

Asana Work Innovation Summit New York: Investor Q&A

Wednesday, October 14, 2026

Hilton New York Midtown, New York, NY

Asana Participants

- Dan Rogers, Chief Executive Officer
- Arnab Bose, Chief Product Officer
- Aziz Megji, Chief Financial Officer
- Eva Leung, Head of Investor Relations

This transcript has been edited for clarity and readability. It may contain forward-looking statements that are subject to risks and uncertainties, and actual results may differ materially. Please refer to Asana's filings with the Securities and Exchange Commission, including its most recent Quarterly Report on Form 10-Q, for a discussion of the risk factors that could affect results. These statements reflect management's views as of October 14, 2026, and Asana undertakes no obligation to update these statements. Reconciliations of non-GAAP financial measures discussed in this session are available on Asana's investor relations website.

Opening Remarks

Eva Leung: Good morning, everyone, and thank you so much for joining us. I'm Eva Leung, Head of Investor Relations at Asana. I hope you enjoyed the keynote. With me today are Dan Rogers, our CEO; Arnab Bose, our Chief Product Officer; and Aziz Megji, our CFO.

A quick housekeeping note: this session is being recorded. Before we begin, a reminder that today's discussion includes forward-looking statements based on our current expectations, which are subject to risks and uncertainties. Actual results may differ materially. The full statement is on the screen behind me, and our risk factors are described in our SEC filings, including our most recent Form 10-Q. We may also refer to non-GAAP measures, and reconciliations are available on our investor relations website. And since we're in the middle of the quarter, we won't be commenting on current quarter trends or updating our outlook today.

Here's how we'll spend our time together. I'll start by asking each of our speakers a question that you've been asking us throughout the quarter, and then we'll open it up. If you'd like to ask a question, please raise your hand, wait for the mic, and share your name and your company.

Moderated Questions

How Asana's opportunity is changing

Eva Leung: Dan, let's start with you. Over the past month, we have introduced AWM and three new agentic applications, and this morning's keynote showcased them in action. What should investors understand about how Asana's opportunity is changing, and how do AWM and the new agentic applications expand the work we can support and the buyers we can reach?

Dan Rogers: Good morning, everyone. What did you think of the keynote? Did you enjoy that? Hopefully some of our differentiation shone through there. What you saw is a bit of our ambition about where we're headed.

The first big idea is that we were amazing at human-to-human collaboration. Could we take this same platform and make it super relevant for human-to-agent collaboration? As we did that, we realized it's not just humans and agents, it's also systems. So we acquired StackAI because it completes that trifecta: work that cuts across people, agents and systems, because that is how all work travels, as you saw in the many examples we shared today.

Agentic Work Management, our flagship product, is the evolution of collaborative work management. Collaborative work management was the single canvas for people to collaborate on projects. With AWM, we extend that to being the single canvas and harness for humans, agents and systems. Why is it a harness? Because our objective isn't just to coordinate and collaborate. It's to complete work: to assign work to people and to agents and drive it through to the end outcome. Take the Cindy example I shared on stage. In CWM, we would have helped Cindy, because she would have seen what all the people were working on and known when the work would be complete. In AWM, our ambition goes much higher. We're not resting until she has that new vendor, and Dash won't rest until we get to that new vendor.

So AWM isn't just CWM with new labeling. It's a fundamental engineering effort that we've been busy with for the last year: how do we bring humans, agents, workflows and our 500 connectors together on the same platform. We believe AWM is what makes LLMs actually work in the enterprise context. In your home life, as we talked about, you're seeing great productivity, and I love some of the new tooling that has come out from new vendors over the last couple of weeks for personal productivity. But when it comes to work, you need context. You need the harness. That's exactly what we've built, and we're ready for it. We can literally translate LLMs into productivity at work.

Then we said, why stop there? We know the jobs to be done for the core audiences we serve today. We know exactly what IT teams are trying to do, and in particular they're trying to do things with tickets. Our platform is now this harness for all teams, so why not make it great at serving tickets as well: resolving tickets against a knowledge base, pushing tickets to the right teams, and creating and enriching projects out of a ticket. It's such a natural fit for us. In fact, many of our customers were, let's say, Frankensteining CWM to make it work for service management. So why not just create a product? We created Asana Service Management. We're in v1 of the product, which we released in mid-September. We'll get a lot of feedback, and v2 will be even better, with really fast iteration cycles. For v1, about 30 design partners helped us get it to the point where it was super useful for them, and now we're going to take it more mainstream.

Similarly, a large percentage of our customer base is in client-facing teams, serving either external clients or internal clients, as in a creative organization. They wanted some very specific things around managing their clients and giving clients visibility into their work. Work trapped in CWM, without the client having visibility into it, wasn't that helpful. They wanted to manage the back and forth with those clients instead of it being trapped in email, phone calls and Slack. Could that happen in a single coordinated portal that captured that work and offloaded some things to agents? So we created Asana Client Management.

We also saw that for R&D teams in particular, the lingua franca is now pull requests, not tickets. So why not purpose-build something to coordinate the PRs that are now being created by all of these agents, so that everyone can stay in flow instead of having to report back to the product teams or the design teams? Hence the creation of Command.

With the acquisition of StackAI, we can also more fully serve another persona, whom you might call the process architect or the transformation leader. Sometimes they call themselves the COO or the CDO. They want to reimagine whole processes at a time, and StackAI allows us to do that.

So now you have this broader ambition of us being the company that brings the agentic enterprise

to life, as well as purpose-built products for the customers we already serve.

What Asana uniquely contributes versus frontier labs and software vendors

Eva Leung: Thanks, Dan. Arnab, investors are seeing agents from the frontier labs and from established software companies. Can you walk through a customer workflow that illustrates what Asana uniquely contributes, and why our customers are choosing our approach?

Arnab Bose: Dan covered a lot of this in his summary of what Agentic Work Management is: a harness with a context layer of who does what, by when and how, and a way to bring AI agents into that same fabric.

The big difference we provide over using any one of the frontier lab models directly is this. When you use a frontier lab model directly, you typically have a one-to-one relationship with that model. You chat with it, it produces content, and you publish that content into some other system where other people can collaborate with you on it. The first problem is how that model gets access not just to a flat set of files, chats or emails, but to structured data about how a process or project should run. It doesn't have access to who does what, by when and how, or to the historical runs of how the customer or employee has completed that workflow within their teams. So there's a very interesting component around context and the graph that we've built over 17-plus years, which is native to our system. So step one is context.

The second is that our agents are entities you train as a team. When you bring in one of these AI Teammates, especially a pre-built one that gets automatically suggested to you as a non-technical knowledge worker who just wants to get their work done, that project manager or campaign manager Teammate is something everyone on your marketing or operations team can use and create memories around. That means whoever is the most AI fluent on your team is actually raising the AI fluency of every other person on the team. Say Dan, Aziz and I were running marketing for Asana. Dan has been a CMO; I've never been one.

Dan Rogers: I'm not going back.

Arnab Bose: The context he brings and his ability to use that agent will far surpass mine. But as he uses it, whenever I invoke the agent, it will be as good as if I were Dan. So that elevates the people as well.

The third thing is that because we have this structured context graph, we can take these memories and convert them into reusable knowledge. Say we decide that from this point forward we are going to optimize for mid-market companies, and our campaign strategy will work in a particular way. That decision could get trapped inside the memory of a single frontier lab agent for a single marketer. In Asana, it gets elevated to a Work Graph object that is reusable by people and agents going forward. It can become a page, a completed task or a completed approval. That builds a self-improving graph that makes AI outcomes that much better going forward.

So: agents you can coach as a team, shared memory, reusable context, and elevating memories to knowledge objects. These are deep concepts we've been working on for the last 18 to 24 months, and they give us that differentiation in the market.

How AI usage and new applications translate into growth

Eva Leung: Thanks, Arnab. Aziz, one for you. We are including AI usage across the paid base and also introducing new applications. How does that translate into growth through retention, expansion and paid consumption, and what signals will tell us the model is working?

Aziz Megji: We've been fairly consistent that long-term growth acceleration is what we're striving for, and everything you saw today in the keynote is focused on growth acceleration.

The company has evolved from a single-product company with collaborative work management to a multi-product company. We now have five product families in the market, which gives us a lot more to expand with customers. If you look at our primary growth lever over the last several years, expansion has been driven by pricing and seats. Now we have consumption-based products, with AI Teammate requests, to expand with customers. So we're not just tied to headcount. We're tied to outcomes and the work we're driving for those customers.

We're also tapping into new buying centers, as Dan laid out: the engineering buying center, the IT buying center, and professional services and agencies with Client Management. So our TAM has expanded greatly. Having more to sell to customers and expand with them is a key lever for growth acceleration.

It's also a key lever for retention. We're getting deeper into business-critical workflows, and as customers drive more outcomes, we become a lot stickier. So both seat retention and the repeatability of consumption become the key drivers of our growth algorithm going forward. The seat is still really important. The subscription is the core and the foundation, with consumption, new buying centers and outcomes on top of that. And we've been clear at Q2 that this is not instant gratification in the second half. It takes time to build pipeline and convert it into revenue. But we believe that in the intermediate term these are important levers to expand and accelerate growth.

Audience Q&A

AI pricing across the industry

Eva Leung: Thanks, Aziz. With that, let's open it up. Who would like to go first?

Audience 1: The biggest question I get is about pricing in the industry for credits, tokens, whatever you call them. If we are cutting prices, how should investors, or customers, not view this as a demand problem for your AI services?

Dan Rogers: Who is "we" that is cutting prices?

Audience 1: Software companies. Application companies in general putting credits into SKUs, and either not charging for them or charging less for them to have an impact.

Dan Rogers: The reason I asked is that I have a slightly different mental model on the question, so let me break it into a couple of pieces.

First, the foundation models are lowering token prices fairly regularly, which is great. When I worked at AWS, we reduced the price of compute pretty much every six months, and I think that trend will continue. So the cost of fuel for all of these AI companies is decreasing. For the same quality of output, the input cost is decreasing, which is wonderful for all the companies building on the shoulders of LLMs. Someone else might ask whether the models themselves are commoditizing, since people can now switch freely between models based on a clear price-performance choice and pick the model that gives the highest quality answer for the lowest cost. I think that's definitely a trend.

I'd characterize what we're doing with requests a little differently. First, we find requests as a meter to be particularly customer friendly. We chose that meter because it's how our customers told us they want to pay, and it's how they perceive the world: I ask the agent to do something, and if it

completes that, you can charge me. That is pretty unique; I think we're the only ones doing something like that. It's a very linear relationship between asking and receiving.

Now, why did we choose to include many requests at no additional charge? Is it part of everything trending to zero? I see it very differently. The model costs are on a price-performance curve that we will all benefit from, just as with cloud computing. The reason we chose to prime the pump on requests is that when I start working with Teammates, with Dash and with AI Studio workflows, my work changes. We didn't want only a small portion of our customers to experience that. We wanted all of our customers to experience it. So we made a bold decision to prime the pump and make sure every one of our customers experiences Agentic Work Management, not just collaborative work management, because we think it is a tectonic shift and we want them to see it. Not just the customers who could imagine what these things are ahead of time and were therefore willing to prepay. The reality is you don't understand the power of a Teammate until you use it, until you get past a certain amount of usage and see it improving and getting to know your habits and your taste. We believe AWM is a movement, and we want it to be a real movement, not just one or two customers.

At the time we made that decision, it was a bold decision for us financially too, because under the hood the models were still expensive. I was adamant we were going to make AWM prevalent. At the time, it was a potentially costly decision. Model prices have become a lot more affordable, perhaps 60% to 70% more affordable. We have also gotten very good at routing and optimizing tasks to different models for different task types. That's how I'd break down your question.

Third-party agents and access to the Work Graph through APIs

Audience 2: A bit more of a product question. You addressed the perceived threat from third-party agents, and there has been a lot of talk about that recently. Asana has the context, the memories and so on, and that all makes sense. But investors still have the idea that you can plug your agent into the Asana API and get all, or 95%, of the context that exists. So the bear might say you can still pull all of the context out. Could you explain why that isn't the case, or how we should think about how transparent the API is to third-party agents?

Arnab Bose: Great question. The context you're talking about is only one of three things that make the Asana platform differentiated. Yes, there is context in the platform, and yes, you can connect to the Work Graph via API and MCP and connect one of your frontier model chats directly to it.

Dan Rogers: If you're a customer.

Arnab Bose: Exactly, if you're a customer. As Dan was saying, that is a good thing and a tailwind for us, because the more you interact with the Work Graph and use the data, the more embedded we become in our customers' workflows, and it also writes updates back in. If you manage your tasks and projects in Asana and you want portfolio-level reporting on goals, even if you use a third-party agent plugged into MCP, you're writing back into Asana, and that makes for a stickier customer.

So yes, you're pulling the context out, but when context goes out through MCP, it flattens everything. You get back, in a single-player way, what historical context exists for that particular project. What are you not getting? First, you're not getting the ergonomics of multiple people and the AI agent working together inside a work canvas. Say you create a report from today's conversation with your team. If you use the Asana MCP and publish that back based on your own personal chat agent, your team sees that you posted a comment or a task with that data. They won't see the context behind your prompt, and they won't be able to go back and forth with your agent. They can't ask, for example, what topics came up in the meeting recording or what KeyBanc

asked. So there's a multiplayer ergonomics perspective that you completely lose.

Second, the deep work we've done on sharing rules and making sure shared memory works correctly. If I reuse a campaign manager agent across projects where some people have access to one project and not the other, that's where the secret sauce lies. We've written white papers about how we support shared memory, how we handle the permission overlaps and the diffing, and how we convert those memories into things that get reused. The ability to work together as a team is very differentiated, and it's something you can't extract just by talking to us through a chat.

Third, going back to what I said earlier, when you record things back into the graph, the AI Teammate can reuse them in a way that a third-party agent plugging in directly via MCP can't, because it won't know that a given memory has to be elevated to knowledge. These are ergonomics you only get when multiple people work with the agent.

I worked with Anthropic on a thought piece. They have an AI fluency team that is trying to educate people about how to build the human-agent teams of the future, and they look to Asana as a model for how that should be done: coaching agents as a team, raising everyone's AI fluency in the same way, and getting reusability out of it. That piece is on their public site. So those are the three ways it works. I don't look at the MCP option as a threat I have to cut off. I look at it as an onboarding step and a way to get more data into the Work Graph and get stickiness out of it.

Dan Rogers: I'd just reiterate that in all scenarios, this only works if you're an Asana customer.

Differentiation versus AI labs building their own context

Audience 3: You talked about the AI labs building their own context for their own teams, and it seems they're moving that way, with OpenAI launching spaces and their own agents. I'd love to hear what differentiates Asana from ChatGPT, OpenAI and Anthropic as these labs move into that space, particularly for newer customers who don't already have their data in Asana. Why would they choose Asana versus starting with a lab?

Arnab Bose: I can start with that, and it's a good progression from the last question. Say the labs copy everything we have: building a context graph and the 17 years of user experience about how these project spaces should work. Even in that scenario, there are some interesting challenges from a customer perspective. Say you're not an Asana customer, and you're evaluating Asana against a theoretical frontier lab product that is an exact copy. A frontier lab is also selling its own book in terms of tokens. Is that something you want to be locked into for your organizational behavior, your most critical business workflows and how your business works?

This is a forward-looking statement: one of the options we're working on is allowing customers to bring their own AI, to bring their own API keys, or to run in specific Bedrock or Vertex instances, whatever their provider of choice is. That comes alongside, as Dan said, a first-in-industry pricing model based on requests and actual value. We believe in value. We believe in being open and transparent about the connections you can make into and out of the system. And we are not the token vendor, which gives us an inherent advantage in being more customer friendly with our product.

Requests versus credits, and early customer conversations on AWM pricing

Audience 4: Thank you for doing the Q&A. A high-level question. You made the new pricing transition with AWM included in new products. Over the past year and a half, there have been some enterprise headwinds in tech, and some PLG trends that haven't worsened but have stayed similar. If we look head to head against some of the AI labs, there is pressure from enterprise customers saying, "I'm spending too much on credits." What drove the decision to price on a use it

or lose its credit basis, which could be seen as similar to some of the AI lab pricing mechanisms that are facing scrutiny? And as a follow-up, how are early conversations going with customers who are seeing this for the first time, who were renewing on CWM seats and now see AWM with the new pricing model? What do those conversations look like beyond the utility of AWM being better?

Dan Rogers: There are quite a few pieces in that. Let me take requests versus credits and what that means for customers. Credits are different from requests, so let's start there. With any of the LLMs, if you ask it to do a job, you have no idea up front how many credits it will use to get there. In fact, the more poorly you provide context in your prompt, the more credits it burns. We don't penalize our users that way. Whether you're good at prompting or not, you're trying to get something done, and only at the point where we've fulfilled that do we charge you. So the request is a unique meter. It's not the same as the labs. It puts the burden on us, not the user, to be efficient at taking the context and passing it in the right way. I think that's a departure.

In terms of budget pressure, we don't see it that way for our customers. They find this meter to be a fair value exchange, and we're not really seeing customers say they need to cut down the consumption piece of their bill. It's more the opposite. On the fixed piece, the seat piece, they're asking how they can transition more into consumption and what that relationship looks like. That's more where we are. Aziz, do you want to add to that?

Aziz Megji: The other thing is that we're tapping into new budgets and conversations. We've historically been very strong with marketing, project management offices, transformation offices and operations. These consumption-based tools are opening us up more to the CIO and to AI transformation offices, so budgets are expanding in terms of how they think about Asana. It's not just going to the marketing department and drawing on that budget. That's a key part of it as well.

It's also early. AWM has been out for less than a month. AWM is a packaging change, in that you're getting seats plus requests, but it's also a core experience change. Our customers are now discovering Teammates in the flow of work, which is accelerating Teammate adoption. We're also giving them requests to try it out and get to what we call an aha moment, where they see the value and want to use more and move into a consumption path. It's too early to comment on that consumption path, and we'll talk about it more on our Q3 earnings call. But so far the receptivity has been strong to how we're evolving not only the packaging but the core experience. We're going from human-to-human to human-to-agent, and not just giving you the requests to do that, but also changing the experience so you can discover these Teammates more efficiently, readily and transparently.

That's really important down market, where you don't have people helping you in the discovery process. The product does that, and that's key. As we think about PLG and the headwinds there, part of it is that we've been selling one product into that base. Now that base has not only AWM but 30 different Teammates to potentially consume and adopt alongside our agentic applications, and the product helps you discover them. That's a powerful adoption tool.

Asana Dash: early feedback and how usage shows up in the business

Audience 5: Asana Dash is a really interesting product because it's almost simple is best. For a lot of your customers, a native chat experience is very familiar, compared with some of the more forward ideas people throw around with AI. Dan, from your perspective, you almost don't have to guess what the customer wants; they'll just build it themselves. What has early feedback and utilization been like? And Aziz, how does that usage manifest in the business? Are you more interested in free-to-paid conversion from things like Dash, or in larger customers adopting it faster?

Dan Rogers: I really like the start of your question. That's definitely the vision for Dash. I'm

probably Asana's number one internal champion of how far Dash can go. Could it be the window into the agentic enterprise for all knowledge workers? That's my ambition for Dash.

I'd characterize it in two ways. First, the engineering effort around Dash. We add new capabilities and performance improvements by the week, so every time you open Asana, Dash has probably gotten better. It's on a very steep curve. Internally, we're on the beta versions of Asana, and I think we update Dash twice a day.

Second, there's a curve around user adoption. When users first use Dash, they probably don't understand the power of what it can do, so they start with "help me find some stuff," like finding a task or a person. A week in, they start experimenting and asking more ambiguous questions. Where we'd like to get to is that you can be as ambiguous as you like, and we'll disambiguate it because we know work. That's the end point for Dash.

One thing I'm excited about, and this is forward-looking, is transposing ambiguous things into workflows. You saw this idea in the keynote. The average company probably has a thousand processes, and most of them are pretty messy, with a lot of back and forth. Which of those can we identify as workflows? I think Dash [over time] can be your companion in that, and in visualizing those workflows. Under the hood we now have AI Studio and StackAI, so we have powerful workflow capability, and Dash could be the translator. So watch this space. Many of you may be Asana users. If you are, give Dash a try next week, and I think you'll see it's moving at a frightening clip.

Arnab Bose: I can add a couple of product vision thoughts. I strongly believe the current AI chat experience falls well short of what engaging with AI should be. Why? Because you have to figure out what you want to get out of the system, and you have no idea what you'll get. It's like a box of chocolates. Sometimes it gives you gold, sometimes it doesn't.

The models are getting really good. If you can connect that system to a context system that knows your work, your priorities and your goals, it becomes much easier for it to become a system of next best action. Instead of having to prompt it, or relying on suggested prompts, and this is forward-looking and something we don't have yet, with that context and the ability to route to cheaper models to produce work automatically in a cost-efficient way, you could start suggesting next best actions to people in a way that's more natural to them. It feels more like a feed of things they accept or reject, versus having to figure out how to complete a task and then asking AI to help produce a brief, a summary or some content. The system has already figured out what you want to do based on your task list or what your team is working on, has created an initial brief, and you simply accept it. Dash gives us a way to express those enterprise actions very well, all built on the platform fundamentals you saw live today.

Here's how my team uses Dash through Command, since we're in R&D and already live. A PM wants to build a feature. They have a customer conversation and jot down a few ideas as bullet points in a page. They meet with designers, researchers and engineers, and record that meeting with their brainstorm. That meeting transcript can then be connected to Dash: "Dash, take this meeting transcript and turn my initial thoughts and our discussion into a full PRD." Dash can do that very effectively, and in a token-efficient way, because it has access to all the historical PRDs we've converted into tickets and approved in the system. Once the PRD is ready, the PM can highlight sections and say, "I know I'm missing security and privacy feedback here. Look at how we've addressed these comments in the past and flesh this out to the level of detail a coding agent needs to understand what to build." That's possible in production today. You can then create effective tickets and hand them off to your coding agent of choice. So next best action, and using context to make these actions highly efficient, are things Dash is capable of. Next best action is not available yet. The coding agent workflows I described are live in Command today.

Aziz Megji: On the last piece, for both Dash requests and Teammate requests, we're really looking for engagement with what's included. If you're adopting and using it on a repeated basis, that will naturally lead to consumption. But through the renewal cycle, it should also drive better retention and expansion outcomes on seats. We see that when customers are well engaged and utilization goes up on their seats, they have better renewal and expansion outcomes. So first and foremost, if you're using it, you'll naturally move from free to paid. But it's really the renewal and engagement outcome, especially in the sub-\$5K cohort we called out on the last earnings call. That cohort is currently the gating factor to getting us to 100%-plus NRR, if you look at the [positive Q2] trends in the \$5K and \$100K cohorts.

In-product discovery and go-to-market for new buying centers

Audience 6: A two-part question on go-to-market. First, on the in-app discovery experience and exposing customers to Teammates that way. I spoke with someone at breakfast who built a Teammate through that exact process. How do you balance that with not disrupting the user experience? Second, on the go-to-market side, with these new agentic applications you're trying to get into new buying centers. How do you gear the go-to-market organization toward that without disrupting existing sales processes?

Dan Rogers: I'll take a crack at both, and Arnab may want to add on the product experience.

You know the saying: if a tree falls in the forest and no one is around to hear it, did it really fall? I don't want that to happen with our innovation. Given the pace of innovation and how useful teams will find this, it can't be like that tree. So for some period of time, I'm willing to accept us showing the way in a slightly more paternal way in the product experience than may be our steady state. I believe the transition to AWM is a crossing of the chasm for our customers. Design is everything and the user aesthetic is everything, but there will be a period where we want our customers to step into the future. How long that period is, I don't know. We haven't heard customers say it's obtrusive. We're being very thoughtful. In the product and in Dash, we'll tastefully recommend the right Teammates and tastefully introduce you to AI Studio. Even with Dash itself, users may ask where the search function went, or where the chat function went. It's all Dash now. As always with Asana, we've been design first. But yes, we want to prime the pump on this experience, not with monetization in mind, but because people deserve this way of working. It's a game changer.

On the sales piece, the practical reality is that our sellers are very good at engaging certain buying centers already, but they're less familiar with some of the newer buying centers we're talking about. So we now have a mix of core sellers, specialist sellers and incubation sellers. Core sellers already speak the language of the customers they're talking to, and yes, they need to be able to have an AWM conversation. Specialist sellers augment those conversations; you can bring in the experts for specific conversations around IT or process architecture. And then we have a different team entirely for incubation. Today we don't want our core or specialist sellers talking to R&D teams, so we have a separate team for that. What the final model looks like, we don't know, but we've fit this to the conversations our customers want to have for now. Over time, people will ask questions about the efficiency of the model, and I think it's too soon for that. What I care about right now is meeting our customers where they are, with the quality of conversation they expect from us.

Arnab Bose: I have no notes on your product answer.

A marketplace for Teammates

Audience 7: It seems like every time someone creates a new Teammate, they're effectively creating a mini SaaS business. I've done this many times: custom MCPs that often sit outside your

environment, custom RAG, maybe even fine-tuned models. What are the plans to create a marketplace for this, and is that the eventual evolution? If I'm giving that agent memory, my resources and my custom tokens, why can't I sell it? What would your take rate be? It seems like that could be the whole business in five years.

Dan Rogers: Let me have a go, and Arnab may add on. To answer this, we probably need to draw a line between what we know and what we speculate, so I'll start with what we know.

We pre-built 30 Teammates to start, instead of letting everyone custom build their own or bring their own third-party agents, because we wanted to show the way. Trying is believing, and we had an opinion on the most important agents for the enterprise. You could look at our 30 Teammates and see that we've thoughtfully identified the most burdensome manual work cyclones and created an agent for each. If you look at the distribution of tedious manual processes, we went after roughly the top 80%. We built those 30 thoughtfully and trained each one carefully. That's what we know.

We also know that beyond that 80%, there are other things customers will want to build themselves. They'll want to custom build agents, which is possible today, and we've given customers plenty of guidance on how to do so. That's the next piece.

We also expect that people will want to bring in agents they've built elsewhere, and that last piece gets into the realm of speculation. To be intellectually honest, there are four or five ways we could do that, and questions of how and when. So the best thing to say for now is that we're still in the first phase: getting our 30 Teammates widely adopted, because we think we've done a thoughtful job of getting rid of the work cyclone in those areas. That needs to be our obsession for the near future. Whether that's two, four or five quarters, we'll figure it out internally.

Arnab Bose: I agree with all of that. We are seeing interest from vendors who build custom agents and want those agents to be part of the Asana Teammate infrastructure and show up inside Asana. I also lead our technology partnerships function, and I'm seeing interesting inbound questions. Say a company specializes in marketing AI agents that do creative or other work. They see those agents producing content that gets plugged into Asana, and instead of a flat MCP integration, they would prefer that their agents show up inside Asana as an AI Teammate, so they get all the shared memory, context and multiplayer experiences. So there already is inbound interest. As Dan said, we want to spend at least the near term getting it right and making sure customers can use the pre-built AI agents. As customers fully understand the value of multiplayer, shared memory and reusable knowledge, that will give us more precision on how third-party agents should show up and what we should be monetizing. The people building third-party agents are also in R&D, so they understand the lingo very well and they're all excited. But we need to make sure we don't get ahead of ourselves, and that our customers are excited and bought in first. That's the sequencing.

Evaluating and rating agents

Eva Leung: We may have time for one last one.

Audience 7: A quick follow-up. In the early days of the app store, a lot of time was spent on reviews to make sure code was safe, and on curation. Are people rating your agents, the way they rate Uber drivers? How are you evaluating agents, and how can I evaluate agents within my organization? If two similar ones are created, how do I know which one is better?

Arnab Bose: Great question. There are a couple of layers. First, inside Asana, whenever you interact with an agent, you can give it thumbs up or thumbs down feedback. That goes into a full eval framework we have on the back end. For the pre-built AI agents, we have a lot of specificity about what that eval is, and that also helps us on the R&D side to optimize model choice and

selection, and to keep improving as models evolve. We get early access to many frontier lab models as well. If you look at our blogs, you'll see we were early on the latest Haiku and Sonnet models, for example. So there's a full self-learning loop built in with feedback.

Second, there are ergonomics in the experience. When you go to the AI Teammate gallery, we first guide you toward a Teammate that is actively being used by your team and your organization. It has social proof: it shows that Aziz, Jay and Eric are using this particular Teammate, so you should use it too, rather than creating your own copy and fragmenting usage. The next step is to use one of the pre-built Teammates, and then you fall back to building a custom one if none of those work. So there's the experience, and then there's the whole eval framework behind the scenes.

Leadership priorities and capital allocation

Audience 8: My question is for Dan. You took over as CEO at a time when software was going through what people call the "SaaS apocalypse." Dustin is still on the board and controls, I think, 73% of the voting rights. Things have changed since you took over. Growth has slowed, multiples have de-rated, and cash flow has grown, but stock-based compensation is still a significant portion of it. When you look at the company today and think about the next five years and the long term, how do you make capital allocation decisions, and what are your top priorities today?

Dan Rogers: I'll break that into a bit about the transition, and then the capital allocation process separately.

On the transition, you might say the "SaaS apocalypse" happened about seven weeks after I started in the role.

Since then, relative to what we've committed to the Street every quarter, the record speaks quite well in terms of what we've done on both the top line and the bottom line. I think we're very predictable, thanks to Aziz, and very knowable in terms of what we're about. And there are plenty of reasons to believe there is a possibility that we become a major player in the agentic enterprise. So we're very predictable, with a steady pace of improvement, and a non-zero chance of major acceleration. We're playing for the big game.

We've hired some amazing talent. We are hiring people who are going for the big game, not for anything incremental. That's the mindset.

On capital allocation, for me it's pretty simple. We're an innovation company. We have a lot of swagger about what we're about and what we're trying to build. No one here is just trying to keep the lights on; we're shooting for big growth. So we are continuing to fund R&D first and foremost, to deliver on the human-agent operating system. We now have five product categories and five TAMs that we're early in, and more fully realizing that vision is what we invest in first. Of course, we want to be efficient as we go, so anything that doesn't meet that charter is under close scrutiny. But first and foremost, we're an R&D company, and let no one forget that. We are trying to be one of the innovators in the AI infrastructure that makes the harness and the context for agents work at work.

Closing

Eva Leung: That's all the time we have. Thank you all so much for coming, and if you have any questions, please reach out to me directly. Thank you.