

**Meta Platforms, Inc. (META)  
Fourth Quarter 2024 Results Conference Call  
January 29<sup>th</sup>, 2025**

**Kenneth Dorell, Director, Investor Relations**

Thank you. Good afternoon and welcome to Meta Platforms fourth quarter and full year 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.atmeta.com](http://investor.atmeta.com).

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

**Mark Zuckerberg, CEO**

Thanks Ken and thanks everyone for joining today.

We ended 2024 on a strong note with now more than 3.3 billion people using at least one of our apps each day. This is going to be a really big year. I know it always feels like every year is a big year, but more than usual it feels like the trajectory for most of our long-term initiatives is going to be a lot clearer by the end of this year. So I keep telling our teams that this is going to be intense, because we have about 48 weeks to get on the trajectory that we want to be on.

In AI, I expect this is going to be the year when a highly intelligent and personalized AI assistant reaches more than 1 billion people, and I expect Meta AI to be that leading AI assistant. Meta AI is already used by more people than any other assistant, and once a service reaches that kind of scale it usually develops a durable long-term advantage. We have a really exciting roadmap for this year with a unique vision focused on personalization. We believe that people don't all want to use the same AI -- people want their AI to be personalized to their context, their interests, their personality, their culture, and how they think about the world. I don't think that there's just going to be one big AI that everyone uses that does the same thing. People are going to get to choose how their AI works and what it looks like for them. I continue to think that this is going to be one of the most transformative products that we've made. We have some fun surprises that I think people are going to like this year.

I think this will very well be the year when Llama and open source become the most advanced and widely used AI models as well. Llama 4 is making great progress in training. Llama 4 mini is done

with pre-training and our reasoning models and larger model are looking good too. Our goal with Llama 3 was to make open source competitive with closed models, and our goal for Llama 4 is to lead. Llama 4 will be natively multimodal -- it's an omni-model -- and it will have agentic capabilities, so it's going to be novel and it's going to unlock a lot of new use cases. I'm looking forward to sharing more of our plan for the year on that over the next couple of months.

I also expect that 2025 will be the year when it becomes possible to build an AI engineering agent that has coding and problem-solving abilities of around a good mid-level engineer. This is going to be a profound milestone and potentially one of the most important innovations in history, as well as over time, potentially a very large market. Whichever company builds this first I think is going to have a meaningful advantage in deploying it to advance their AI research and shape the field. So that's another reason why I think that this year is going to set the course for the future.

Our Ray-Ban Meta AI glasses are a real hit, and this will be the year when we understand the trajectory for AI glasses as a category. Many breakout products in the history of consumer electronics have sold 5-10 million units in their third generation. This will be a defining year that determines if we're on a path towards many hundreds of millions and eventually billions of AI glasses -- and glasses being the next computing platform like we've been talking about for some time -- or if this is just going to be a longer grind. But it's great overall to see people recognizing that these glasses are the perfect form factor for AI -- as well as just great, stylish glasses.

These are all big investments -- especially the hundreds of billions of dollars that we will invest in AI infrastructure over the long term. I announced last week that we expect to bring online almost 1GW of capacity this year, and we're building a 2GW and potentially bigger AI datacenter that is so big that it'll cover a significant part of Manhattan if it were placed there.

We're planning to fund all this by at the same time investing aggressively in initiatives that use these AI advances to increase revenue growth. We've put together a plan that will hopefully accelerate the pace of these initiatives over the next few years. That's what a lot of our new headcount growth is going towards. And how well we execute on this will also determine our financial trajectory over the next few years.

There are a number of other important product trends related to our family of apps that I think we're going to know more about this year as well. We're going to learn what's going to happen with TikTok, and regardless of that I expect Reels on Instagram and Facebook to continue growing. I expect Threads to continue on its trajectory to become the leading discussion platform and eventually reach 1 billion people over the next several years. Threads now has more than 320 million monthly actives and has been adding more than 1 million sign-ups per day. I expect WhatsApp to continue gaining share and making progress towards becoming the leading messaging platform in the US like it is in a lot of the rest of the world. WhatsApp now has more than 100 million monthly actives in the US. Facebook is used by more than 3 billion monthly actives and we're focused on growing its cultural influence. I'm excited this year to get back to some OG Facebook.

This is also going to be a pivotal year for the metaverse. The number of people using Quest and Horizon has been steadily growing -- and this is the year when a number of the long-term investments that we've been working on that will make the metaverse more visually stunning and inspiring will really start to land. So I think we're going to know a lot more about Horizon's trajectory by the end of this year.

This is also going to be a big year for redefining our relationship with governments. We now have a US administration that is proud of our leading companies, prioritizes American technology winning, and that will defend our values and interests abroad. And I'm optimistic about the progress and innovation that this can unlock.

So this is going to be a big year. I think this is the most exciting and dynamic that I have ever seen our industry. Between AI, glasses, massive infrastructure projects, doing a bunch of work to try to accelerate our business, and building the future of social media, we have a lot to do. I think we're going to build some awesome things that shape the future of human connection. As always, I'm grateful for everyone who is on this journey with us. Thank you and 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.

Q4 total revenue was \$48.4 billion, up 21% on both a reported and constant currency basis.

Q4 total expenses were \$25.0 billion, up 5% compared to last year. Before I cover the specific cost lines, I would note that our fourth quarter expense growth rate reflects a 13 percentage point favorable impact from legal accrual reductions in Q4 and lower year-over-year restructuring costs.

In terms of the specific line items:

Cost of revenue increased 15%, driven mostly by higher infrastructure costs.

R&D increased 16%, primarily driven by higher employee compensation and infrastructure costs, which were partially offset by lower restructuring costs.

Marketing & Sales were approximately flat year-over-year.

G&A decreased 67% driven mostly by lower legal-related expenses due to a \$1.55B reduction in legal accruals related to certain legal proceedings.

We ended the year with over 74,000 employees, up 10% year-over-year, with growth primarily driven by hiring in priority areas of monetization, infrastructure, generative AI, Reality Labs, as well as regulation and compliance.

Fourth quarter operating income was \$23.4 billion, representing a 48% operating margin.

Our tax rate for the quarter was 12%.

Net income was \$20.8 billion or \$8.02 per share.

Capital expenditures, including principal payments on finance leases, were \$14.8 billion, driven by investments in servers, data centers and network infrastructure.

Free cash flow was \$13.2 billion. We paid \$1.3 billion in dividends to shareholders, ending the year with \$77.8 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, and we estimate more than 3.3 billion people used at least one of our Family of Apps on a daily basis in December.

Q4 Total Family of Apps revenue was \$47.3 billion, up 21% year over year.

Q4 Family of Apps ad revenue was \$46.8 billion, up 21% on both a reported and constant currency basis.

Within ad revenue, the online commerce vertical was the largest contributor to year-over-year growth.

On a user geography basis, ad revenue growth was strongest in Rest of World at 27%, followed by Asia-Pacific and Europe at 23% and 22%, respectively. North America grew 18%.

In Q4, the total number of ad impressions served across our services increased 6% and the average price per ad increased 14%. Impression growth was mainly driven by Asia-Pacific. Pricing growth benefited from increased advertiser demand, in part driven by improved ad performance. This was partially offset by impression growth, particularly from lower-monetizing regions and surfaces.

Family of Apps other revenue was \$519 million, up 55%, 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 Q4, Family of Apps expenses were \$19.0 billion, representing 76% of our overall expenses. Family of Apps expenses were up 5%, primarily due to growth in infrastructure costs and employee compensation, which were partially offset by lower legal-related expenses.

Family of Apps operating income was \$28.3 billion, representing a 60% operating margin. Within our Reality Labs segment, Q4 revenue was \$1.1 billion, driven by hardware sales and up 1% year-over-year.

Reality Labs expenses were \$6.0 billion, up 6% year-over-year driven primarily by higher infrastructure costs and employee compensation, partially offset by lower restructuring costs.

Reality Labs operating loss was \$5.0 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, daily actives continue to grow across Facebook, Instagram and WhatsApp year-over-year, both globally and in the United States. In Q4, global video time grew at double digit

percentages year-over-year on Instagram, and we're seeing particular strength in the US on Facebook, where video time spent was also up double digit rates year-over-year. We see continued opportunities to drive video growth in 2025 through ongoing optimizations to our ranking systems. We're also making several product bets that are focused on setting up our platforms for longer-term success.

Creators are one of our central focuses. On Instagram, we continue to prioritize original posts in recommendations to help smaller creators get discovered. We also want to ensure creators have a place to experiment with their content, so we introduced a new feature in Q4 that allows creators to first share a Reel with people who don't follow them. This allows them to test content and see what performs best before deciding to share it with their followers, and also helps introduce them to entirely new audiences. Creative tools is another area we're investing in. In the coming weeks, we'll launch a new standalone app called Edits that provides a full suite of creative tools to make it easier for creators to make great Reels on their phone.

Another focus is making it easier for people to connect over content. Reels are already reshared over 4.5 billion times a day, and we've been introducing more features that bring together the social and entertainment aspects of Instagram. In the US, we recently launched a new destination in Reels that consists of content your friends have left a note on or liked. We're seeing very positive early results and will look to expand this globally in the coming months.

On Threads, we made tremendous progress in 2024 and our focus this year is establishing Threads as the place people come to keep up with what they care about. We're making a number of updates to our recommendation systems to prioritize more recent posts, surface content from top creators, and ensure people see more of the content from accounts they follow. We will also continue improving custom feeds so people can build personalized feeds on topics they're interested in.

Finally, Meta AI usage continues to scale, with more than 700 million monthly actives. We're now introducing updates that will enable Meta AI to deliver more personalized and relevant responses by remembering certain details from people's prior queries and considering what they engage with on Facebook and Instagram to develop better intuition for their interests and preferences.

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

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

We continue to grow supply on lower monetizing surfaces, like video, while optimizing ad supply on each of our surfaces to deliver ads at the time and place they will be most relevant to people. For example, we are continuing to better personalize when ads show up, including the optimal locations in the depth of someone's feed, to introduce ad supply when it's most optimal for the user and revenue. This is enabling efficient supply growth.

Longer term, we also see impression growth opportunities on unmonetized surfaces like Threads, which we are beginning to test ads on this quarter. We expect the introduction of ads on Threads will be gradual and don't anticipate it being a meaningful driver of overall impression or revenue growth in 2025.

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

The ongoing enhancements to our ads ranking systems are an important driver of this work.

In the second half of 2024, we introduced an innovative new machine learning system in partnership with Nvidia, called Andromeda. This more efficient system enabled a 10,000x increase in the complexity of models we use for ads retrieval, which is the part of the ranking process where we narrow down a pool of tens of millions of ads to the few thousand we consider showing someone. The increase in model complexity is enabling us to run far more sophisticated prediction models to better personalize which ads we show someone. This has driven an 8% increase in the quality of ads that people see on objectives we've tested. Andromeda's ability to efficiently process larger volumes of ads also positions us well for the future as advertisers use our generative AI tools to create and test more ads.

Another way we're delivering value for advertisers is through increased automation of their ad campaigns with Advantage+.

Adoption of Advantage+ shopping campaigns continues to scale, with revenue surpassing a \$20 billion annual run-rate and growing 70% year-over-year in Q4. Given the strong performance and interest we're seeing in Advantage+ Shopping and our other end-to-end solutions, we're testing a new streamlined campaign creation flow so advertisers no longer need to choose between running a manual or Advantage+ Sales or App campaign. In this new setup, all campaigns optimizing for sales, app or lead objectives will have Advantage+ turned on from the beginning. This will allow more advertisers to take advantage of the performance Advantage+ offers while still having the ability to further customize aspects of their campaigns when they need to. We plan to expand to more advertisers in the coming months before fully rolling it out later in the year.

Advantage+ creative is another area where we're seeing momentum. More than 4 million advertisers are now using at least one of our generative AI ad creative tools, up from one million six months ago. There has been significant early adoption of our first video generation tool that we rolled out in October, Image Animation, with hundreds of thousands of advertisers already using it monthly.

Next, I would like to discuss our approach to capital allocation. Our primary focus remains investing capital back into the business, with infrastructure and talent being our top priorities.

On the first, we expect compute will be central to many of the opportunities we're pursuing as we advance the capabilities of Llama, drive increased usage of generative AI products and features across our platform, and fuel core ads and organic engagement initiatives.

We're working to meet the growing capacity needs for these services by both scaling our infrastructure footprint and increasing the efficiency of our workloads. Another way we're pursuing efficiencies is by extending the useful lives of our servers and associated networking equipment. Our expectation going forward is that we'll be able to use both our non-AI and AI servers for a longer period of time before replacing them, which we estimate will be approximately five and a half years. This will deliver savings in annual capex and resulting depreciation expense, which is already included in our guidance.

Finally, we're pursuing cost efficiencies by deploying our custom MTIA silicon in areas where we can achieve a lower cost of compute by optimizing the chip to our unique workloads. In 2024 we started deploying MTIA to our ranking and recommendation inference workloads for ads and organic content. We expect to further ramp adoption of MTIA for these use cases throughout

2025 before extending our custom silicon efforts to training workloads for ranking and recommendations next year.

From a hiring standpoint, our focus continues to be on adding technical talent to support our strategic priorities.

In the fourth quarter, nearly 90% of our year-over-year headcount growth was within the R&D function. The remaining growth was primarily in cost of revenue as we added infrastructure headcount to support our data center operations.

In 2025, we expect headcount growth will continue to be primarily driven by technical roles across our priority initiatives within infrastructure, monetization, Reality Labs, generative AI, as well as regulation and compliance. We anticipate headcount growth in our business functions will remain relatively limited.

To achieve our ambitions in these areas, we will need to continue executing at a rapid pace. We're supporting this by building tools to help our engineering base be more productive. As part of our efficiency focus over the past two years, we've made significant improvements in our internal processes and developer tools and introduced new tools like our AI-powered coding assistant, which is helping our engineers write code more quickly. Looking forward, we expect that the continuous advancements in Llama's coding capabilities will provide even greater leverage to our engineers, and we are focused on expanding its capabilities to not only assist our engineers in writing and reviewing our code, but also to begin generating code changes to automate tool updates and improve the quality of our code base.

Finally, we expect our strong financial position will enable us to support these investments while continuing to return capital to shareholders through share repurchases and dividends.

Moving to our financial outlook.

We expect first quarter 2025 total revenue to be in the range of \$39.5-41.8 billion. This reflects 8-15% year-over-year growth, or 11-18% growth on a constant currency basis as our guidance assumes foreign currency is an approximately 3% headwind to year-over-year total revenue growth, based on current exchange rates. This also reflects the effect of lapping leap day in the first quarter of 2024. While we are not providing a full year 2025 revenue outlook, we expect the investments we're making in our core business this year will give us an opportunity to continue delivering strong revenue growth throughout 2025.

Turning now to the expense outlook. We expect full year 2025 total expenses to be in the range of \$114-119 billion.

We expect the single largest driver of expense growth in 2025 to be infrastructure costs, driven by higher operating expenses and depreciation. We expect employee compensation to be the second-largest factor as we add technical talent in the priority areas that I referenced earlier.

Turning now to the capex outlook. We anticipate our full year 2025 capital expenditures will be in the range of \$60-65 billion. We expect capex growth in 2025 will be driven by increased investment to support both our generative AI efforts and our core business. The majority of our capex in 2025 will continue to be directed toward our core business.

On to tax. Absent any changes to our tax landscape, we expect our full year 2025 tax rate to be in the range of 12-15%.

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

In closing, this was a good year for our company, with investments across our priority areas delivering strong business performance and innovative new products for our community. We have a compelling set of opportunities to invest in this year, which we expect will help us drive continued strong growth and develop transformative technologies that shape the future of our company and of the industry.

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 our first question comes from the line of Brian Nowak with Morgan Stanley. Please go ahead.

Brian Nowak: Thanks for taking my questions. Mark, I appreciate all the excitement about this year and all the innovation to come. I know there's a lot of announcements over the course of the year, but I wonder if you could just share a few sort of high-level examples of your vision on new potential use cases and offerings that could drive utility for your users and value for your advertisers as you sort of think about Llama 4 and Meta AI changing throughout 2025?

And then the second one on custom silicon, maybe a question for either of you. Just any learnings on the difference between your custom silicon and third-party chips in your ranking models and results? And how should we think about the main gating factors as to how quickly you'd be able to move a higher percentage of your engagement to your custom silicon?

Mark Zuckerberg: On the first one, I tried to lay this out in my opening comments a bit. I mean we're very focused on Meta AI as a highly intelligent and personalized assistant that you can access across our apps. There's a website, you can access it outside of our apps, too.

I think that the quality of this is just - it's going to keep on improving and improved a lot over the last year. We're also finding more ways that it's useful to integrate it into our services to help more people discover it.

I think that that's undoubtedly why so many hundreds of millions of people are using it today, is because it's kind of easy to discover what we're doing and then keep using it.



I want to keep some surprises and fun for the stuff that we're going to release this year. I gave a bit of detail on what we're planning to do with Llama 4 that I'm sure technical people will enjoy because we haven't talked about that before.

But I'm going to refrain from adding a whole lot more on what we're launching this year. But it's the different things that I talked about. It's Meta AI. I do expect Llama 4 to be a very exciting set of releases. It's not just one thing. Just like with Llama 3, there were kind of a few different models at different dates, I think we'll see that with Llama 4 too. And then the AI engineer piece, I'm really excited about it.

I mean I don't know that that's going to be an external product anytime soon. But I think for what we're working on, our goal is to advance AI research and advance our own development internally. And I think it's just going to be a very profound thing.

So I mean that's something that I think will show up through making our products better over time. But -- and then as that works, there will potentially be a market opportunity down the road.

But I mean for now and this year, we're really -- I think this is -- I don't think you're going to see this year like an AI engineer that is extremely widely deployed, changing all of development.

I think this is going to be the year where that really starts to become possible and lays the groundwork for a much more dramatic change in '26 and beyond. I don't know, yeah, that's kind of -- that's kind of it.

Susan Li:

Brian, I'm happy to take your second question about custom silicon. So first of all, we expect that we are continuing to purchase third-party silicon from leading providers in the industry.

And we are certainly committed to those long-standing partnerships, but we're also very invested in developing our own custom silicon for unique workloads, where off-the-shelf silicon isn't necessarily optimal and specifically because we're able to optimize the full stack to achieve greater compute efficiency and performance per cost and power because our workloads might require a different mix of memory versus network, bandwidth versus compute, and so we can optimize that really to the specific needs of our different types of workloads.

Right now the in-house MTIA program is focused on supporting our core ranking and recommendation inference workloads. We started adopting MTIA in the first half of 2024 for core ranking and recommendations inference.

We'll continue ramping adoption for those workloads over the course of 2025 as we use it for both incremental capacity and to replace some GPU-based servers when they reach the end of their useful lives. Next year, we're hoping to

expand MTIA to support some of our core AI training workloads and over time some of our Gen AI use cases.

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

Eric Sheridan: Thank you so much for taking the question. Maybe I can go back to your comments on open source. Can you help us understand how your views continue to evolve with respect to the competitive dynamic around your approach with open source versus others in the industry? And how your approach to open source could possibly bend the cost curve and improve return on capital for AI over the medium to long term? Thanks so much.

Mark Zuckerberg: Yeah, I mean on open source, I think the best analogy for us is what we did with open compute, where we weren't first to building the system. So then by the time that we got around to building it, it wasn't really a big advantage to have it be proprietary.

So we shared it. And then a lot of the industry adopted what we were doing, contributed innovations back to it. By standardizing on it, that meant that a bunch of supply chain standardized on building it, which made prices more efficient for everyone.

I think what we see here is as Llama becomes more used, it's more likely, for example, that silicon providers and others -- other APIs and developer platforms will optimize their work more for that and basically drive down the costs of using it and drive improvements that we can, in some cases, use too.

So I think that the strategy will continue to be effective, and yeah, I mean I continue to be optimistic on this. I think it's kind of -- I think it's working.

I also just think in light of some of the recent news, the new competitor DeepSeek from China, I think it also just puts -- it's one of the things that we're talking about is there's going to be an open source standard globally.

And I think for our kind of national advantage, it's important that it's an American standard. So we take that seriously and we want to build the AI system that people around the world are using and I think that if anything, some of the recent news has only strengthened our conviction that this is the right thing for us to be focused on.

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

Mark Shmulik: Thank you for taking my questions. Mark, appreciate we may get an answer this year. But looking out, as you kind of track the progress of smart glasses, Orion and so forth, do you view that as a better form factor to get the most out of the Meta AI assistant you highlighted in your opening remarks? Or is it more complementary to kind of the in-app experience in the way you've seen people use it today?

And then, Susan, the last few quarters, we've kind of seen pricing growth as the dominant driver of ad revenue growth. Given the efforts you've highlighted around driving deeper, more commercial engagement and better advertiser ROI, how do we just think about the contribution of the formula for ad revenue growth going forward? Thank you.

Mark Zuckerberg: Yeah. I mean I can talk about glasses. I mean it's -- I mean I've said for a while that I think that glasses are the ideal form factor for an AI device because you can let an AI assistant on your glasses see what you see and hear what you hear, which gives it the context to be able to understand everything that's going on in your life that you would want to talk to it about and get context on.

So -- but look, I think that glasses are going to be a very important computing platform in the future. When phones became the primary computing platform, it's not like computers went away.

I think we'll have phones for some time. But there are a lot of people in the world who have glasses.

It's kind of hard for me to imagine that a decade or more from now all the glasses aren't going to basically be AI glasses as well as a lot of people who don't wear glasses today, finding that to be a useful thing.

So I'm incredibly optimistic about this. And like I shared last year, I think one of the big surprises last year was I previously thought that glasses weren't going to become a major form factor until we got these -- the full kind of holographic displays that we started showing in the prototype for Orion.

But now I think it's pretty clear that AI is actually going to drive at least as much of the value as the holographic AR is. So that's a cause to be excited.

But look, the Ray-Ban Metas were a hit. We still don't know what the long-term trajectory for this is going to be. And I think we're going to learn a lot this year. So I think that this is a really important year for that.

Susan Li: And I can take the second question on pricing growth. So first of all, what I would say is over the long term, we think we have continued opportunity to drive revenue growth across both pricing and impression growth, so both sort of supply and demand dimensions. When we look at pricing, our reported growth can be influenced by different factors such as supply because of the auction dynamics by the mix shift of the different types of surfaces where ads show up.

For example, services like video are lower monetization efficiency, relatively speaking, and then, of course, broader macro factors.

But we generally expect that we are going to be able to deliver ongoing ad performance improvements through a lot of the ongoing work that we're doing across our monetization roadmap and that will have the sort of effect of

benefiting pricing overall. And part of what I think is kind of important to think about here when we think about price growth is, we really -- the average price per ad as we report it, is really blending, it's an output metric.

It's blending a lot of things that are happening including what are advertisers bidding for, what are their bids for those things? What is the average cost of their actions? So given that there are so many different objectives that advertisers can optimize for that have different values, it's a very complex metric that tries to distill that into one thing.

Overall, we are seeing healthy cost per action trends for advertisers for whatever is the action that they are optimizing for. And we believe we'll continue to get better at driving conversions for advertisers. And when we do, that will have the effect of continuing to lift CPMs over time because we're delivering more conversions per impressions served, resulting in higher value impressions.

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

Justin Post: Great, thanks. Maybe one for Mark and one for Susan. Mark, you mentioned political changes in the U.S. and better positioning maybe for U.S. companies abroad.

But how do you think about it in the U.S. as far as usage and advertiser adoption, you got rid of fact checking. So is -- do you think content could change? Could it appeal to more users? Will that impact advertising at all?

And then Susan, on Meta AI, I know people are pretty excited about the use case, but also thinking about the revenue case. How do you think about monetizing that? Could it be CPC ads? Or how are you thinking about that? Thank you.

Mark Zuckerberg: The question was about fact checking and our content policies. I mean look, I think we're trying to build the service that we think is the best for people.

I've believed in free expression for quite a while. People don't want to see misinformation, but you need to build an effective system that gives people more context. And I think what we found over time is that the community notes system, I think, is just going to be more effective than the system that we had before.

And I'm not afraid to admit when someone does something that's better than us. I think it's sort of our job to go and just do the best work and implement the best systems.

So I think that there's been a lot of people who have read this announcement as if we somehow don't care about adding context to things that are on our platform that are misinformation, that's not right.

I actually think that the community notes system, like what X has had for a while is actually just more effective than what we were doing before. And I think our product is going to get better because of it.

Susan Li: I would add to that, just to say, we also haven't seen any noticeable impact from our content policy changes on advertiser spend.

So we're continuing to see strong advertiser demand. Again, particularly for AI-powered tools that are helping businesses maximize the value of their ad spend. So our commitment to brand safety is unchanged, and we expect that we will invest in our suite of tools to meet the needs of advertisers.

On your second question in terms of monetizing Meta AI, our initial focus for Meta AI is really about building a great consumer experience, and that's frankly, where all of our energies are kind of directed to right now.

There will, I think, be pretty clear monetization opportunities here over time including paid recommendations and including a premium offering, but that's really not where we are focused in terms of the development of Meta AI today.

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

Douglas Anmuth: Thanks for taking the questions. One for Mark, one for Susan. Mark, just following up on open source as DeepSeek and other models potentially leverage Llama or others to train faster and cheaper, how does this impact Meta in your view? And what could it mean for the trajectory of investment required over a multiyear period?

And then, Susan, just as we think about the \$60 billion to \$65 billion CapEx this year, does the composition change much from last year when you talked about servers as the largest part followed by data centers and networking equipment? And how should we think about that mix between like training and inference just following up on Yann's post this week? Thanks.

Mark Zuckerberg: I can start on the DeepSeek question. I think there's a number of novel things that they did that I think we're still digesting. And there are a number of things that they have advances that we will hope to implement in our systems. And that's part of the nature of how this works, whether it's a Chinese competitor or not.

I kind of expect that every new company that has an advance -- that has a launch is going to have some new advances that the rest of the field learns from. And that's sort of how the technology industry goes.

I don't know -- it's probably too early to really have a strong opinion on what this means for the trajectory around infrastructure and CapEx and things like that. There are a bunch of trends that are happening here all at once.

There's already sort of a debate around how much of the compute infrastructure that we're using is going to go towards pretraining versus as you get more of these reasoning time models or reasoning models where you get more of the intelligence by putting more of the compute into inference, whether just it will mix shift how we use our compute infrastructure towards that.

That was already something that I think a lot of the -- the other labs and ourselves were starting to think more about and already seemed pretty likely even before this, that -- like of all the compute that we're using, that the largest pieces aren't necessarily going to go towards pre-training.

But that doesn't mean that you need less compute because one of the new properties that's emerged is the ability to apply more compute at inference time in order to generate a higher level of intelligence and a higher quality of service, which means that as a company that has a strong business model to support this, I think that's generally an advantage that we're now going to be able to provide a higher quality of service than others who don't necessarily have the business model to support it on a sustainable basis.

The other thing is just that when we're building things like Meta AI, but also how we're implementing AI into all the feeds and ad products and things like that, we're just serving billions of people, which is different from, okay you start to pretrain a model, and that model is sort of agnostic to how many people are using it.

Like at some level, it's going to be expensive for us to serve all of these people because we are serving a lot of people. And so I'm not sure what the kind of net effect of all of this is. The field continues to move quickly. There's a lot to learn from releases from basically everyone who does something interesting, not just the ones over the last month.

We'll continue to kind of incorporate that into what we do as well as making novel contributions to the field ourselves. And I continue to think that investing very heavily in CapEx and infra is going to be a strategic advantage over time.

It's possible that we'll learn otherwise at some point, but I just think it's way too early to call that. And at this point, I would bet that the ability to build out that kind of infrastructure is going to be a major advantage for both the quality of the service and being able to serve the scale that we want to.

Susan Li: I'm happy to add a little more color about our 2025 CapEx plans to your second question.

So we certainly expect that 2025 CapEx is going to grow across all three of those components you described. Servers will be the biggest growth driver, that remains the largest portion of our overall CapEx budget.

We expect both growth in AI capacity as we support our gen AI efforts and continue to invest meaningfully in core AI, but we are also expecting growth in

non-AI capacity as we invest in the core business including to support a higher base of engagement and to refresh our existing servers.

On the data center side, we're anticipating higher data center spend in 2025 to be driven by build-outs of our large training clusters and our higher power density data centers that are entering the core construction phase. We're expecting to use that capacity primarily for core AI and non-AI use cases.

On the networking side, we expect networking spend to grow in '25 as we build higher-capacity networks to accommodate the growth in non-AI and core AI-related traffic along with our large Gen AI training clusters.

We're also investing in fiber to handle future cross-region training traffic. And then in terms of the breakdown for core versus Gen AI use cases, we're expecting total infrastructure spend within each of Gen AI, non-AI and core AI to increase in '25 with the majority of our CapEx directed to our core business with some caveat that that is -- that's not easy to measure perfectly as the data centers we're building can support AI or non-AI workloads and the GPU-based servers we procure for gen AI can be repurposed for core AI use cases and so on and so forth.

But overall, I would reiterate what Mark said. We are committed to building leading foundation models and applications. We expect that we're going to make big investments to support our training and inference objectives, and we don't know exactly where we are in the cycle of that yet.

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

Ronald Josey: Hey, thanks for taking the question. Mark, I want to get back to your comment on getting back to the OG Facebook, and I want to understand a little bit more on the use cases and how that could expand. Video is clearly a benefit. Local, Marketplace, Groups have all been positive.

So any insights on the OG Facebook? And then back to Meta AI, given the adoption we're seeing on the 600-plus MAUs, just how has the user experience evolved to? What are people doing with Meta AI? Thank you.

Mark Zuckerberg: Okay. So for Facebook. A lot of people use Facebook every day, and it's an important part of their lives. And I think that there are a lot of opportunities to make it way more culturally influential than it is today. And I think that that's sort of a fun and interesting goal that will take our product development in some interesting directions that we maybe haven't had a focus on it as much over the last several years.

So I don't know that I have anything much more specific on this other than that this is going to be one of my focus areas for this year. I mean I think it's an investment area and something I'm going to spend some time on.

It might mean that in the near term, we make some trade-offs to kind of focus on some product areas of what we're doing ahead of just kind of maximizing business results in the near term on it.

But overall, I'm really excited about doing some exciting stuff here. And I'm not going to get into many specifics now but we'll get -- we'll follow up on this over the next, I don't know call it, half year or a year as we start rolling stuff out and I think some of this will kind of get back to how Facebook was originally used back in the day. So I think it will be fun.

Susan Li: I'm happy to share a little bit more about Meta AI and what people are doing with it. We are in a phase where we are really learning a lot from the way that people engage with Meta AI.

So from an app perspective, WhatsApp continues to see the strongest Meta AI usage across our Family of Apps. People there are using it most frequently for information seeking and educational queries along with emotional support use cases. Most of the WhatsApp engagement is in one-on-one threads, though we see some usage in group messaging.

And on Facebook, which is the second largest driver of Meta AI engagement, we're seeing strong engagement from our feed deep dives integration that lets people ask Meta AI questions about the content that is recommended to them.

So across, I would say, all query types, we continue to see signs that Meta AI is helping people leverage our apps for new use cases.

We talked about information gathering, social interaction and communication. Lots of people use it for humor and casual conversation. They use it for writing and editing, research, recommendations. And as we look forward to 2025 in our Meta AI roadmap, we are really focused on doing more to make it feel more personalized.

So I would say some of the most exciting features we're working on including improving sort of the memory dimension of the Meta AI experience, where it will be able to remember certain details that people share in one-on-one chats, for example, and use those details to personalize its responses and then really increasing its ability to deliver great content recommendations and enhance really what makes Facebook and Instagram so valuable for people today.

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

Kenneth Gawrelski: Thank you very much. Two for me, please. First, could you talk a little bit -- I know you talked a little bit on the capital intensity side and the recent developments, and it's hard to see, it's hard to tell yet where things are going.

But maybe you could just talk a little bit more near term, '25, the CapEx budget you laid out or the CapEx forecast. Could you talk a little bit about the



constraints you're seeing or where you're seeing constraints, either internally resources planning or externally and any one -- any parts of the ecosystem?

And then on the second one, I'm curious, as you think about the -- as you think about your needs for hiring and we just think about -- we know you gave the OpEx guide for this year.

But as we think about future needs for hiring, could you just give us a sense of how we should think about that? You announced the performance-related reductions earlier this -- for early this year. Could you just talk about how we should be thinking about that '26, '27 and beyond? Thank you.

Susan Li:

Sure. I'm happy to take both of those. So on your first question on just where do we see constraints in our ability to execute against our CapEx plans, obviously we are staying on top of supply availability. That is certainly one of the factors that will influence our CapEx spend in 2025, but we don't really have any updates to share on supply availability right now.

We are planning to significantly ramp up deployment of GPUs in 2025, and we'll continue to engage with our vendors and invest in our own silicon to meet those needs.

When you asked how to think about capital intensity, we're not really -- as both Mark and I alluded to in our prior comments, I think it is really too early to determine what long-run capital intensity is going to look like. There are so many different factors.

The pace of advancement in underlying models, how efficient can they be, what is the adoption and use case of our Gen AI products, what performance gains come from next-generation hardware innovations, both our own and third party and then ultimately, what monetization or other efficiency gains our AI investments unlock.

So again, I think we are -- we're sort of early in the journey here, and we don't have -- I would say we don't have kind of anything to share about long-run capital intensity yet. Your second question was about thinking about hiring needs.

So it's a good segue after infrastructure, employee compensation is the next largest driver of expense growth in 2025. And here, growth in employee comp and headcount more broadly is primarily driven by those areas that I mentioned, infrastructure monetization, generative AI, Reality Labs, and regulation and compliance.

And those generally are more technical organizations. That means that it is a higher cost base relative to business functions where we are also expecting to keep headcount growth constrained. And I would say we are -- we're focused on running the company efficiently.

But at the same time it is -- we feel like we're in a critical period in terms of making sure that we are investing to win, and we want to make sure that we staff those priority areas in a way that really positions us to best do that.

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

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

Ross Sandler: Yeah. One for Mark, on agents. So we all saw OpenAI's operator demo last week. So Mark, as the industry moves from chat to agentic behavior and more commercial intent moves into these AI products, I guess how are you thinking about monetization potential for Meta AI? And then how might Llama 4 reasoning help drive some of these new agentic experiences for Meta AI? Thank you.

Mark Zuckerberg: Yeah. So I guess a couple of things that I'd say on this. One is when you're thinking about agents and reasoning, a lot of this is about being able to perform multistep tasks. So right now the way that a lot of these systems work is you kind of say something and then it responds and it's almost chat like.

But I think that the direction that it's going is you're going to be able to give it an intent or a task and it's going to be able to go off and use sort of an arbitrary amount of compute as much as you want to use on it to be able to do a task. Some of the tasks might be pretty simple for people, go buy a specific thing.

Some of them might be really hard, like go write an app or optimize this code and like really make it as good as possible. And that type of thing, I think, is just going to start becoming more and more prevalent over the next year or two. So I think it's very exciting.

It sort of will feel in some ways like the current products are just getting smarter and others, it will feel like sort of a new form factor because it won't be as much like chat. But it's sort of another generation of these products.

So I think it's just in general, there's a lot to build and be excited about. I guess my note of caution or just my kind of periodic reminder on our product development process, if you will, is we build these products.

We try to scale them to reach usually a billion people or more. And it's at that point once they're at scale that we really start focusing on monetization. So sometimes we'll experiment with monetization before, we're running some experiments with Threads now for example.

But we typically don't really ramp these things up or see them as meaningfully contributing to the business until we reach quite a big scale. So the thing that I think is going to be meaningful this year is the kind of getting of the AI products to scale.

Last year was sort of the introduction and starting to get it to be used. This year my kind of expectation and hope is that we will be at a sufficient scale and have sufficient kind of flywheel of people using it and improvement from that that this will have a durable advantage.

But that doesn't mean that it's going to be a major contributor to the business this year. This year, the improvements to the business are going to be taking the AI methods and applying them to advertising and recommendations and feeds and things like that.

So the actual business opportunity for Meta AI and AI Studio and business agents and people interacting with these AIs remains outside of '25 for the most part. And I think that's an important thing for us to communicate and for people to internalize as you're thinking about our prospects here.

But nonetheless, we've run a process like this many times. We built a product. We make it good. We scale it to be large. We build out the business around it. That's what we do. I'm very optimistic. But it's going to take some time.

Kenneth Dorell: Thank you, everyone, 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.