's using Llama, but it makes our products better, right, rather than if we were just on an island building a model that no one was kind of standardizing around in the industry.

So that's some of what we're seeing around Llama and why I think it's good business for us to do this in an open way.

In terms of scaling infra, I mean when I talk about our teams executing well, some of that goes towards delivering more engaging products and some of it goes towards delivering more revenue.

On the infra side, it goes towards building out the expenses faster, right? So I think part of what we're seeing this year is the infra team is executing quite well. And I think that's why, over the course of the year, we've been able to build out more capacity.

I mean going into the year, we had a range for what we thought we could potentially do. And we have been able to do, I think, more than, I think, we'd kind of hoped and expected at the beginning of the year.

And while that reflects as higher expenses, it's actually something that I'm quite happy that the team is executing well on. And I think that will -- so that execution makes me somewhat more optimistic that we're going to be able to keep on building this out at a good pace but that's part of this whole thing. This is -- this part of the formula around kind of building out the infrastructure is maybe not what investors want to hear in the near term that we're growing it out.

But I just think that the opportunities here are really big. We're going to continue investing significantly in this. And I'm proud of the teams that are doing great work to stand up a large amount of capacity so that way we can deliver world-class models and world-class products.

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

Ronald Josey: Great, thanks for taking the question. Maybe a bigger picture one as well, Mark, just this time on Threads. Now one of the core apps, and on it's way to becoming the next major social app and 275 million MAUs.

I wanted to hear your thoughts on how this product evolves over time, specifically from a monetization perspective, but also next steps on users? And then, Susan, with pricing up 11% in the quarter, I wanted to hear more about the pricing dynamics on the platform.

I think you talked about just pricing increasing due to greater advertising demand and improved ad performance. So help us understand that a little bit more. Thank you.

Susan Li: Thanks, Ron. So your first question was about Threads. We're making good progress there. We are continuing to launch more features and make improvements to our ranking stack. We feel very good about the continued user growth on Threads.

We're bringing on an increasing number of users each quarter and depth of engagement also continues to grow. And in Q3, we saw especially strong user

growth in key markets like the U.S., Taiwan and Japan. And we -- we've added a number of new features over the course of Q3 including account insights for businesses and creators to see how they're posts perform, the ability to save multiple drafts, continuing to deliver on our commitment to integrate Threads with the Fediverse and basically, we're very focused on continuing to build out the sort of functionality of Threads over time and being responsive to what users tell us that they're interested in.

Specifically, as it pertains to monetization, we don't expect Threads to be a meaningful driver of 2025 revenue at this time.

We've been just pleased with the growth trajectory and again, are really focused on introducing features that the community finds valuable and working to deepen growth and engagement. Your second question was about the increase in average price per ad.

So that grew 11% year-over-year, driven by strong advertiser demand and part of that is because of better ad performance over time also. And we saw that CPM growth accelerate slightly from 10% in Q2, in part because we experienced lower impression growth in Q3.

But more broadly, as we think about pricing growth and this metric, the year-over-year growth in reported price per ad, there's a lot that goes into that including the auction dynamics resulting from fluctuations and impression growth. And one of the things that we feel like we're very focused on is really the input metrics.

What are the conversions that we are delivering to advertisers? Are they getting more value over time? The sort of blended reported price per ad is complicated because all of those things get rolled up into it.

There are so many different objectives that advertisers are optimizing for, those objectives have very different values that make them hard to compare on an apples-to-apples basis.

But we care a lot about conversion growth, which is growing, continues to grow faster than impression growth. And are we seeing healthy cost per action or cost per conversion trends, which we are. And as long as we continue to get better at driving conversions for advertisers, that should have the effect of lifting CPMs over time, because we're delivering more conversions per impressions served and that will result in higher value impressions.

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

Kenneth Gawrelski: Thanks for the opportunity. I appreciate that. I have a bigger picture kind of like ecosystem question here. Is -- when we think -- I'm curious, how far do you think we are from seeing a proliferation of third-party AI applications, specifically on the kind of the consumer side? I know we're seeing more and more on the enterprise side, agents, et cetera.

But when do we see -- how far out until we see a proliferation of consumer applications in the AI space? And how would you think about -- and how does Meta think of itself as one of those key applications in the mobile internet and the desktop internet? But now you're also seemingly an infrastructure player as well. So I'd love to hear your thoughts there. Thank you.

Mark Zuckerberg: Yes. I mean there are a lot of consumer products that we're working on. And with Llama, I would expect that app developers will be able to build a lot of really good things, too.

I've touched on Meta AI and AI Studio and Business AIs a bunch, and I expect those to be important parts of the consumer experience. Another part that I haven't talked about quite as much yet is the opportunity for AI to help people create content that just makes people's feed experiences better.

If you look at the big trends in Feeds over the history of the company, it started off as friends, right? So all the updates that were in there were basically from your friends posting things. And then we went into this era where we added in creator content too, where now a very large percent of the content on Instagram and Facebook is not from your friends.

It may not even be from people that you're following directly. It could just be recommended content from creators that we can algorithmically determine is going to be interesting and engaging and valuable to you.

And I think we're going to add a whole new category of content, which is AI generated or AI summarized content or kind of existing content pulled together by AI in some way. And I think that that's going to be just very exciting for the - - for Facebook and Instagram and maybe Threads or other kind of Feed experiences over time.

It's something that, we're starting to test different things around this. I don't know if we know exactly what's going to work really well yet. Some things are promising. I don't know that -- this isn't going to be a big impact on the business in '25 would be my guess.

But I think that there is -- I have high confidence that over the next several years, this is going to be an important trend and one of the important applications. But you're going to get that, you're going to get Meta AI, AI Studio, business AIs and a whole lot of things that developers would do with Llama too.

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

Youssef Squali: Yes, thank you very much. Mark, it appears that Meta AI now crawls the web and provides conversational answers about pretty much anything including current events. And so with over ten million advertisers and one of the best ROAS offerings out there in your core business, just wondering if there are any



plans to start maybe testing ads on commercial queries, and move Meta AI closer to becoming a real answer engine for the billions of queries that you guys are already seeing?

And then Susan, one of the biggest areas of pushback we get is around Reality Labs and the ongoing losses there, I think \$16 billion last year, probably north of \$20 billion this year. The question is, are we getting any closer to peak losses there or alternatively, what products do you think have the biggest potential there over the next couple of years? Thanks.

Susan Li: I'm happy to take both of these, Youssef. So your first question was on plans to provide ads on commercial queries. I think I alluded to this maybe in a much earlier question. Right now, we're really focused on making Meta AI as engaging and valuable a consumer experience as possible.

Over time, we think there will be a broadening set of queries that people use it for. And I think that the monetization opportunities will exist when -- over time as we get there.

But right now, I would say we are really focused on the consumer experience above all and this is sort of a playbook for us with products that we put out in the world where we really dial in the consumer experience before we focus on what the monetization could look like. The second part of your question is about Reality Labs.

We aren't sharing expectations beyond 2024 at this point. And we are certainly as we think about the 2025 budgeting process for Reality Labs, we're certainly thinking about where we want to make sure we're putting our sort of focus and energy.

We are very excited, again, about the progress that we've seen with our smart glasses as well as the sort of strong consumer interest in them. And so, we're kind of thinking about where we want to make sure that we are investing appropriately behind the consumer momentum that we see.

Overall, I'd say Reality Labs is clearly one of our strategic long-term priorities, and we expect it will be an area of significant investment in -- as we build out towards the very ambitious product roadmap that we have there.

Kenneth Dorell: Krista, we have time for one last question.

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

Mark Mahaney: Let me throw in two questions, please. One, this is a year in which we've had a lot of unusual events that could be driving ad revenue, major elections in not just the U.S., but Europe and in India and major sports events like World Cup and I know there's some other things.

Is there any -- just in thinking about comps for next year and maybe in the future, anything Susan, you would call out, like how much of an impact there may have come from these one in every four or five-year events? And then secondly, could you just talk a little bit more about WhatsApp monetization and where you are with that now?

It sounds like the business messaging part is really feeding in nicely into other revenue. But help us think about where the monetization levels of WhatsApp are now versus where they can be two or three years down the road, how far away we are from optimization? Thank you.

Susan Li:

Thanks, Mark. So your first question was about sort of the revenue backdrop in 2024, you mentioned events that occur once every four or five years. So I imagine we're talking about the Olympics. We historically have not seen meaningful incremental contribution from events like the Olympics. We believe that was largely the case this year.

So when we think about the Q4 outlook and when we think about going into next year, we generally expect growth to continue to benefit from the healthy global advertising demand that we've seen.

We think that our investments in improving our ads performance will continue to accrue benefits to advertisers. But obviously, there's a big range of possible macro backdrops and that's something that we try to reflect in the range of revenue guidance that we give.

So -- but I don't know that there are a lot of specific events that we would say had a material sort of idiosyncratic to 2024 type of revenue impact. Your second question was around WhatsApp monetization and where we are. And right now what I would say, again, is Click-to-Message is really the big focus area for us here.

We're seeing continued traction in this area. And in particular, growth in Click-to-WhatsApp Ads remain particularly strong. And so we're continuing to focus both on scaling Click-to-WhatsApp Ads in more markets where WhatsApp has strong user adoption like Brazil, for example.

It's obviously earlier in the U.S., but we're seeing good growth in Click-to-WhatsApp Ads and are continuing to invest in scaling and consumer adoption of WhatsApp in the U.S. also, which will create bigger opportunities down the line.

So -- and then, of course, a lot of work that we're doing to make the Click-to-Messaging ads more effective and helping to focus for the particular -- helping advertisers optimize, sorry, for the particular conversion events that they care about.

The other element of revenue on WhatsApp, I would say, is paid messaging. That continues to grow at a strong pace again this quarter.

It remains, in fact, the primary driver of growth in our Family of Apps other revenue line, which was up 48% in Q3, and we're seeing generally a strong increase in the volume of paid conversations driven both by growth in the number of businesses adopting paid messaging as well as in the conversational volume per business.

Kenneth Dorell: Great. Thank you for joining us today. We appreciate your time. And we look forward to speaking with you again soon.

Operator: This concludes today's conference call. Thank you for your participation, and you may now disconnect.