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

Ken 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, 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 their content and catalog into an AI agent that drives sales and saves them money. When this is working at scale, I expect it to dramatically accelerate our business messaging revenue.

There are a lot of other new opportunities 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 our lives faster, and will 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 develop by ourselves. 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 leading models, I'm confident we'll 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 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 we'll be able to meet 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 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.