

Meta Platforms, Inc. (META)
First Quarter 2025 Results Conference Call
April 30th, 2025

Kenneth Dorell, Director, Investor Relations

Thank you. Good afternoon and welcome to Meta's first quarter 2025 earnings conference call. Joining me today to discuss our results are Mark Zuckerberg, CEO and Susan Li, CFO.

Our remarks today will include forward-looking statements, which are based on assumptions as of today. Actual results may differ materially as a result of various factors including those set forth in today's earnings press release, and in our annual report on Form 10-K filed with the SEC. We undertake no obligation to update any forward-looking statement.

During this call we will present both GAAP and certain non-GAAP financial measures. A reconciliation of GAAP to non-GAAP measures is included in today's earnings press release. The earnings press release and an accompanying investor presentation are available on our website at investor.atmeta.com.

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

Mark Zuckerberg, CEO

Thanks Ken, thanks everyone for joining today.

We've had a strong start to the year. Our community keeps growing with more than 3.4 billion people now using at least one of our apps each day. Our business is also performing very well -- and I think we're well-positioned to navigate the macro economic uncertainty.

The major theme right now of course is how AI is transforming everything we do. And as we continue to increase our investments and focus more of our resources on AI, I thought it would be useful today to lay out the five major opportunities that we're focused on: improved advertising, more engaging experiences, business messaging, Meta AI, and AI devices. These are each long term investments that are downstream from us building general intelligence and leading AI models and infrastructure. Even with our significant investments, we don't need to succeed in all of these areas to have a good ROI. But if we do, then I think that we will be wildly happy with the investments that we're making.

The first opportunity is improved advertising. Our goal is to make it so that any business can basically tell us what objective they're trying to achieve -- like selling something or getting a new customer -- and how much they're willing to pay for each result, and then we just do the rest. Businesses used to have to generate their own ad creative and define what audiences they wanted to reach. But AI has already made us better at targeting and finding the audiences that will be interested in their product than many businesses are themselves, and that keeps improving. And now AI is generating better creative options for many businesses as well. I think that this is really redefining what advertising is into an AI agent that delivers measurable business results at scale. And if we deliver on this vision, then over the coming years I think that the increased productivity from AI will make advertising a meaningfully larger share of global GDP than it is today.

In just the last quarter, we're testing a new ads recommendation model for Reels, which has already increased conversion rates by 5%. And we're seeing 30% more advertisers are using AI creative tools in the last quarter as well.

The second opportunity is more engaging experiences. This will come in two forms: better recommendations for existing content types, and better, new types of content.

In the last six months, improvements to our recommendation systems have led to a 7% increase in time spent on Facebook, 6% increase on Instagram, and 35% on Threads. Threads now also has more than 350 million monthly actives and continues to be on track to become our next major social app.

In addition to better recommendations for existing content types, AI is also enabling the creation of better content as well. Some of this will be helping people produce better content to share themselves. Some of this will be AI generating content directly for people that is personalized for them. Some of this will be in existing formats like photos and videos, and some of this will be increasingly interactive. I've often talked about this long term trend of content becoming richer over time. Our feeds started mostly with text, and then became mostly photos when we all got mobile phones with cameras, and then became mostly video when mobile networks became fast enough to handle that well. We're now in the video era, but I don't think that this is the end of the line. In the near future I think we're going to have content in our feeds that you can interact with and it'll interact back with rather than you just watching it.

Over the long term, as AI unlocks more productivity in the economy, I also expect that people will spend more of their time on entertainment and culture, which will create an even larger opportunity to create more engaging experiences across all of these apps.

The third opportunity is business messaging. Right now the vast majority of our business is advertising in feeds on Facebook and Instagram. But WhatsApp now has more than 3 billion monthly actives, with more than 100 million people in the US and growing quickly there. Messenger is also used by more than a billion people each month, and there are now as many messages sent each day on Instagram as there are on Messenger. Business messaging should be the next pillar of our business.

In countries like Thailand and Vietnam where there is a low cost of labor, we see many businesses conduct commerce through our messaging apps. There is actually so much business through messaging that those countries are both in our top 10 or 11 by revenue even though they're ranked in the 30s in global GDP. This phenomenon hasn't yet spread to developed countries because the cost of labor is too high to make this a profitable model before AI -- but AI should solve this.

In the next few years, I expect that just like every business has an email address, social media account, and website, they'll also have an AI business agent for customer support and sales. And they should be able to set that up very easily given all the context they've already put into our business platforms.

We're going to have more to share on upcoming calls about our progress in this area.

The fourth opportunity is Meta AI. Across our apps, there are now almost a billion monthly actives using Meta AI. Our focus for this year is deepening the experience and making Meta AI the leading personal AI -- with an emphasis on personalization, voice conversations, and entertainment. I think

that we're all going to have an AI that we talk to throughout the day -- while we're browsing content on our phones, and eventually as we're going through our days with glasses -- and I think this will be one of the most important and valuable services that has ever been created.

In addition to building Meta AI into our apps, we just released our first Meta AI standalone app. It's personalized so you can talk to it about interests you've shown while browsing Reels or different content across our apps. And we built a social feed so you can discover entertaining ways that others are using Meta AI, and initial feedback has been good so far.

Over time, I expect the business opportunity for Meta AI to follow our normal product development playbook. First we build and scale a product, and then once it's at scale then we focus on revenue. In this case, I think there will be a large opportunity to show product recommendations or ads, as well as a premium service for people who want to unlock more compute for additional functionality or intelligence. But I expect that we're going to be largely focused on scaling and deepening engagement for at least the next year before we'll really be ready to start building out the business here.

The fifth opportunity is AI devices, which is increasingly how we're thinking about our work on the next generation of computing platforms. Glasses are the ideal form factor for both AI and the metaverse. They enable you to let an AI see what you see, hear what you hear, and talk to you throughout the day. And they let you blend the physical and digital worlds together with holograms. More than a billion people worldwide wear glasses today, and it seems highly likely that these will become AI glasses over the next 5 to 10 years. Building the devices that people use to experience our services lets us deliver the highest quality AI and social experiences. This will serve as an amplifier on all the opportunities I mentioned so far, as well as unlocking some new opportunities as well.

Ray-Ban Meta AI glasses have tripled in sales in the last year and people who have them are using them a lot. We've got some exciting new launches with our partner EssilorLuxottica later this year as well that should expand that category and add some new technological capabilities to the glasses.

On Quest, we're also seeing deeper engagement as Quest 3S makes VR accessible to more people, and more people are creating experiences in Horizon with AI tools.

Everything that I've talked about today is built on top of our AI models and our infrastructure. We released the first Llama 4 models earlier this month -- and they are some of the most intelligent, best multi-modal, lowest latency, and most-efficient models that anyone has built. We have more models on the way, including the massive Llama 4 Behemoth model.

Overall, we are focused on building full general intelligence. All of the opportunities that I've discussed today are downstream of delivering general intelligence and doing so efficiently. The pace of progress across the industry and the opportunities ahead for us are staggering. I want to make sure that we're working aggressively and efficiently, and I also want to make sure that we are building out the leading infrastructure and teams we need to achieve our goals.

To that end, we're accelerating some of our efforts to bring capacity online more quickly this year as well as some longer term projects that will give us the flexibility to add capacity in the coming years as well. And that has increased our planned investment for this year.

More broadly, this has been a good start to what I expect will continue to be an intense year. We've got a lot more exciting work in the pipeline that I'm looking forward to sharing soon. I continue to think that this year is going to be a pivotal moment for our industry. I'm grateful for everyone who is working so hard at the company to deliver all this amazing technology and new experiences. As always, thank 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.

Q1 total revenue was \$42.3 billion, up 16% or 19% on a constant currency basis.

Q1 total expenses were \$24.8 billion, up 9% compared to last year.

In terms of the specific line items:

Cost of revenue increased 14%, driven primarily by higher infrastructure costs and payments to partners, partially offset by a benefit from the previously announced extension of server useful lives.

R&D increased 22%, mostly due to higher employee compensation and infrastructure costs.

Marketing & Sales increased 8%, driven mainly by an increase in professional services related to our ongoing platform integrity efforts.

G&A decreased 34% driven primarily by lower legal-related costs.

We ended Q1 with over 76,800 employees, up 4% quarter-over-quarter.

First quarter operating income was \$17.6 billion, representing a 41% operating margin.

Our tax rate for the quarter was 9%, as we recognized excess tax benefits from share based compensation due to the increase in our share price versus prior periods.

Net income was \$16.6 billion or \$6.43 per share.

Capital expenditures, including principal payments on finance leases, were \$13.7 billion, driven by investments in servers, data centers and network infrastructure.

Free cash flow was \$10.3 billion. We repurchased \$13.4 billion of our Class A common stock and paid \$1.3 billion in dividends to shareholders, ending the quarter with \$70.2 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 3.4 billion people used at least one of our Family of Apps on a daily basis in March.

Q1 Total Family of Apps revenue was \$41.9 billion, up 16% year over year.

Q1 Family of Apps ad revenue was \$41.4 billion, up 16% or 20% on a 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 and North America at 19% and 18%, respectively. Europe and Asia-Pacific grew 14% and 12%.

In Q1, the total number of ad impressions served across our services increased 5% and the average price per ad increased 10%. 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 \$510 million, up 34%, driven mostly by business messaging revenue growth from our WhatsApp Business Platform as well as Meta Verified subscriptions. We continue to direct the majority of our investments toward the development and operation of our Family of Apps. In Q1, Family of Apps expenses were \$20.1 billion, representing 81% of our overall expenses. Family of Apps expenses were up 10%, mainly due to growth in employee compensation and infrastructure costs, which were partially offset by lower legal-related expenses.

Family of Apps operating income was \$21.8 billion, representing a 52% operating margin. Within our Reality Labs segment, Q1 revenue was \$412 million, down 6% year-over-year due to lower Meta Quest sales, which were partially offset by increased sales of Ray-Ban Meta AI glasses.

Reality Labs expenses were \$4.6 billion, up 8% year-over-year driven primarily by higher employee compensation.

Reality Labs operating loss was \$4.2 billion.

Turning now to the business outlook. There are two primary factors that drive our revenue performance: our ability to deliver engaging experiences for our community, and our effectiveness at monetizing that engagement over time.

On the first, we're focused both on enhancing our core Family of Apps today and building the next generation of devices and experiences through Reality Labs. I'll start with our Family of Apps.

In the first quarter, we saw strong growth in video consumption across both Facebook and Instagram, particularly in the US where video time spent grew double digits year-over-year.

This growth continues to be driven primarily by ongoing enhancements to our recommendation systems, and we see opportunities to deliver further gains this year. We're also progressing on longer-term efforts to develop innovative new approaches to recommendations.

A big focus of this work will be on developing increasingly efficient recommendation systems so that we can continue scaling up the complexity and compute used to train our models while avoiding diminishing returns. There are promising techniques we're working on that will incorporate the innovations from LLM model architectures to achieve this.

Another area that is showing early promise is integrating LLM technology into our content recommendations systems. For example, we're finding that LLMs' ability to understand a piece of content more deeply than traditional recommendation systems can help better identify what is interesting to someone about a piece of content, leading to better recommendations. We began testing using Llama in Threads recommendation systems at the end of last year given the app's text-based content, and have already seen a 4% lift in time spent from the first launch. It remains early here, but a big focus this year will be on exploring how we can deploy this for other content types, including photos and videos.

We also expect this to be complimentary to Meta AI as it can provide more relevant responses to people's queries by better understanding their interests and preferences through their interactions across Facebook, Instagram, and Threads. Earlier this year, we began testing the ability for Meta AI to better personalize its responses by remembering certain details from people's prior queries and considering what that person engages with on our apps. We are already seeing this lead to deeper engagement with people we've rolled it out to, and it is now built into Meta AI across Facebook, Instagram, Messenger, and our new standalone Meta AI app in the US and Canada.

We're also continuing to focus on helping people connect over content. In Q1, we launched a new experience on Instagram in the US that consists of a Feed of content your friends have left a note on or liked, and we're seeing good results. We also just launched Blend, which is an opt-in experience in direct messages that enables you to blend your Reels algorithm with your friends to spark conversations over each other's interests. These features all lean into Instagram's position at the intersection of entertainment and social connection. WhatsApp remains at its core a private messaging app, but it has evolved to also become a place people come to get updates from accounts they are connected to or follow. Today, there are tens of billions of views of Status posts on WhatsApp each day, and we continue to invest in the Updates tab as a place people can go to do more.

Creators remain another big focus for us, and we're investing in tools to help them produce the best original content on our platforms. Last week, we launched our standalone Edits app, which supports the full creative process for video creators - from inspiration and creation to performance insights. Edits has an ultra-high resolution short-form video camera and includes generative AI tools that enable people to remove the background of any video or animate still images, with more features coming soon.

Moving to Reality Labs. We're seeing very strong traction with Ray-Ban Meta AI glasses, with over 4x as many monthly actives as a year ago, and the number of people using voice commands is growing even faster as people use it to answer questions and control their glasses. This month, we fully rolled out live translations on Ray-Ban Meta AI glasses to all markets for English, French, Italian, and Spanish. Now, when you are speaking to someone in one of these languages, you'll hear what they say in your preferred language through the glasses in real time.

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 optimize ad supply across each surface to better deliver ads at the time and place they are most relevant to people. We are also starting to introduce ads on unmonetized surfaces, like Threads, which we opened up to all eligible advertisers this month to reach people in over 30 different markets to start, including the US. As we do for any newly monetized surface, we expect to gradually ramp ad supply as we optimize the ad formats and ensure they feel native to the app. We don't expect Threads to be a meaningful driver of overall impression or revenue growth in 2025.

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

We're continuing to improve our ads systems by developing new modeling technologies to more efficiently predict the right ad to show. In Q1, we introduced our new Generative Ads Recommendation model, or GEM, for ads ranking. This model uses a new architecture we developed that is twice as efficient at improving ad performance for a given amount of data and compute. This efficiency gain enabled us to significantly scale up the amount of compute we use for model training, with GEM trained on thousands of GPUs, our largest cluster for ads training to date. We began testing the new model for ads recommendations on Facebook Reels earlier this year and have seen up to a 5% increase in ad conversions. We're now rolling it out to additional surfaces across our apps.

On the ads product side, we're seeing continued momentum with our Advantage+ suite of AI-powered solutions.

We've been encouraged by the initial tests of our streamlined campaign creation flow for Sales, app and lead campaigns, which starts with Advantage+ turned on from the beginning for advertisers. In April, we rolled this out to more advertisers and expect to complete the global roll out later this year.

We're also seeing strong adoption of Advantage+ creative. This week, we are broadening access of Video Expansion to Facebook Reels for all eligible advertisers, enabling them to automatically adjust the aspect ratio of their existing videos by generating new pixels in each frame to optimize their ads for full screen surfaces. We also rolled out image generation to all eligible advertisers and this quarter we plan to continue testing a new virtual try-on feature that uses genAI to place clothing on virtual models, helping customers visualize how an item may look and fit.

Last, we continue to evolve our ads platform to drive results that are optimized for each business' objectives and the way they measure value. One example of this is our Incremental Attribution feature, which enables advertisers to optimize for driving incremental conversions, or conversions we believe would not have occurred without an ad being shown. We're seeing strong results in testing so far, with advertisers using Incremental Attribution in tests seeing an average 46% lift in incremental conversions compared to their business-as-usual approach. We expect to make this available to all advertisers in the coming weeks.

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.

Starting with headcount. Our hiring continues to be targeted at technical roles within our company priorities. In the first quarter, the significant majority of the roughly 2,800 employees we

added were to support our priorities of monetization, infrastructure, generative AI, regulation and compliance, and Reality Labs.

On infrastructure, we have two primary focuses to meet the growing compute needs of our services and AI initiatives.

The first way is by significantly scaling up our infrastructure footprint. Our capex growth this year is going toward both generative AI and core business needs, with the majority of overall capex supporting the core. We expect the significant infrastructure footprint we're building will not only help us meet the demands of our business in the near-term, but also provide us an advantage in the quality and scale of AI services we can deliver. We continue to build this capacity in a way that grants us maximum flexibility in how and when we deploy it to ensure we have the agility to react to how the technology and industry develop in the coming years.

The second way we're meeting our compute needs is by increasing the efficiency of our workloads. In fact, many of the innovations coming out of our ranking work are focused on increasing the efficiency of our systems. This emphasis on efficiency is helping us deliver consistently strong returns from our core AI initiatives.

For example, we shared on the Q3 2024 call that improvements to our AI-driven Feed and video recommendations drove a roughly 8% lift in time spent on Facebook and a 6% lift on Instagram over the first 9 months of last year. Since then, we've been able to deliver similar gains in just six months time, with improvements to our AI recommendations delivering 7% and 6% time spent gains on Facebook and Instagram, respectively.

Before moving to our financial guidance, I want to acknowledge the dynamic macro environment and note that our range reflects the potential for a wider set of outcomes. We continue to feel good about the fundamental drivers of revenue growth and believe the past work we've done to streamline our operations and cost profile puts us in a strong position to navigate a variety of outcomes.

Moving to our financial outlook. We expect second quarter 2025 total revenue to be in the range of \$42.5-45.5 billion. Our guidance assumes foreign currency is an approximately 1% tailwind to year-over-year total revenue growth, based on current exchange rates.

Turning now to the expense outlook. We expect full year 2025 total expenses to be in the range of \$113-118 billion, lowered from our prior outlook of \$114-119 billion.

Turning now to the capex outlook. We anticipate our full year 2025 capital expenditures, including principal payments on finance leases, will be in the range of \$64-72 billion, increased from our prior outlook of \$60-65 billion. This updated outlook reflects additional data center investments to support our AI efforts as well as an increase in the expected cost of infrastructure hardware. The majority of our capex in 2025 will continue to be directed to 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. The European Commission recently announced its decision that our subscription for no

ads model is not compliant with the DMA. Based on feedback from the European Commission in connection with the DMA, we expect we will need to make some modifications to our model, which could result in a materially worse user experience for European users and a significant impact to our European business and revenue as early as the third quarter of 2025. We will appeal the Commission's DMA decision but any modifications to our model may be imposed before or during the appeal process.

In closing, this was another solid quarter for our business. We believe the investments we're making across our company priorities will position us well in the coming years to continue delivering engaging services for our community, compelling results for advertisers, and strong business performance.

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

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

Brian Nowak: Great, thanks for taking my questions. I have two. The first one is on Llama. Mark, can you -- the LLM landscape continues to sort of evolve and be somewhat competitive.

Can you sort of talk us through some of the key areas of advancement you are most focused on and excited about as we sort of think about Behemoth and next versions of Llama to come?

And then the second one on Meta AI, almost 1 billion users globally. Any help on sort of how you're seeing U.S. traction there and the types of recurring user behaviors that you're seeing in the early Meta AI use cases? Thanks.

Mark Zuckerberg: Sure. I can talk about the LLMs. On the Meta AI usage, I'm not sure if we have more stats to share on that now. Yes. It's -- I'll defer to Susan on if there's anything that we're ready on that.

On the LLM, yes, there's a lot of progress being made in a lot of different dimensions. And the reason why we want to build this out is one, is that we think it's important that for kind of how critical this is for our business that we sort of have control of our own destiny and are not depending on another company for something so critical.

But two, we want to make sure that we can shape the development to be optimized for our infrastructure and the use cases that we want.

So to that end, Llama 4, the shape of the model with 17 billion parameters per expert was designed specifically for the infrastructure that we have in order to provide the low latency experience to be voice optimized.

One of the key things, if you're having a voice conversation with AI is it needs to be low latency. That way when you're having a conversation with it, there's isn't a large gap between when you stop speaking and it starts.

So everything from the shape of the model to the research that we're doing to the techniques that go into it are kind of fit into that.

Similarly, another thing that we focused on was context window length. And in some of our models, we have really -- we're industry-leading on context window length.

And part of the reason why we think that that's important is because we're very focused on providing a personalized experience. And there are different ways that you can put personalization context into an LLM, but one of the ways to do it is to include some of that context in the context window. And having a long context window that can incorporate a lot of the background that the person has shared across our apps is one way to do that.

So that's like -- it kind of is giving you a flavor of the products that we're trying to build and then some specific technical architecture decisions and research prioritization that we basically have made in order to deliver the specific experience that we're going for.

I could go on and add a lot more. The reason I think it's also very important to deliver big models like Behemoth, not because we're going to end up serving them in production, but because of the technique of distilling from larger models, right. The Llama 4 models that we've published so far and the ones that we're using internally and some of the ones that we'll build in the future, are basically distilled from the Behemoth model in order to get the 90%, 95% of the intelligence of the large model in a form factor that is much lower latency and much more efficient.

So these things are all very important. Obviously we wouldn't be able to do that kind of distillation from other closed models.

So that kind of gives you a flavor for how we're thinking about the development of this and then, of course, the models and the infrastructure that we're building out power all of the opportunities that I mentioned before.

Susan Li:

Brian, I'm happy to answer your second question about Meta AI. The top use case right now for Meta AI from a query perspective is really around information gathering as people are using it to search for and understand and analyze information followed by social interactions from, ranging from casual chatting to more in-depth discussion or debate.

We also see people use it for writing assistance, interacting with visual content, seeking help. And we see Meta -- people engage with Meta AI from several different entry points. WhatsApp continues to see the strongest Meta AI usage across our Family of Apps.

Most of that, WhatsApp engagement is in one-on-one threads, followed by Facebook, which is the second largest driver of Meta AI engagement, where we're seeing strong engagement from our feed deep dives integration that lets people ask Meta AI questions about the content that's recommended to them. And we're obviously excited about the launch of the Meta AI standalone app.

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

Eric Sheridan: Thanks so much for taking the question. Maybe following up on Brian's question and coming at it from a different angle. I appreciate the color on the use cases you're seeing today for Meta AI. How would you suspect those use cases evolve with a standalone app?

Can you bring us into a little bit of the decision process to do a standalone app, what that might change in terms of utility, frequency or scale relative to what you see inside Family of Apps today? And how do you think about positioning Meta AI as a standalone app, I guess, the competitive landscape today of other standalones for the consumer AI apps? Thank you.

Mark Zuckerberg: Yes. I can talk about that. We're going to focus on both integrating it into our Family of Apps in more ways and building a standalone experience.

I think some people want faster access to it or a more built out feature set than you can build into an app like WhatsApp. So the standalone app will be valuable for that.

I also think the standalone app is going to be particularly important in the United States because WhatsApp, as Susan said, is the largest surface that people use Meta AI in which makes sense if you want to text an AI, having that be closely integrated and a good experience in the messaging app that you use makes a lot of sense.

But we're -- while we have more than 100 million people use WhatsApp in the United States, we're clearly not the primary messaging app in the United States at this point, iMessage is.

We hope to become the leader over time. But we're in a different position there than we are in most of the rest of the world on WhatsApp.

So I think that the Meta AI app as a standalone is going to be particularly important in the United States to establishing leadership in as the main personal AI that people use. But we're going to keep on advancing the experiences across the board in all of these different areas.

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

Justin Post: Great, thank you. A couple of questions. Just on the guide in the second quarter, there are reports of potential supply issues in eCommerce, how you thought about that in the guide and maybe how you're thinking about it for the back half. And then on a bigger picture question.

Your CapEx spend is now close to some hyperscalers with very big client bases. Just help us conceptualize the kind of ecosystem you're building with your CapEx. I know you gave a lot of help on the intro, but maybe the ROI works without direct enterprise spend to drive revenues. How you're thinking about that? Thank you.

Susan Li: Thanks, Justin. On the Q2 guide, there's uncertainty, obviously in how the macro environment will evolve over time and how that could impact different segments of our business.

Our Q2 revenue outlook aims to factor that in. And partly -- that's partly why the \$3 billion range reflects the potential for a wider range of outcomes.

Specifically, we have seen some reduced spend in the U.S. from Asia-based e-commerce exporters, which we believe is in anticipation of the de minimis exemption going away on May 2nd. A portion of that spend has been redirected to other markets, but overall spend for those advertisers is below the levels prior to April.

But our Q2 outlook reflects the trends we're seeing so far in April, which have generally been healthy. So it's very early. Hard to know how things will play out over the quarter and certainly harder to know that for the rest of the year.

Your second question is about why we're investing more in CapEx. And we really believe that our ability to build world-class infrastructure gives us a meaningful advantage in both developing the leading AI technology and services over the coming years. And there are a lot of opportunities also for us to improve our core business by putting more compute against our ads and recommendation work.

So even with the capacity that we're bringing online in 2025, we are having a hard time meeting the demand that teams have for compute resources across the company.

So we are going to continually invest meaningfully here across our infrastructure footprint, but we are also really looking to build this capacity in a way that gives us the maximum flexibility in how and when we deploy it over the coming years, so we can respond to how the market and technology develop.

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

Douglas Anmuth: Thanks for taking the questions. I just wanted to follow up on CapEx and infrastructure spending. Just on the higher range for CapEx.

Can you just help us understand how much of that is tied to the additional data center investments versus the increased hardware costs and really what's driving those higher hardware costs?

And then separately, there have been some articles suggesting that you've been looking to partner to share some of the cost of the AI infrastructure build-out. Can you just help us understand your thought process there and some of the pros and cons of going alone versus partnering? Thanks.

Susan Li: Thanks, Doug. So our increased CapEx outlook reflects both of those updates, the increased data center spend this year as we have made some adjustments to flex our build strategy that will enable us to really stand up capacity more quickly, both in '25 and '26.

We haven't broken down sort of the exact drivers. The higher costs we expect to incur for infrastructure hardware this year really comes from suppliers who source from countries around the world. And there's just a lot of uncertainty around this, given the ongoing trade discussions.

And so that is both reflected in the wider range that we are giving. And we're also working on our end on mitigations by optimizing our supply chain, and our outlook is really trying to reflect our best understanding of the potential impact this year across all of that uncertainty.

On the second part of your question, we are -- we're pleased to have partners investing alongside us and bringing Llama to market like AWS and Azure who are helping us host Llama.

We're always looking for opportunities to continue deepening or expanding those partnerships. But we are funding the infrastructure that is being used to train Llama, and we don't have any expectation that that will change at this point.

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

Mark Shmulik: Yes, thanks for taking the questions. Mark, in your conversation last night with Satya, I think you both discussed a bit around kind of the portion of code being written internally by AI. Kind of back to some of your previous comments around this being the year where we might see an AI kind of the place of a mid-level engineer.

With the world evolving so quickly, can you share some places where you've seen strong traction there? And are we progressing kind of faster, slower or as you expected towards this milestone?

And then, Susan, with the expense guidance coming down just a touch, how should we think about just the overall cadence of expected spending really as it

relates to kind of core business performance and the realities of the day-to-day world we're living in?

Mark Zuckerberg: I can talk about the coding agent work. I don't think that there's been any real change in our prediction for the timing of this.

So I'd say it's basically still on track for something around a mid-level engineer kind of starting to become possible sometime this year, scaling into next year. So I'd expect that by the middle to end of next year, AI coding agents are going to be doing a substantial part of AI research and development. So we're focused on that.

Internally, we're also very focused on building AI agents or systems that can help run different experiments to increase recommendations across our other AI products like the ones that do recommendations across our feeds and things like that.

So I think that, if it works, should just accelerate our progress in those areas, that's the basic bet that we're making.

Susan Li: On your second question about our lowered expense outlook. Really, we are four months into the year. The lowered outlook reflects more refined forecasts including updated expectations for both employee compensation as well as some other non-headcount-related operating expenses this year.

And that's partially offset by higher expected infrastructure costs related to our increased CapEx outlook as well as higher expected Reality Labs cost of goods sold. And we've maintained our \$5 billion range just given the more dynamic operating environment that we're in.

And what I would say is our investment posture today reflects the significant opportunities that we see across each of the company and priorities that we're investing in this year. We will obviously continue evaluating depending on how macro conditions more broadly evolve.

But we really feel like these are big strategic priorities for us and are critical for us to continue investing in. And in fact, I think one of the aims of our efficiency work over the last two years was to put us in a stronger financial position so that we can continue investing in key priorities through tougher financial cycles.

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

Ross Sandler: Great. Mark, yesterday in one of your many kind of podcast or keynote presentations, you had mentioned that like a bunch of projects that your teams want to or aspire to do are kind of bottlenecked by the AI capacity, which Susan just talked about earlier and that even some of the testing that the ad ranking team wants to run is just getting kind of delayed.

So I guess looking out either this year or next year or whatever, when do you kind of see some of this constraint being eased back? And more broadly, we're kind of three years past the IDFA impact to your business.

So where do you think we are in terms of just the overall improvements to the ad ranking system, the ROI that you guys are able to deliver? And like what inning are we in on that in your opinion? Thank you very much.

Susan Li:

I can take a shot at both of those and Mark, you can obviously chime in. On the first question, the capacity landscape we are in is pretty dynamic, both in terms of the many moving parts in terms of us bringing capacity online, but also in terms of the demand from different product groups in our company, whether they are in the gen AI teams or whether they're doing more of the core AI work around ranking and recommendations.

So both the supply and demand are quite fluid. And so we don't have a sort of fixed answer in terms of when we expect that we will sort of have enough supply to meet all demand, but that's something that we are working very hard to alleviate and it's part of why we accelerated bringing more data center space online this year. And also, we're very focused on increasing the efficiency of our workloads over the course of the year. On your second question about ads performance, ads ranking.

We have invested for many years and continue to invest in driving ad performance improvements. Year-over-year conversion growth remains strong. And in fact, we continue to see conversions grow at a faster rate than ad impressions in Q1, so reflecting increased conversion rates, and ads ranking and modeling improvements are a big driver of overall performance gains.

We have a lot of innovations in model architecture in both the ads retrieval and ranking stages of the ads delivery process to serve more relevant ads to people.

We talked about the introduction of the new GEM ads recommendation model in Q1. And we have talked about some of the prior model architecture improvements like Lattice and Andromeda in past quarters.

For us, we really believe, first and foremost, that advertising is a relative performance game. That's especially important for us because the vast majority of our business is direct response advertising.

So we feel good about how the prior investments are paying off and we continue to invest in a lot of different work to constantly improve our ads ranking and recommendations performance.

Operator:

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

Kenneth Gawrelski: Thank you so much. Two for me, please. First, maybe, Mark, how should we think about the timing of AI capabilities necessary to drive WhatsApp for business adoption in higher labor -- cost labor markets?

What is Meta doing to accelerate that adoption? And do you see this as mostly incremental to SME ad spend that you're already capturing? And then for Susan, one, please.

What is the revised CapEx outlook for this year for '25 mean about future years? Does it mean anything? Or you talked about this being an acceleration in your revised outlook statement.

Should we think about this as a new starting point for -- to think about '26 and beyond? Or should we just start fresh in '26 and think about the needs in capacity at that point? Thank you.

Susan Li: I'm happy to take -- I'll go ahead and take both of those. And Mark, you should feel free to chime in wherever you would like.

So Mark talked a little bit about our general vision that every business will soon have an AI that is an expert on their business for their customers to talk to you in the same way that today, they've got e-mail and websites, social media presences, et cetera.

We are currently testing business AIs with a limited set of businesses in the U.S. and a few additional countries on WhatsApp, Messenger and ads on Facebook and Instagram.

We've been starting with small business and focusing first on helping them sell their goods and services with business AIs.

But ultimately, we are working on tools to support businesses at every stage of the customer funnel from lead generation to order management and customer service. And a core area that we're addressing right now is really the ability for businesses to customize and control the agent to achieve the outcome that they want.

So we've launched a new agent management experience and dashboard that makes it easier for businesses to train their AI based on existing information on their website or WhatsApp profile or their Instagram and Facebook pages, and we're starting with the ability for businesses to activate AI in their chats with customers.

We are also testing business AIs on Facebook and Instagram ads that you can ask about product and return policies or assist you in making a purchase within our in-app browser.

So again, the ultimate vision is to build an experience that serves customers across all of these different services and apps. No matter where you engage

with the business AI, it should be one agent that recalls your history and your preferences.

And we're hearing encouraging feedback, particularly that adopting these AIs are saving the business that we're testing with a lot of time in helping to determine which conversations make sense for them to spend more time on. And then your second question, right, was about 2026 CapEx.

Infrastructure, as I alluded to earlier, just is a very dynamic planning area given the continued advances in AI and also for us, the fact that we continue to find a lot of good use cases to put capacity toward in our core AI ranking and recommendations work. So I would say it's too early to discuss plans beyond 2025.

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

Youssef Squali: Great, thank you guys for taking the questions. So Mark, in a world where we now have maybe 5 to 10 chatbots including Meta AI on our smartphones doing virtually the same thing. Do you think this is a market much like Search where the winner takes most or is it likely to be much more fragmented?

But in either case, what would you say are the top two or three competitive advantages of Meta AI? And then Susan, on the EU decision in connection with the DMA, what kind of modifications will you need to make to the apps? And can you maybe just help us gauge the potential financial fallout, understanding that it may obviously be too early. Thank you.

Mark Zuckerberg: Yes. On Meta AI, I mean I think that there are going to be a number of different agents that people use, just like people use different apps for different things.

I'm not sure that people are going to use multiple agents for the same exact things, but I'd imagine that something that is more focused on kind of enterprise productivity might be different from something that is somewhat more optimized for personal productivity and that might be somewhat different from something that is optimized for entertainment and social connectivity.

So I think that there will be different experiences. One of the trends that I think we're starting to see now is personalization across these. Right now if the experience is unpersonalized then you can kind of just go to different apps and get reasonably similar answers to different questions.

But once an AI starts getting to know you and what you care about and context and can build up memory from the conversations that you've had with it over time, I think that will start to become somewhat more of a differentiator.

So that's one thing that we think will matter. And then, of course, there's all the different modalities being able to not just answer questions about -- in text, but being able to do voice and multimodal and be able to produce images and

videos and understand all those things and have good conversations about that, I think, is going to be important overall.

So yes. I mean I think Meta AI is well positioned, but we have a lot of work to do in order to make it the leading personal AI.

Susan Li: And Youssef, on your second question, it is really too early to speak about what those changes could be because we are in the process of engaging with the European Commission.

I think maybe the most useful sort of metric I could give you is just that our advertising revenue in the European economic area and Switzerland, which would be the geographies impacted here, was 16% of our worldwide total revenue in 2024. Again, we are continuing to engage actively with the European Commission further on this, so we hope to have more clarity by next quarter's call.

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

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

Mark Mahaney: Thanks, I'll just throw in two. I think you called out that the China-based retailers as one sort of potentially soft advertising vertical. Anything else you'd call out? And I would just suggest autos.

Is that an area of any softness? And then on the Reality Labs and on the losses associated with Reality Labs, they've been very consistent \$4 billion a quarter for quite some time.

Is there light at the end of the tunnel? Is there a reason to think, is there a factor that would occur that would cause those losses to come down and when would that be? But maybe more importantly, what is going to cause those losses to come down? Thank you very much.

Susan Li: Mark, let me take your first question about other verticals. We generally saw healthy growth in most verticals in Q1. We did see some weakness in gaming and politics.

So year-over-year growth in gaming was negative in Q1 as we lapped a period of strong spend from China-based advertisers that were promoting a larger volume of game titles in Q1 of 2024. And then year-over-year growth in the government and politics vertical dropped sharply as expected with the conclusion of U.S. elections. But that continues to just be a very small vertical overall. And then your second question on Reality Labs.

Mark Zuckerberg: Yes. I can take the Reality Labs one. I mean we're basically focused on doing the work more efficiently. But as the AI glasses have really taken off, I've talked about this on a number of calls. There are more investments that I think makes sense to make around making sure that we can distribute this and grow it very quickly.

I mean some of the -- if you look at some of the leading consumer electronics products of other categories, by the time they get to their third generation, they're often selling 10 million units and scaling from there. And I'm not sure if we're going to do exactly that, but I think that that's like the ballpark of the opportunity that we have. And that's something that I think we're kind of focused on scaling to that and then scaling beyond that for the generations after that.

So I think some of the effort that we're doing is going to -- we're going to get more efficient in some parts of the work that we do.

But then as a bunch of the products start to hit and start to grow even bigger than the number that I just said is just sort of like the near-term milestone, then I think we'll continue scaling in terms of distribution.

And then at some point, just like the other products that we build out, we will feel like we're at a sufficient scale that we're going to primarily focus on making sure that we're monetizing and building an efficient business around it.

But -- that's kind of where we're at on it. We're definitely focused on doing the work more efficiently, but also very optimistic about what we're seeing in the results, especially on the AI glasses side.

Kenneth Dorell: Thank you everyone for joining us today, and we look forward to speaking to you again soon.