Dell Technologies

March 5, 2024
3:35 PM PST

Erik Woodring:	Perfect. So, why don't we get started?
	Good afternoon, guys. My name is Erik Woodring. I lead the hardware coverage here at Morgan Stanley. I am delighted to have Yvonne McGill, CFO of Dell Technologies, and Arthur Lewis, President of the ISG business, here with us.
	Yvonne has been a mainstay at Dell for...you've been at Dell.
Yvonne McGill:	Twenty-six-26-plus years.
Erik Woodring:	I didn't want to say it. But obviously, recently stepped into the CFO role. And Arthur, president of the core business operations of ISG, responsibility for sales, strategy, general management of the business. So, both of you, thank you for joining us today.
Yvonne McGill:	Thank you for having us.
Arthur Lewis:	Thank you.
Erik Woodring:	So, before we start, I have two things I need to read here.
	So, for important disclosures, please see the Morgan Stanley Research disclosure website, at www.morganstanley.com/researchdisclosures. If you have any questions, please reach out to your Morgan Stanley sales representative.
	And then, from the Dell side, Dell Technologies statements that relate to future results and events are forward-looking statements based on the company's current expectations. Actual results and events could differ materially due to a number of risks and uncertainties, including those discussed in the company's SEC filings. The company assumes no obligation to update its forward-looking statements.
	Perfect. Well, thank you, guys, very much for joining us.
	Yvonne, you're sitting closer to me. So, I am going to start with you. I think what would be maybe most helpful to start this conversation is just to look back at your January quarter earnings report, Fiscal '24 as a whole, and maybe give us a bit of a post-mortem on the year. Obviously, you reported late last week. And then we'll get into it, but I'll follow up with kind of speaking more about the factors for '25. But maybe let's just start with a post-mortem on last year.

Yvonne McGill: Sure. Well, I think last year was a very exciting year for us. And it made a few things very clear to us. First, data and technology are critical, crucial to everything that's going on. And so, that was a great validation year.

And it was also that things don't always go as you expect them to during the year. And one of the things that I love about Dell and I love about how we navigate is our ability to navigate and be successful regardless of what the environment is.

FY24 – so, the year we just finished – certainly was a great evidence point there. We delivered $88.4 billion in revenue; EPS of $7.13; $8.7 billion of cash flow from operations. So, really great opportunity – great results there.

In Q4, we entered that quarter wondering exactly how it would come out, but really finished quite strong, delivering $22.3 billion of revenue and $2.20 of EPS. So, again, a solid performance there.

But I think, again, what was really exciting in the fourth quarter and as we enter this year is we saw some recovery, if you will, in some signs of recovery, like server. We saw a third quarter of server, traditional server, growth. AI, we can certainly talk about AI, and I know we'll talk about AI; we've been talking about AI a lot. But it was great to see that traditional server growth for the third quarter in a row.

We still haven't seen corporate and enterprise spending come back. And you all know from our portfolio of where we play, we are more weighted towards that corporate and enterprise side of the business customer set. And so, we haven't really seen that yet.

When I look through to FY25, which we talked about, we've just started. I guess we're in Month Two now with FY25. Really excited and really optimistic about the year that we're in. Now, why am I so optimistic? Well, AI is part of it, but it's great to return to growth. Growing is so much more exciting than not growing. And so, we are calling for growth this year and super excited to deliver against that expectation holistically. And we entered the year with $2.9 billion in AI backlog and pipeline growing. That's just one indicator of the year that we're in right now. So, super excited.

Erik Woodring: And Arthur will be our guy to dig into that. So, you stay put. We'll get to you in a second.

In terms of putting some numbers around Fiscal '25, you guide to revenue growth of about 5%, EPS growth of about 5%. Compared to your long-term targets you laid out in October, that was about 3% to 4% revenue growth, 8%-plus EPS growth. So, we're outperforming on the revenue side, a little underperformance versus a long-term target; and again, long-term multiyear target on the EPS side. Based on what you just talked about, why do you think kind of the shape of Fiscal '25 will turn out the way that it is – that you guided for at the last (inaudible)?

Yvonne McGill: That we've guided to it? Our strategy is consistent. So, we did lay it out in our Security Analysts meeting in October. So, good news, we haven't changed our strategy since October, but we're starting to see it play out already. And I would say in October we were even newer into AI and the opportunity ahead of us and how fast it was going to ramp. And so, I really feel confident about that.

I feel confident about our strategy. We've been consistent with that over a number of years. So, leveraging our advantages and taking and extending our competitive advantages holistically.

I think the tailwind, I could say "AI," like, a thousand times and it probably wouldn't be enough. We've probably said "AI" a thousand times already today. And I think that's just a great opportunity. But it's more than just the AI-optimized servers. It's really a change. It's a dynamic change in how we, how companies are thinking about the future and enabling the future.

So, that's really what I'm super excited about for FY25, the year we're in right now, but for many years to come. Because this isn't a one-and-done. Everyone won't be through an AI transformation this fiscal year. And so, it's just the beginning.

Erik Woodring: Perfect. Perfect. So, why don't we start with you and talk about the PC business, the CSG business? Then we'll eventually transition to ISG and kind of round out the conversation.

PCs, despite the fact how much everybody's using the term "AI" this week, PCs are still, at least in my conversations, top of investor minds. Can you help us understand how Dell is thinking about the PC TAM this year and what drives that kind of 2% to 3% CSG growth in 2025? And there's a number of factors that go into that, right? It's demand, it's pricing, it's catalysts, it's channel inventory, it's seasonality, it's where you're seeing spend kind of ebb and flow right now. But help us think about this CSG business for this year, again 2% to 3% growth.

Yvonne McGill: Sure. So, we think that this is the year that we'll start to see growth in the PC business. We've been in the longest digestion cycle – so, eight-quarter digestion cycle – in the history of PCs. And so, we know it's back-to-back years of double-digit decline. So, pretty amazing, never-seen-before results.

But it's time for a refresh. We've got drivers out there, whether it's Win10 refresh, whether it's AI – and I think that's less of a driver right now; I think more the age of the install base. But it's time to get the capabilities and get the environment refreshed from a productivity. That's the primary productivity driver for employees, for companies. And so, we want to make sure that we're on top of it. We're expecting it to start in the second half of the year, and that's what's embedded within our holistic guidance (ph).

Erik Woodring: And then, just the last one was in terms of pricing there, how we should be thinking about CPC pricing, consumer versus commercial, how that impacts (inaudible).

Yvonne McGill: I think from a pricing standpoint, we have a number of influential areas there. We have a commodity – component costs are expected to increase. That's certainly one of the drivers to price. We have a competitive environment. So, what we've seen is less opportunities in, let's say, the large corporate and enterprise space. We've seen lots of competition for there. So, I've got increased input costs. I've got more competition. And so, that's some of the drivers that we're expecting to see in PCs. So, I think that will have an impact on margins holistically. That is embedded and expected within our guide.

And you have then other areas that could drive change and influence, like maybe AI. Maybe you need, if you're thinking forward, which I hope all companies are, about what capabilities they need within their device two years from now, three years from now, they need to make sure that they're buying the most ready technology for the future.

Erik Woodring: Perfect. So, you mentioned margins. So, I figure maybe we go there for a second. And so, can you talk about the kind of cost and margin landscape, how that evolves over the next 12 months? You told us last week gross margins in Fiscal '25 will be down 100 basis points year over year. That kind of implies low incremental margins kind of below your

normal model. So, how should we all be thinking about those most significant headwinds that you're facing in light of growing? And that includes some of the more cyclical stuff. What would you do to try to offset those factors?

Yvonne McGill: Well, I mentioned some of it already, right? The higher input costs. We will expect to pass those on to our customers. And usually, we talk about a cycle that takes us 90 days or so to get that all priced. But we have visibility. We are expecting an inflationary cycle right now. So, we're being thoughtful about the pricing that we're putting out there.

I think of another element I already talked about, too, is the competitive environment. So, everyone's going after the same transactions. And so, we leverage the strong relationships we have with our customers, but there's some price aggression that we're expecting. And that's also embedded into our guide.

So, I think those are probably two of the key drivers holistically that we're seeing.

And then, I think of around AI servers, for example. We've talked about them being margin dollar-accretive, but rate-dilutive.

Erik Woodring: Mix shift.

Yvonne McGill: Yes, some mix shift there. So, you'll see that pressure holistically on our operating margin rate. And I think that's probably the bigger areas.

Erik Woodring: Okay. So, maybe, Arthur, I'll turn to you and dive into the ISG business a bit here. And there's a lot to unpackage between kind of the traditional business and AI. If we go back to October, you guided ISG to 6% to 8% annual revenue growth. That's up from 3% to 5% from your Analyst Day three years ago. Last week, you told us it was going to grow mid-teens in Fiscal '25. So, there's a clear tailwind.

What's interesting is that the shift of workloads to the cloud has historically been a headwind for a business that sells into on-premise data centers. But clearly, there's strength in ISG. So, why is this business perhaps better positioned than history and expected to see stronger growth this year than maybe you even talked about longer term?

Arthur Lewis: Good question, Erik, and thanks for having us.

When we met back in October, we talked about the fact that the ISG TAM was roughly $265 billion and growing at 7%. And embedded in that outlook was an AI TAM of about $124 billion projected out to Calendar Year '27. Since we met, not surprisingly, IDC has recast their forecast and now have it at $152 billion, growing at a 20% CAGR. Obviously, growing faster than the average. And obviously, a lagging indicator, right? Because customer demand continues to grow, and we would expect that demand to grow. And therefore, AI is the significant tailwind that we called back in October, which really underpins the mid-teen growth that we're calling for FY25.

In terms of deployments, I think it's important to start with the fundamental fact that 83% of the world's data sits on premises. And there's gravity to that data. And the reason why there's gravity to that data is because of cost, performance, and security. Eighty-two percent of the customers that we've surveyed definitely trends towards a preference on-prem for the deployment of generative AI. When you have information that touches proprietary company information, proprietary financial information, proprietary IP, CIOs and CEOs and boards really don't want to adjust their data strategy. So, I think there's

probably a stronger argument to make for repatriation than for more data going to the cloud because of cost, security, and performance.

What we see is a lot of customers are testing out workloads in the cloud just to prove out concepts. But when it comes to putting AI in production, they're coming to us to help them think through an on-prem solution. Given the breadth of our portfolio, our services capability, and our other durable advantages, we feel like we're very well positioned to capture the opportunity, and you heard a little bit about that last Thursday.

Erik Woodring: Perfect. And then, so maybe let's start on the traditional business, then we'll get into everything GenAI-related. So, Jeff has told us a number of times that the kind of typical historical cycle for storage and servers is kind of four to six quarters. This down cycle has been one of the longest in history for the traditional business. But as you mentioned, Yvonne, you are seeing green shoots. Traditional servers have grown sequentially for three consecutive quarters. You had a strong ending to the year in the storage business.

So, maybe simplistically, my question is, are we past the worst of the down cycle? What gives you confidence that the traditional business can grow this year, just given something you started the conversation with, which was large enterprise spending is still weak at the moment?

Arthur Lewis: I mean, there's no question that we've been in an elongated period of digestion: eight quarters. We started to see this period begin in Q1, Q2 of Calendar Year '23. And we've navigated geo conflicts, we've navigated global downturns. But the beauty of our model is that we never let a good crisis go to waste, and we prepare for the eventual rebound. And that's what we've been doing.

And as we look to FY25, we exited the year on the server business with three consecutive quarters of sequential growth and the first year-on-year growth. And I want to reiterate what Yvonne said: growing is a lot better than not growing. We see really strong adoption with our 16G portfolio, given the performance, security, and embedded intelligent automation that we have in that. And we finished strong in storage, above our seasonality, and it was good to see growth in our PowerScale portfolio and triple-digit sequential growth in our PowerFlex portfolio.

And look, I mean, at the end of the day, I know we're talking about the core business, but AI, generally, and generative AI, specifically, is going to drive a modernization around infrastructure that is not only relevant to generative AI systems, but is also going to bode very well for our core server and storage business.

Erik Woodring: Okay. Perfect. And Yvonne, I'll turn quickly to you because, obviously, as the CFO of Dell, you manage the finance department. You have your own internal conversations about what can maybe unlock spend in your business. And so, as a business leader, as a large enterprise leader, what are the catalysts, again as we think about weak near-term enterprise spending, what are the catalysts that you think would help you unlock spending within Dell, again as kind of a corollary to what we could see from the larger enterprise community?

Yvonne McGill: When I think of that, I really – I start with the macro environment. I think there's uncertainty in the macro environment, and that's not helpful for people to lean into spending. So, I think that would be one area: a little bit more clarity with the macro environment.

Customers are cautious, especially in the large customers, large enterprises. And so, we're working through that with them. We are them. And so, we're navigating through it together. And so, that's been really helpful. And really being more thoughtful about how to spend and when to spend and where to spend, how do you prioritize all of that.

That has been lengthening the sales cycle. And I do joke about they've inserted finance people into the approval process that didn't used to be there. I love that, kind of. But it does elongate the process when you put more approvals in there.

So, I don't think it's lost on anyone that refreshes are coming, that we all need more technology to deliver on the future. So, it's a matter of "when," not "if." And we're working on – we're ready when our customers are ready. And we're trying to help them see that now is the time to get started and get moving, moving forward on their technology journey.

So, that's kind of what I would think.

AI has been top of mind, and you have to allocate your capital spend. And so, I feel very comfortable with our holistic portfolio, that regardless of where our customers prioritize their spend, we will benefit from that holistically.

Erik Woodring: Okay. So, let's turn to the topic du jour with you, Arthur, generative AI. Maybe before we go into each specific opportunity, you mentioned the TAM earlier from IDC originally $124 billion by '27; now $152 billion by 2027; probably going to go higher than that. For you guys, Dell, you've recognized about $1.5 billion of AI server revenue in the year that just ended. You have a $3 billion backlog of AI servers at the end of the quarter. You have a pipeline that's even larger. I guess, my long lead into this question is, are we really, like, only in the first inning of this opportunity? Like, where is this growth coming from for you guys? I know it's a hard question to answer because it's such a fast-moving market, but take a stab at it. Where is this opportunity today?

Arthur Lewis: I think as we kind of take a look at what's happening in the business, the way I think about it is, to keep the baseball analogy going, there's sort of two games that are being played right now. I think in the Tier 2 CSP community – and these are the CoreWeave, the Denvr Dataworks of the world that are building out large training and inferencing infrastructure to support GPU-as-a-service to their end customers – we're in the early innings of the baseball game. But when it comes to the enterprise, we're in the "car pulling up into the stadium" and "the teams are still trying to figure out the rules of the game."

But what's incredibly exciting is that there's something different that's happening here. Because one of our differentiators is our capability and understanding of all things artificial intelligence. So, we're having conversations with enterprise customers that start way before the infrastructure conversation. We're talking about, "Hey, Dell, what use cases are you guys seeing?" "Hey, Dell, how do you think about model selection?" "Hey, Dell, how do you think about data preparation?" "How do you think about fine-tuning?" "How do you think about RAG?" And then you get into an architecture conversation and then you get into an infrastructure conversation.

And what's really cool about that is we're way up the stack in the conversation, which gives us great opportunity, a much higher opportunity, in order to capture that system that the enterprise customer is deploying.

And we're able to leverage what we're doing inside of Dell in Dell Digital and use that externally with customers to let them know, hey, we understand exactly what you're going through and this is what we're doing about it. In fact, Dell Digital is customer zero for us from an ISG perspective. Everything that we do, we consult with Dell Digital. We understand what they're doing. And we're all learning as we go. Every day we learn something new and we impart that information to a lot of the enterprise customers, and they really appreciate it.

I mean, at the end of the day, this technology, we believe, is a once-in-a-lifetime opportunity, one of the most significant advancements probably in human history. And so, you see these productivity numbers that are atmospheric in nature. And so, this now becomes a board-level, CEO-level conversation. And the aspirations are incredibly high.

So, folks are crossing T's and dotting I's and running a lot of POC's against a myriad of different use cases to really understand where are the productivity benefits. And it's opening them up to a lot of other ideas around, "Hey, I should really be thinking about modernizing my operations of the business in order to make generative AI that much better. I need to be thinking about streamlining my processes. I need to be thinking about automating my processes. And then I need to be thinking about generative AI."

I think this is going to drive a massive change in our industry. It's going to drive a massive change in infrastructure. This goes well beyond AI-optimized servers that we're talking about today in the context of Tier 2 cloud service providers.

Erik Woodring: Okay. Something I think we all are very interested in is – I'll make this question simple – why is Dell winning? Your commentary obviously stands out very positively amidst the market. What is your competitive advantage in this market? Again, why are you winning right now?

Arthur Lewis: I think it's – I would kind of boil it down to sort of five things. The foundation of it is the breadth of the portfolio and the thought that went into the creation of the portfolio. And it all starts with the 9680. We spent a lot of time on the architecture of this product. We had some pretty aggressive design points. Design point number one, it had to be the densest platform in the industry. Design point number two, it had to offer silicon diversity for all of the various GPU suppliers. Design point number three is that it offered diversity for networking and connectivity. And design point number four, it had to offer complete flexibility for thermals, moving from air cooling to liquid-assisted to direct-to-chip cooling. And the team did a phenomenal job and hit every single one of those design points. And we were first to market when we launched the 9680 with the H100 in it.

So, it starts with that. The 9640, the 750xa, the 760xa, the AI fabric that goes along with it, and the diversity among that.

And then, we add the storage portfolio on top of it. And today, a lot of people like to talk about the unstructured portfolio because of text, video, imaging, and whatnot. And so, we launched the new F210 and F710, which have increased performance, increased density, increased I/O. All of the things that provide for a very performance storage class product for generative AI. But down the road, you can see how structured databases are also going to be very relevant. So, while people talk about unstructured storage today, structured databases are going to be there as well.

And what's also important here is that we have the ability to co-engineer this system. Because generative AI is a system. It comprises compute, networking, and storage. And the ability to engineer that under one roof is incredibly important. Because if you're

trying to create this system through a myriad of partnerships, that level of collaboration that is required assigns a very high beta to that project. So, we have a very differentiated portfolio that we collaborate and design as a system.

And then we take that portfolio and we offer it as a solution and we create Dell-validated designs that are targeted at a bunch of different use cases. So, if a customer comes and says, "Hey, I'm looking at a sales chat bot" or "I'm looking for a customer service" or "I'm looking for something content creation," we have the ability to talk to them about a validated solution that comprises compute, networking, and storage.

That's the foundation of our differentiation.

On top of that, we have the ability to build out the ecosystem. We have the only partnership in the industry with Hugging Face and we have the only partnership with Meta and Llama2, and that's going to become incredibly important, as indexing data is incredibly important when you think about the generative AI system.

On top of that, we're able to build out our professional and consulting services. Professional services is what you would typically know Dell for in terms of support and deployment services, which are incredibly valuable not just to the enterprise, but to the Tier 2 CSPs, as well. But we've also taken the extra step of providing these consulting services to have the conversations that we talked about in the previous question.

We also have the ability to provide financing through Dell Financial Services, a one-stop shop for product and financing.

And then, lastly, but also extremely importantly, we have the industry's largest sales team – not only largest, largest and most knowledgeable sales team – and we have world-class supply chain that is wielding the power of Dell's scale.

You combine those five things together, and we very much like the hand that we're playing.

Erik Woodring: That was great. And maybe just to piggyback on that, just how do we think –? So, in the traditional server business, if we put the ODMs to the side and think about your TAM, the largest server provider in the world, how do we think about the market share and this specific opportunity with Dell in mind? I mean, there's only a few players in this market. Is this something where you can remain on the forefront – again not just in '23 and '24, but beyond – for the reasons that we just talked about?

Arthur Lewis: I mean, look, if you think about sort of my answer to the previous question, we think that we've created a pretty strong barrier to entry for competitors to come in, including ODMs. We are absolutely targeting equal to or greater share of this space that we have in our traditional server and storage market. And again, given all of the things that are working in our favor – the tailwind of the market, the progress we're making with the Tier 2 CSPs, what we're seeing in the enterprise – we're very comfortable with that.

Erik Woodring: Perfect. So, something I've been trying to better understand is how many kind of CoreWeave- and MBU- (ph) and Denvr Dataworks-like customers are out there. Meaning, these AI-as-a-service companies are buying AI servers in large volumes. That customer base is expanding, like you told us last week. But how much kind of follow-on demand and sustainability is there really with these customers? I know you talked about the early innings, but if you could maybe give us a bit more detail on the sustainability and kind of follow-on demand, that would be helpful.

Arthur Lewis:	I mean, it's tough, because I don't have my crystal ball with me. But I can tell you that we had more of the Tier 2 CSPs customers in Q4 than we had in Q3. And the customers that existed in Q3 that also existed in Q4 continue to tell us that they need more infrastructure because they don't have enough to support the needs of their customers. So, I keep thinking that Wave One is going to crest any moment, but it doesn't seem to be cresting anytime soon.
Erik Woodring:	Okay. And then, on the enterprise side, you said we're still driving to the ball park.
Arthur Lewis:	We're pulling into the stadium.
Erik Woodring:	Right. Can you talk about maybe qualitatively, again within the backlog and the pipeline that you have, how you've seen that enterprise either mix or spend trend over the last couple quarters? And what industries are buying these AI servers? What are the use cases that they're telling you that they need these servers for?
Arthur Lewis:	So, from a revenue perspective, the mix is not going to change dramatically just because of the scale of the Tier 2 CSPs. So, what we're looking for is the absolute number of customers that we're engaged with quarter over quarter and the mix of customers that we see quarter over quarter, which would include enterprise customers versus CSP customers, and then sort of the percent of the bookings and the percent of the backlogs. Everything we saw in Q4 was trending in a very positive direction.
	And it's really cool because we're seeing it across, really, every industry. I don't think there's any industry that stands out. It's financial services, it's manufacturing, it's retail, it's healthcare. There's a very healthy cross-section of vertical participation in really wanting to understand what's going on with generative AI.
	I think sort of the leading use cases are kind of like the four that we're really focused on. Software development comes up a whole lot. "Hey, Dell, how are you thinking about software development? Where in the software development process do you think generative AI plays the best role?" They're thinking about it from a sales chatbot perspective. They're thinking about it from a customer service perspective.
	And something as mundane as content creation. You'd be surprised, and I certainly was surprised, how much money we at Dell spent on this and how much money customers say they spent on this, and the amount of savings that they could realize through deploying generative AI.
	And then there's a lot of other core functions in the company, like finance and supply chain, that they're looking at deploying generative AI.
	But this is where the cloud kind of comes into play. Because they've got so many different use cases, they're running a bunch of different models to see, "Hey, where am I going to get the biggest bang for the buck?" Because they really want to think about, "I want to deploy this against the thing that's going to provide the biggest return on my investment."
	But again what's really cool, and I think I said this a little bit earlier, is there's a short-term, "Hey, I understand that generative AI is a system and it's a different workload from sort of traditional workloads." But customers' eyes are starting to open to say, "Hey, generative AI is going to drive a different AI-optimized infrastructure, going forward." Jensen in the video refers to this as an AI factory, and that's what we see transpiring over

time. If you think about the majority of the world's data today is cold, sitting in backup and archive, we think the equation of that's going to flip, where more of the data is hot/warm and is constantly circulating, feeding the AI engine, which is going to cause a significant retooling of data centers around the world.

So, again, this is, like, an incredibly exciting time to be in technology, and I think we're extremely well positioned to capture the opportunity.

Erik Woodring: Cool. And the last question I'll ask on this topic is maybe the most important and maybe the most underappreciated, which is, at earnings, you guys made the comment that for every dollar of AI server spend, there's two to three dollars of storage and services and networking attach that comes alongside of that. And so, maybe help us better understand exactly what that means for your non-AI server solution, because that could ultimately suggest the non-AI server TAM is bigger than the AI server TAM or that hardware TAM. And so, does the breadth of offerings that you guys have give you a chance to really supercharge what you've already talked about on the AI server side?

Arthur Lewis: Absolutely. Today – and not surprisingly, as folks are really trying to understand how the system operates – a lot of the focus is really on the compute side, because the model is sitting in memory and they're trying to kind of, like, really figure out what the use case is.

The storage component, you're going to need attached storage that's highly performing. We talked about density, performance, and I/O. But ultimately, vector databases are going to sit there, and you're going to need a superior ingest engine to really get the benefit of the generative AI system.

You can see that's a near to midterm solution. That's not going to scale. You're going to have to take that type of performance storage and basically kind of rethink your entire storage strategy. So, as time goes on, you're going to see the generative AI really drag not just compute systems, but the traditional core servers as well as traditional storage, because everything is going to be optimized for AI.

Erik Woodring: Okay. Yvonne, I'll shift back to you. One more PC question. Again, PCs and AI, we'll bring them together. But a number of your peers have said publicly AI PCs will be, I think it's, 50% of global shipments by 2026. Even as a PC bull, that seems high given we've just really started the launches of these devices. We might not even be in the first inning. How do you guys think about the progression of AI PC mix? And when we think about the real use cases to drive adoption, there's some skepticism. So, how should we think about those use cases? What's top of mind when it comes to these devices?

Yvonne McGill: I mean, I think it's interesting, and it's interesting to say what is an AI PC. And so, we think a PC with an accelerator is an AI PC. So, it could be an MPU, it could be a GPU. We have those now; we'll have more in the future.

How do you really take advantage of it? What capabilities do you need? I think that's really what we're, what everyone's talking about and what we're expecting to accelerate, going forward. But when is forward? And when does acceleration start? And so, we're thinking that we'll start to see more of AI capabilities being embedded in PCs that are starting the refresh cycle, which I already said we think is going to be towards the latter half of this year.

I think it's – I expect a lot more capabilities. You're going to want to have those capabilities embedded in your device. Because when the technology is delivered, you don't want to have to refresh your PC again. And so, if you think of a notebook, I'm going

to give a three-year – a notebook should be a three-year refresh cycle. I know we're elongated right now, but I'm going to plant the seed that's a three-year refresh cycle.

Erik Woodring: (inaudible)

Yvonne McGill: Yes, I'm going to make it – it's going to happen because we're going to talk about it.

But you're going to have these devices when you're refreshing them. All of the capabilities that are going to be delivered, you need to be able to take advantage of them with the investments you're going to be making. And so, I think that's part of what we're talking to our customers about, is what capabilities do you need and how can we make sure when you're refreshing your install base that you're refreshing it with the technology needed not just for today, but for the future.

Erik Woodring: Okay. Perfect. So, with the last time we have, I kind of want to touch on another very important aspect of this story, which is capital returns and capital allocation. An important part of your Analyst Day last year was the new capital return commitment that the company made, returning 80%-plus – emphasis on the "plus" – of free cash flow to shareholders, growing your dividend at least 10% annually.

So, if we rewind to last week, you announced a 20% annual dividend increase. You're going to have to buy back a lot of stock, as well. You obviously showed that in the January quarter. So, first, obviously, just confirming that's the right way to think about capital returns in Fiscal '25? There's no change versus that long-term guide?

And then, second, how do you think about balancing dividends and buybacks and liquidity of the stock and some of the actions that we've seen Michael and Silver Lake take that he talked about in October in terms of trying to help increase the float (ph). How do we think of all of those together in terms of capital returns?

Yvonne McGill: It's certainly complex. But we did roll out our long-term framework, with the "pluses" in there. So, we didn't feel limited to that framework.

We did announce the dividend increase of 20% last week, which was really, to us, a commitment and confidence in the capital returns we will be driving. And so, we wanted to make that statement. Is 20% more than 10%? Yes. And that's what we said we would do. And again, it's confidence. You don't take your dividend down after you raise it. So, just wanting to speak to the confidence there.

As we look through, since the dividend inception, we've delivered $7 billion returned to shareholders. So, that's a pretty sizable number. And we see that continuing. I think it's the – you talk about the stock, the repurchase that we're doing. Certainly, with the stock performing like it does right now, we'll navigate through it. We still think our stock is undervalued. And so, it is to our advantage to continue the repurchase cycle. So, we're very committed to that and to our overall capital framework.

Erik Woodring: Perfect. As we have 30 seconds, I know just to touch on free cash flow again – maybe I'm doing this out of order – but you've talked about 100%-plus free cash flow conversion. It seems like you still might have some working capital efficiencies you can capture. Maybe my question is, has anything structurally changed such that you couldn't capture some incremental working capital efficiencies? Is the emphasis on AI, does that change the capital intensity of the business at all? Or how should we think about that conversion rate?

Yvonne McGill: No, we're committed to our capital returns, to our cycle. And I don't think of anything that's going to change that, going forward. I am confident that the AI momentum will only improve that, going forward.

Erik Woodring: Okay. We're out of time. So, Yvonne, Arthur, thank you very much for joining us today.

Arthur Lewis: Thank you.

Yvonne McGill: Thank you.