Meta Platforms, Inc. (META)
Second Quarter 2024 Results Conference Call
July 31st, 2024

Kenneth Dorell, Director, Investor Relations

Thank you. Good afternoon and welcome to Meta Platforms second 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.

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

Mark Zuckerberg, CEO

Thanks Ken. And hey everyone -- thanks for joining today.

This was a strong quarter for our community and business. We estimate that there are now more than 3.2 billion people using at least one of our apps each day. The growth we're seeing here in the US has especially been a bright spot. WhatsApp now serves more than 100 million monthly actives in the US, and we're seeing good year-over-year growth across Facebook, Instagram, and Threads as well -- both in the US and globally.

I'm particularly pleased with the progress that we're making with young adults on Facebook. The numbers we're seeing, especially in the US, really go against the public narrative around who's using the app. A couple of years ago we started focusing our apps more on 18-29 year olds, and it's good to see that those efforts are driving good results.

Another bright spot is Threads, which is about to hit 200 million monthly actives. We're making steady progress towards building what looks like it's going to be another major social app. We're seeing deeper engagement and I'm quite pleased with the trajectory here.

The big theme right now is of course AI, and I'll focus my comments today on three areas: what AI means for our family of apps and core business, what new AI experiences and opportunities we see, and how AI is shaping our metaverse work.

So, let's start: across Facebook and Instagram, advances in AI continue to improve the quality of recommendations and drive engagement. We keep finding that as we develop more general recommendation models, content recommendations get better. This quarter we rolled out our full-screen video player and unified video recommendation service across Facebook -- bringing Reels, longer videos, and Live into a single experience. This has allowed us to extend our unified AI systems, which had already increased engagement on Facebook Reels more than our initial move from CPUs to GPUs did. Over time, I'd like to see us move towards a single, unified recommendation system that powers all of the content including things like People You May Know across all of our surfaces. We're not there, so there's still upside -- and we're making good progress here.

AI is also going to significantly evolve our services for advertisers in some exciting ways. It used to be that advertisers came to us with a specific audience they wanted to reach -- like a certain age group, geography, or interests. Eventually we got to the point where our ads system could better predict who would be interested than the advertisers could themselves. But today advertisers still need to develop creative themselves. In the coming years, AI will be able to generate creative for advertisers as well -- and will also be able to personalize it as people see it. Over the long term, advertisers will basically just be able to tell us a business objective and a budget, and we're going to go do the rest for them. We're going to get there incrementally over time, but I think this is going to be a very big deal.

Moving onto some of the brand new experiences that AI enables, last quarter we started broadly rolling out our assistant, Meta AI, and it is on track to achieve our goal of becoming the most used AI assistant by the end of the year. We have an exciting roadmap ahead of things that we want to add, but the bottom line here is that Meta AI feels on track to be an important service and it's improving quickly both in intelligence and features. Some of the use cases are utilitarian, like searching for information or role-playing difficult conversations before you have them with another person, and other uses are more creative, like the new Imagine Yourself feature that lets you create images of yourself doing whatever you want in whatever style you want. And part of the beauty of AI is that it's general, so we're still uncovering the wide range of use cases that it's valuable for.

An important part of our vision is that we're not just creating a single AI, but enabling lots of people to create their own AIs. This week we launched AI Studio, which lets anyone create AIs to interact with across our apps. I think that creators are especially going to find this quite valuable. There are millions of creators across our apps -- and these are people who want to engage more with their communities and their communities want to engage more with them -- but there are only so many hours in the day. So now they're going to be able to use AI Studio to create AI agents that can channel them to chat with their community, answer people's questions, create content, and more. So I'm quite excited about this. But this goes beyond creators too. Anyone is going to be able to build their own AIs based on their interests or different topics that they are going to be able to engage with or share with their friends.

Business AIs are the other big piece here. We're still in alpha testing with more and more businesses. The feedback we're getting is positive so far. Over time I think that just like every business has a website, a social media presence, and an email address, in the future I think that every business is also going to have an AI agent that their customers can interact with. Our goal is to make it easy for every small business, and eventually every business, to pull all of their content and catalog into an AI agent that drives sales and saves them money. When this is working at scale, I think that this is going to dramatically accelerate our business messaging revenue.

There are a lot of other new opportunities here that I'm excited about too, but I'll save those for another day when we're ready to roll them out.

The engine that powers all these new experiences is the Llama family of foundation models. This quarter we released Llama 3.1, which includes the first frontier-level open source model, as well as new and industry-leading small and medium-sized models. The 405B model has better cost performance relative to the leading closed models, and because it's open, it is immediately the best choice for fine-tuning and distilling your own custom models of whatever size you need. I think we're going to look back at Llama 3.1 as an inflection point in the industry where open source AI started to become the industry standard, just like Linux is.

I often get asked why I'm so bullish on open source. I wrote a letter along with the Llama 3.1 release explaining why I believe that open source is better for developers, for Meta, and for the world more broadly. My view is that open source will be safer, will enable innovation that improves all of our lives faster, and will also create more shared prosperity.

For Meta's own interests, we're in the business of building the best consumer and advertiser experiences. To do that, we need to have access to the leading technology infrastructure and not get constrained by what competitors will let us do. But these models are ecosystems, they're not just isolated pieces of software that we can develop by ourselves. So if we want the most robust ecosystem of tools, efficiency improvements, silicon optimizations, and other integrations to develop around our models, then we need them to be widely used by developers across the industry. Once we know that we're going to have access to the leading models, then I am confident that we're going to be able to build the best social and advertising experiences.

Part of why I'm so optimistic about this is that we have a long track record of success with open source. We've saved billions of dollars with Open Compute Project by having supply chains standardize on our infra designs. Open sourcing tools like PyTorch and React has led to real benefits for us from all of the industry's contributions. This approach has consistently worked for us, and I expect it will work here too.

Another major area of focus is figuring out the right level of infra capacity to support training more and more advanced models. Llama 3 is already competitive with the most advanced models, and we're already starting to work on Llama 4, which we're aiming to be the most advanced in the industry next year. We're planning for the compute clusters and data we'll need for the next several years. The amount of compute needed to train Llama 4 will likely be almost 10x more than what we used to train Llama 3 -- and future models will continue to grow beyond that. It's hard to predict how this will trend multiple generations out into the future, but at this point I'd rather risk building capacity before it is needed, rather than too late, given the long lead times for spinning up new infra projects. And as we scale these investments, we're of course going to remain committed to operational efficiency across the company.

The last area that I want to discuss is how AI is shaping our metaverse work, which continues to be our other long-term focus. Last quarter I discussed how advances in AI have pulled in the timelines for some of our products. A few years ago I would have predicted that holographic AR would be possible before smart AI, but now it looks like those technologies will actually be ready in the opposite order. We're well-positioned for that because of the Reality Labs investments that we've already made.

Ray-Ban Meta glasses continue to be a bigger hit sooner than we expected -- thanks in part to AI. Demand is still outpacing our ability to build them, but I'm hopeful that we'll be able to meet that demand soon. EssilorLuxottica has been a great partner to work with on this, and we're excited to team up with them to build future generations of AI glasses as we continue to build our long term partnership.

Quest 3 sales are also outpacing our expectations -- and I think that's because it's not just the best MR headset for the price, but it's the best headset on the market, period. In addition to gaming, people are increasingly taking advantage of Quest's capabilities as a general computing platform, spending time watching videos, browsing websites, extending their PC via virtual desktop, and more. Horizon also continues to grow across VR, mobile, and desktop -- and I expect that it will become an increasingly important part of that ecosystem as well.

We're hosting our annual Connect conference on September 25, and we'll have lots of exciting updates around all of our AI and metaverse work -- so I encourage you to tune in to that.

At the end of the day, we are in the fortunate position where the strong results we're seeing in our core products and business gives us the opportunity to make deep investments for the future -- and I plan to fully seize that opportunity to build some amazing things that will pay off for our community and our investors for decades to come. The progress we're making on both the foundational technology and product experiences suggests that we're on the right track. I'm proud of what our team has accomplished so far and I'm optimistic about our ability to execute on the opportunities ahead. As always, thank you to our teams who are pushing all of this important work forward, and thanks to all of you for being on this journey with us. 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.

Q2 total revenue was $39.1 billion, up 22% or 23% on a constant currency basis.

Q2 total expenses were $24.2 billion, up 7% compared to last year.

In terms of the specific line items:

Cost of revenue increased 23%, driven primarily by higher infrastructure and Reality Labs inventory costs.

R&D increased 13%, primarily driven by higher headcount-related expenses and infrastructure costs, which were partially offset by lower restructuring costs.

Marketing & Sales decreased 14%, due mainly to lower restructuring and headcount-related costs.

G&A decreased 12%, mostly due to lower legal-related expenses.

We ended the first quarter with almost 70,800 employees, up 2% from Q1.

Second quarter operating income was $14.8 billion, representing a 38% operating margin.

Our tax rate for the quarter was 11%.

Net income was $13.5 billion or $5.16 per share.

Capital expenditures, including principal payments on finance leases, were $8.5 billion, driven by investments in servers, data centers and network infrastructure.

Free cash flow was $10.9 billion. We repurchased $6.3 billion of our Class A common stock and paid $1.3 billion in dividends to shareholders, ending the quarter with $58.1 billion in cash and marketable securities and $18.4 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 approximately 3.27 billion people using at least one of our Family of Apps on a daily basis in June.

Q2 Total Family of Apps revenue was $38.7 billion, up 22% year over year.

Q2 Family of Apps ad revenue was $38.3 billion, up 22% or 23% on a constant currency basis.

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

On a user geography basis, ad revenue growth was strongest in Rest of World and Europe at 33% and 26%, respectively. Asia-Pacific grew 20% and North America grew 17%. On an advertiser geography basis, total revenue growth continued to be strongest in Asia-Pacific at 28%, though growth was below the first quarter rate of 41% as we lapped a period of stronger demand from China-based advertisers.

In Q2, the total number of ad impressions served across our services and the average price per ad both increased 10%. 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 $389 million, up 73%, 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 Q2, Family of Apps expenses were $19.4 billion, representing approximately 80% of our overall expenses. Family of Apps expenses were up 4%, mostly due to higher infrastructure and headcount-related expenses, which were partially offset by lower restructuring costs.

Family of Apps operating income was $19.3 billion, representing a 50% operating margin.

Within our Reality Labs segment, Q2 revenue was $353 million, up 28% driven primarily by Quest headset sales.

Reality Labs expenses were $4.8 billion, up 21% year-over-year driven mainly by higher headcount-related expenses and Reality Labs inventory costs.

Reality Labs operating loss was $4.5 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.

To deliver engaging experiences, we remain focused on executing our priorities, including video and in-feed recommendations.

On Instagram, Reels engagement continues to grow as we make ongoing enhancements to our recommendations systems. Part of this work has been focused on increasing the share of original posts within recommendations so people can discover the best of Instagram, including content from emerging creators. Now, more than half of recommendations in the US come from original posts.

On Facebook, we are seeing encouraging early results from the global roll-out of our unified video player and ranking systems in June. This initiative allows us to bring all video types on Facebook into one viewing experience, which we expect will unlock additional growth opportunities for short-form video as we increasingly mix shorter videos into the overall base of Facebook video engagement. We expect the relevance of video recommendations will continue to increase as we benefit from unifying video ranking across Facebook and integrating our next generation recommendation systems. These have already shown promising gains since we began using the new systems to support Facebook Reels recommendations last year. We expect to expand these new systems to support more surfaces beyond Facebook video over the course of this year and next year.

We're also seeing good momentum with our longer-term engagement priorities, including Generative AI and Threads.

People have used Meta AI for billions of queries since we first introduced it. We're seeing particularly promising signs on WhatsApp in terms of retention and engagement, which has coincided with India becoming our largest market for Meta AI usage. You can now use Meta AI in over 20 countries and eight languages, and in the US we're rolling out new features like Imagine edit, which allows people to edit images they generate with Meta AI.

Beyond generative AI, the Threads community also continues to grow and deepen their engagement as we ship new features and enhance our content recommendation systems.

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, including within Facebook as the mix of overall video engagement shifts more to shorter videos over time, which

creates more ad insertion opportunities. More broadly, we are continuing to get better at determining the best ads to show and when to show them during a person's session across both Facebook and Instagram. This is enabling us to drive revenue growth and conversions without increasing the number of ads or in some cases even reducing ad load.

The second part of improving monetization efficiency is enhancing marketing performance. We continue to be pleased with progress here, with AI playing an increasingly central role.

We're improving ad delivery by adopting more sophisticated modeling techniques made possible by AI advancements, including our Meta Lattice ad ranking architecture, which continued to provide ad performance and efficiency gains in the second quarter.

We're also making it easier for advertisers to maximize ad performance and automate more of their campaign set up with our Advantage+ suite of solutions. We're seeing these tools continue to unlock performance gains, with a study conducted this year demonstrating 22% higher return on ad spend for US advertisers after they adopted Advantage+ Shopping campaigns. Advertiser adoption of these tools continues to expand, and we're adding new capabilities to make them even more useful.

For example, this quarter we introduced Flexible Format to Advantage+ Shopping, which allows advertisers to upload multiple images and videos in a single ad that we can select from and automatically determine which format to serve, in order to yield the best performance.

We have also now expanded the list of conversions that businesses can optimize for using Advantage+ Shopping to include an additional 10 conversion types, including objectives like "add to cart".

Looking forward, we believe generative AI will play a growing role in how businesses market and engage with customers at scale.

We expect this technology will continue to make it easier for businesses to develop customized and diverse ad creative. We've seen promising early results since introducing our first generative AI ad features – image expansion, background generation, and text generation - with more than one million advertisers using at least one of these solutions in the past month. In May, we began rolling out full image generation capabilities into Advantage+ creative, and we're already seeing improved performance from advertisers using the tool.

Finally, we expect AI will help businesses communicate with customers more efficiently through messaging. We're starting by testing the ability for businesses to use AI in their chats with customers to help sell their goods and services and to generate leads. While we are in the early stages, we continue to expand the number of advertisers we're testing with and have seen good advances in the quality of responses since we began using Llama 3.

Next, I would like to discuss our approach to capital allocation, which remains unchanged. We continue to invest both in enhancing our core experiences in the near-term and developing technologies that we believe will transform how people engage with our services in the years ahead. We expect that having sufficient compute capacity will be central to many of these opportunities, so we are investing meaningfully in infrastructure to support our core AI work, in content ranking and ads, as well as our generative AI and advanced research efforts.

Our ongoing investment in core AI capacity is informed by the strong returns we've seen, and expect to deliver in the future, as we advance the relevance of recommended content and ads on our platforms. While we expect the returns from generative AI to come in over a longer period of time, we are mapping these investments against the significant monetization opportunities that we expect to be unlocked across customized ad creative, business messaging, a leading AI assistant, and organic content generation.

As we scale generative AI training capacity to advance our foundation models, we will continue to build our infrastructure in a way that provides us with flexibility in how we use it over time. This will allow us to direct training capacity to gen AI inference, or to our core ranking and recommendation work when we expect that doing so would be more valuable. We will also continue our focus on improving the cost efficiency of our workloads over time.

Reality Labs remains our other long-term initiative that we continue to invest meaningfully in. Quest 3 is selling well and Ray-Ban Meta smart glasses are showing very promising traction, with the early signals that we are seeing across demand, usage, and retention increasing our confidence in the long-run potential of AR glasses.

Finally, as we pursue these investments across near- and long-term priorities, we will remain focused on operating the business efficiently.

Turning now to the revenue outlook.

We expect third quarter 2024 total revenue to be in the range of $38.5-41 billion. Our guidance assumes foreign currency is a 2% headwind 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 $96-99 billion, unchanged from our prior outlook. 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.

While we do not intend to provide any quantitative guidance for 2025 until the fourth quarter call, we expect infrastructure costs will be a significant driver of expense growth next year as we recognize depreciation and operating costs associated with our expanded infrastructure footprint.

Turning now to the capex outlook. We anticipate our full-year 2024 capital expenditures will be in the range of $37-40 billion, updated from our prior range of $35-40 billion. While we continue to refine our plans for next year, we currently expect significant capex growth in 2025 as we invest to support our AI research and our product development efforts.

On to tax. Absent any changes to our tax landscape, we expect our full-year 2024 tax rate to be in the mid-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, Q2 was another good quarter. We continue to execute well across our business priorities and have exciting opportunities in front of us to deliver more value to the people and businesses using our products around the world.

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

Operator: Thank you. We will now be -- 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. And 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 the line of Brian Nowak from Morgan Stanley. Please go ahead.

Brian Nowak: Great, thanks for taking my question. I have two, one for Mark, one for Susan. Mark, I wanted to sort of go back to some of the new generative AI-enabled use cases for users and advertisers. You talked about Meta AI, Studio, chatbots, diffusion models.

If you could just sort of hone in on one or two of those that you're most excited about, we're seeing good signal that could be a real driver for the business in '25, '26 just so we sort of know where are you most focused on all those opportunities, it'd be helpful. And the second one, Susan, you have a lot of CapEx priorities from building new infrastructure for next-generation models, compute capacity.

Just walk us through again on the CapEx philosophy and any guardrails you have around ensuring you generate a healthy return on invested capital for investors from all the CapEx. Thanks.

Mark Zuckerberg: I can take the first one. So I think the things that will drive the most results in '25 and '26 are actually the first category of things that I talked about in my comments, which are the ways that AI is shaping the existing products.

So the ways that it's improving recommendations and helping people find better content as well as making the advertising experiences more effective. I think there's a lot of upside there. Those are already products that are at scale. The AI work that we're doing is going to improve that.

It will improve the experience and the business results. The other areas that we're working on, I mean I think you all know this from following our business for a while, but we have a relatively long business cycle of starting a new product, scaling it to something that reaches a billion people or more and only then really focusing on monetizing at scale.

So realistically, for things like Meta AI or AI Studio, I mean these are things that I think will increase engagement in our products and have other benefits that will improve the business and engagement in the near term.

But before we're really talking about monetization of any of those things by themselves, I mean I don't think that anyone should be surprised that I would expect that that will be years, right?

It's just -- I think that that's like what we've seen with Reels. It's what we saw with all these things.

But I think for those who have followed our business for a long time, you can also get a pretty good sense of when things are going to work years in advance. And I think that the people who bet on those early indicators tend to do pretty well which is why I wanted to share in my comments the early indicator that we had on Meta AI, which is, I mean look, it's early. Last quarter, we -- I think it just started rolling it out a week or two before our earnings call.

This time, we're a few months later. And what we can say is I think we are on track to achieve our goal of being the most used AI assistant by the end of this year. And I think that's a pretty big deal.

Is that the only thing we want to do? No. I mean we obviously want to kind of grow that and grow the engagement on that to be a lot deeper, and then we'll focus on monetizing it over time.

But the early signals on this are good, and I think that that's kind of all that we could reasonably have insight into at this point. But I do think that part of what's so fundamental about AI is it's going to end up affecting almost every product that we have in some way.

It will improve the existing ones and will make a whole lot of new ones possible. So it's why there are all the jokes about how all the tech CEOs get on these earnings calls and just talk about AI the whole time. It's because it's actually super exciting, and it's going to change all these different things over multiple time horizons.

Susan Li: And Brian, I can take the second question. On the ROI part of your question, I would broadly characterize our AI investments into two buckets, core AI and gen AI. And the two are really at different stages as it relates to driving revenue for our businesses and our ability to measure returns.

On our core AI work, we continue to take a very ROI-based approach to our investment here. We're still seeing strong returns as improvements to both engagement and ad performance have translated into revenue gains, and it makes sense for us to continue investing here.

Gen AI is where we're much earlier, as Mark just mentioned in his comments. We don't expect our gen AI products to be a meaningful driver of revenue in '24.

But we do expect that they're going to open up new revenue opportunities over time that will enable us to generate a solid return off of our investment while we're also open sourcing subsequent generations of Llama.

And we've talked about the four primary areas that we're focused here on the gen AI opportunities -- to enhance the core ads business, to help us grow in business messaging, the opportunities around Meta AI, and the opportunities to grow core engagement over time. The other thing I would say is, we're continuing to build our AI infrastructure with fungibility in mind so that we can flex capacity where we think it will be put to best use.

The infrastructure that we build for gen AI training can also be used for gen AI inference. We can also use it for ranking and recommendations by making certain modifications like adding general compute and storage. And we're also employing a strategy of staging our data center sites at various phases of development, which allows us to flex up to meet more demand and less lead time if needed while limiting how much spend we're committing to in the outer years.

So while we do expect that we are going to grow CapEx significantly in 2025, we feel like we have a good framework in place in terms of thinking about where the opportunities are and making sure that we have the flexibility to deploy it as makes the most sense.

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

Mark Shmulik: Yes, hi thanks for taking the question. Just as we look at the revenue guidance and the outlook, Susan, any color you can share on just kind of the state of the overall digital ad market? And you've highlighted some areas where you're seeing strength versus kind of some of the idiosyncratic efforts you've made to kind of improve the efficacy of the ad product. Thank you.

Susan Li: Hi, Mark. We're continuing to see healthy global advertising demand, and we're also delivering ongoing ad performance improvements just related to all of the investments that we've continued to make over time, and improving the sort of ads, targeting, ranking, delivery, all of the fundamental infrastructure there. And we expect that all of that will continue to benefit ad spend in Q3.

We do expect year-over-year growth to slow in Q3 as we are lapping strong growth from China-based advertisers as well as strong Reels impression growth from a year ago. And we also expect modestly larger FX headwinds in Q3 based on current rates.

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. I'll just ask one. You called out building community size and what's happened in the United States as well as Threads. How are you thinking about those newer, faster-growing elements of either Messaging or Threads as a platform and the mix between the potential for engagement growth and overall monetization longer term of either the messaging layer or Threads and what you're most excited about there to build

to sort of capitalize on scale but bring it back towards monetization? Thank you.

Mark Zuckerberg: I can start, and Susan can jump in if she has anything else that she wants to add. So the WhatsApp stat, I think, is really important as a business trend just because the United States punches above its weight in terms of, it's such a large percent of our revenue.

So before, WhatsApp was sort of the leading messaging app in many countries around the world but not in the U.S. And I think now that we're starting to make inroads into leading in the U.S. as more and more people use the product and realize that, hey, this is a really good experience, it's the best experience for cross-platform communication and groups and on all these different things, I think that that's going to just mean that all of the work that we're doing to grow the business opportunity there over time is just going to have a big tailwind if the U.S. ends up being a big market. So I think that that's one reason why it's really relevant.

It's obviously also personally somewhat gratifying to see all the people around us starting to use WhatsApp, so I think that's pretty fun but maybe somewhat less relevant from a business perspective. Threads, I think it's another example of something that it got off to about as good of a start of any app that I can think of. I think it was the fastest-growing app to 100 million people.

And it's a good reminder that even when you have that start, the path from there to 1 billion people using an app is still multiple years. And that's our product cycle. And I think that that's something that is a little bit different about Meta in the way we build consumer products and then business around them than a lot of other companies that ship something and start selling it and making revenue from it immediately.

So I think that's something that our investors and folks thinking about analyzing the business, if needed, to always grapple with is all these new products, we ship them and then there's a multiyear time horizon between scaling them and then scaling them into not just consumer experiences but very large businesses.

But the thing that I think is just super exciting about Threads is that we've been building this company for 20 years, and there are just not that many opportunities that come around to grow a billion-person app.

I mean there are, I don't know maybe a dozen of them in the world or something, right? I mean there are certainly more of them outside the company than inside the company, but we do pretty well and being able to add another one to the portfolio if we execute really well on this is just really exciting to have that potential.

Now obviously there's a ton of work between now and there. I mean we're almost at 200 million. So it's a really good milestone, I'm excited about that. A lot of work between this and it being a large part of the business.

	But I do think that these kind of opportunities are pretty rare and that's something that we're just really excited about. I think the team is doing great work on it.
Susan Li:	Eric, I would just add to that in terms of nearer-term sources of impression growth, we really expect that video is going to remain a source of impression growth for us in the second half.
	On Instagram, we expect Reels to continue to drive growth, while on Facebook, we expect to grow overall video time while increasing the mix of short-form video, which creates more impression growth opportunities. And generally, we expect continued community growth across our apps.
Operator:	Your next question comes from the line of Doug Anmuth with J.P. Morgan. Please go ahead.
Douglas Anmuth:	Hi, thanks for taking the questions. One for Mark, one for Susan. Mark, just in terms of infrastructure and CapEx, you've talked about currently building out not just for Llama 3 and 4 but really out to 7 perhaps and then Llama 4, 10x the compute required versus Llama 3. Just given how much you're building ahead, how does that influence the shape of the CapEx curve over a multiyear period? And then Susan, if you could talk a little bit more about the 3Q outlook.
	I know you're talking about tougher comps, but at the same time, it really suggests only a point of FX neutral decel at the high end. So just curious if there's anything else you can point to more specifically that's driving the expected strength here? Thanks.
Susan Li:	Thanks, Doug. I can go ahead and talk about both of those. So your first question was sort of about the longer-term CapEx outlook. We haven't really shared an outlook sort of on the longer-term CapEx trajectory.
	In part, infrastructure is an extraordinarily dynamic planning area for us right now. We're continuing to work through what the scope of the gen AI roadmaps will look like over that time.
	Our expectation, obviously again, is that we are going to significantly increase our investments in AI infrastructure next year, and we'll give further guidance as appropriate.
	But we are building all of that CapEx, again with the factors in mind that I talked about previously thinking about both how to build it flexibly so we can deploy to core AI and gen AI use cases as needed and making sure that we both feel good about the returns that we're seeing on the core AI investments, which we're able to measure more immediately.
	And then we feel good about the opportunities in the gen AI efforts. Your second question was about the Q3 revenue outlook. Again, I mentioned this earlier. We have seen healthy global advertising demand on our platform.

We're delivering ongoing ad performance improvements, which again, we feel like is the result of many, many quarters of effort that have accrued and will continue to accrue value to our platform. And we saw basically in Q2 where revenue grew 22%, that there was broad-based strength across regions and verticals including particular strength among smaller advertisers, and we expect that generally to continue into Q3.

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

Justin Post: Great, thank you. I just want to get back to the comment on U.S. young adult user growth, especially maybe on Facebook and Instagram. I know you made a big change with Reels a couple of years ago. But what are those users doing on Facebook and Instagram? And can you give us any quantification of the usage growth?

Susan Li: Thanks, Justin. So building products with young adults in mind has been a core priority area for the Facebook team in recent years, and we've been very encouraged to see these efforts translate into engagement growth with this cohort.

We've seen healthy growth in young adult app usage in the U.S. and Canada for the past several quarters. And we've seen that products like Groups and Marketplace have seen particular traction with young adults.

Posting to groups in the U.S. and Canada has been growing. That's been boosted mainly by young adults. And we also see that they are active users of Marketplace, which has benefited from product improvements and strong demand for secondhand products in the U.S.

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

Mark Mahaney: Yeah, I was going to ask about Marketplace so that's a nice segue. It's a great, somewhat undermonetized or arguably very under-monetized asset. I know you indirectly monetize it and it's a very large marketplace.

It may even be bigger than eBay. Your thoughts on what you may want to do in the future in terms of monetizing it, in part, maybe even improve the quality of the Marketplace. And then secondly, I just want to ask you about headcount.

It's down about 1% year-over-year. You're pretty much back at par with where the employee headcount was prior to significant reduction. How should we think about headcount growth going forward? You talked about a significant growth in CapEx.

Should we expect a moderate growth in headcount? Significant? Any thoughts on that would be helpful.

Susan Li: Thanks, Mark. On your first question about Marketplace, again, we're obviously excited that it's been one of the drivers of strength in young adults.

I would probably just say that more generally, Marketplace is one prong in a broader commerce strategy that we have, which continues to be focused on basically creating the best shopping experience on our platform. Marketplace is obviously consumer-oriented. The broader part of the commerce strategy is about making it easier for businesses to advertise their products, for buyers to find and purchase relevant items on our platform.

And to that end, I would say that we feel quite happy with also the investments we've been making in Shops ads. Shops ads revenue is growing at a strong year-over-year pace.

We're seeing Shops ads drive incremental performance for advertisers, and it's also working well in combination with some of our other products like Advantage+ Shopping. Your second question was about headcount.

We continue to be disciplined about where we're allocating new headcount to ensure that it's really focused on our core company priorities, but we're also working down a prior hiring underrun.

And as we further close that hiring underrun over the course of this year, I do expect that we will end 2024 with in-seat reported headcount that is meaningfully higher than where we ended 2023.

We aren't providing sort of 2025 headcount growth expectations yet as we haven't started our budgeting process yet. But again, I expect that we will primarily target our hiring to focus on priority areas, and we will be running a very disciplined headcount process.

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

Youssef Squali: Great, thank you very much. So the AI assistant using Llama 3.1 has been incorporated in different variations and looks really impressive and seems to be getting closer to becoming a full search engine for virtually everything except for commercial queries so far.

So are there any plans to open it up to the broader web? Kind of like what maybe OpenAI is alpha testing, maybe link it to third-party marketplaces for commercial search, et cetera. And then on Ray-Ban, can you maybe talk a little bit more about the opportunity to deepen your relationship with EssilorLuxottica? What would that look like? What kind of areas are the most exciting to you, Mark, in that relationship? Thank you.

Mark Zuckerberg: Yes. I'm very excited with how Llama 3.1 landed. I think the team there is doing really great work going from the first version of Llama, to Llama 2 last year that was a generation behind the frontier and now Llama 3.1, which is basically competitive and, in some ways, leading the other top closed models. Meta AI

uses a version of Llama 3.1 as well as a bunch of other services that we've built to kind of build a cohesive product.

And when I was talking before about we have the initial usage trends around Meta AI but there's a lot more that we want to add, things like commerce and you can just go vertical by vertical and build out specific functionality to make it useful in all these different areas are eventually, I think, what we're going to need to do to make this just as -- to fulfill the potential around just being the ideal AI assistant for people.

So it's a long roadmap. I don't think that this stuff is going to get finished in the next couple of quarters or anything like that. But this is part of what's going to happen over the next few years as we build something that will, I think, just be a very widely used service.

So I'm quite excited about that. And we're going to continue working on Llama, too. So I mean you mentioned Llama, and I think the question was a little more about Meta AI but they're both -- I mean they're related. Llama is sort of like the engine that powers the product and it's open source, and I'm just excited about the progress that we're making on both of those.

On the smart glasses, EssilorLuxottica is a great partner. We are now in the second generation of the Ray-Ban Meta glasses. They're doing well, better than I think we had expected, and we expected them to grow meaningfully from the first generation so that's been a very positive surprise.

And I think part of that is that it's just well positioned to dovetail well with the AI revolution that we're seeing and offering all kinds of new functionality there. So that was great.

But EssilorLuxottica is a great company that has a lot of different products that we hope to be able to partner with to just continue building new generations of the glasses and deepen the AI products and make it better and better.

I think there's a lot more to go from here. And compared to what we thought at this point, it's doing quite well. Compared to, I think, what it needs to be, to be like a really leading piece of consumer electronics, I think we're still early but all the signs are good.

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

Ronald Josey: Great, thanks for taking the question. I want to get back, Mark, to the commentary on open source and Llama 3. Totally understand that Meta is not offering a public cloud, and so what I wanted to hear from you is maybe a little bit more on the product vision of products to come out of Llama 3.

And meaning potentially offering some of these products to other companies, call it, for customer service or call center offerings or other verticals. And so any

insights on just how you envision maybe the open source and Llama 3.1 can sort of offer greater enterprise services for others to benefit from? Thank you.

Mark Zuckerberg: So Llama is the foundation model that people can shape into all kind of different products.

So whether it's Meta AI for ourselves or the AI Studio or the business agents or like the assistant that's in the Ray-Ban glasses, like all these different things are basically products that have been built with Llama. And similarly, any developer out there is going to be able to take it and build a whole greater diversity of different things as well.

Like I talked about, I think -- the reason why open sourcing this is so valuable for us is that we want to make sure that we have the leading infrastructure to power the consumer and business experiences that we're building. But the infrastructure, it's not just a piece of software that we can build in isolation.

It really is an ecosystem with a lot of different capabilities that other developers and people are adding to the mix, whether that's new tools to distill the models into the size that you want for a custom model, or ways to fine-tune things better or make inference more efficient or all different other kinds of methods that we haven't even thought of yet, the silicon optimizations that the silicon companies are doing, all the stuff.

It's an ecosystem. So we can't do all that ourselves. And if we built Llama and just kind of kept it within our walls, then it wouldn't actually be as valuable for us to build all the products that we're building as it's going to end up being.

So that's the business strategy around that. And that's why we don't feel like we need to necessarily build a cloud and sell it directly in order for it to be a really positive business strategy for us. And part of what we're doing is working closely with AWS, I think, especially did great work for this release.

Other companies like Databricks, NVIDIA, of course, other big players like Microsoft with Azure, and Google Cloud, they're all supporting this. And we want developers to be able to get it anywhere.

I think that that's one of the advantages of an open source model like Llama is -- it's not like you're locked into one cloud that offers that model, whether it's Microsoft with OpenAI or Google with Gemini or whatever it is, you can take this and use it everywhere and we want to encourage that.

So I'm quite excited about that. The enterprise and business applications that we're going to be most focused on, though, in addition to just optimizing the advertiser experience like I talked about in my comments earlier, it is the business agent piece.

I just think that there's a huge potential like I said earlier. I think pretty much every business today, it has an e-mail address. They have a website. They have social media accounts.

I think in the future, they're going to have at least one, if not multiple business agents that can do the whole range of things from interacting to help people buy things to helping support the sales that they've done if they have issues with the product, if they need to get in touch with you for something. And we already see people interacting with businesses over messaging working quite well in countries that have low cost of labor.

But the thing is that in order to have someone answering everyone's questions is quite expensive in a lot of countries. And I think that this is like a thing that AI, I think, is just going to be very well suited towards doing.

And when we can make that easy for the hundreds of millions of businesses that use our platforms to pull in all their information and their catalogs and the history and all the content that they've shared and really just quickly stand up an agent, I think that's going to be awesome.

So we can combine Llama with a lot of custom work that we're doing in our business teams and couple that with all the other investment that the rest of the ecosystem is doing to make Llama good. And I think that's going to be huge, but that's just one area.

I mean this really goes across all of the different products that we're building, both consumer and business. So it's a lot of exciting stuff.

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

Operator: Thank you. That question will come from the line of Ross Sandler with Barclays. Please go ahead.

Ross Sandler: Great, Mark, so on Monday in your interview with Jensen, you said something along the lines of if scaling ended up stopping one day, you'd have five years of product work to do and ahead of you.

So outside of agents or AI assistants, what other areas in AI are you thinking about or looking at in that five-year roadmap? And then the second question is maybe one more stab at the capacity question. You guys said that Llama 3.1 was trained on 16,000 H100s. You've also said that you're going to have 600,000 available by year-end.

So even if we kind of go up to 160,000 GPUs for Llama 4, we have plenty of extra capacity for inference and future products. I guess how are you guys modeling out this entire kind of CapEx roadmap between training, inference, and future things? That's all. Thanks a lot.

Mark Zuckerberg: I can start with the first one, and then I'll let Susan answer the second one. It's an interesting question. It's a little hypothetical because I mean I do think it's true that if, let's say, there were no future foundation models.

I think there would just be a huge amount of product innovation that the industry would bring to bear, and that just takes time.

But then at the same time, there are going to be future foundation models, and they're going to be awesome and unlock new capabilities and we're planning our products around those. So I'm not really planning our product roadmap assuming that there isn't future innovation.

To the contrary, we are planning what's going to be in Llama 4 and Llama 5 and beyond based on what capabilities we think are going to be most important for the roadmap that I just laid out, for having the breadth of utility that you're going to need in something like Meta AI, making it so that businesses and creators and individuals can stand up any kind of AI agents that they want, that you're going to have these kind of real-time, multimodal glasses with you all the time that will just be increasingly useful for all the things that you're doing.

And that -- I guess this kind of dovetails with what I'd expect Susan to talk about next -- we do have this huge set of use cases already about people wanting to discover content and interact with their friends and businesses reaching people, and all that stuff is getting better with this, too.

So I -- there's -- I guess my point there was it's just -- there's some lag between the technology becoming available and the products becoming kind of fully explored in the space.

And I just think that that was kind of my way of saying that I think that this is just a very exciting area where there's just going to be a lot of innovation for a long time to come. I'll let Susan take a stab at the numbers around the GPUs and all that.

Susan Li: Thank you. We are clearly in the process of building out a lot of capacity, and Mark has alluded to that in his comments about what we have needed to train prior generations and the next generation of Llama.

And we are -- and that's driving sort of what we've talked about in terms of the significant growth in CapEx in 2025. And we aren't really in the position now to share a longer-term outlook. When we think about sort of any given new data center project that we're constructing, we think about how we will use it over the life of the data center.

We think about the amount of capacity we would use in terms of training whatever the subsequent generations of Llama are and it's architected around that.

But then we also look at how we might use it several years into its lifetime towards other use cases across our core business, across what we think might be future needs for inference, for generative AI-based products.

So there's sort of a whole host of use cases for the life of any individual data center ranging from gen AI training at its outset to potentially supporting gen

AI inference to being used for core ads and content ranking and recommendation and also thinking through the implications, too, of what kinds of servers we might use to support those different types of use cases.

So we are really mapping across a wide range of potential use cases when we undertake any given project. And we're really doing that with both a long time horizon in mind, again, because of the long lead times in spinning up data centers, but also recognizing that there are multiple decision points in the lifetime of each data center in terms of thinking through when to order servers and what servers to order and what you'll put them towards.

And that gives us flexibility to make the sort of the best decisions based on the information we have in the future.

Kenneth Dorell: Great. Thank you all 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.