

DELL Earnings Call Transcript

Date: 2024-11-26

Quarter: 3

Operator: Ladies and gentlemen, good afternoon, and welcome to the Fiscal Year 2025 Third Quarter Financial Results Conference Call for Dell Technologies Inc. I'd like to inform all participants this call is being recorded at the request of Dell Technologies. This broadcast is copyrighted property of Dell Technologies Inc. Any rebroadcast of this information in whole or part without the prior written permission of Dell Technologies is prohibited. Following prepared remarks, we will conduct a question-and-answer session. [Operator Instructions] I'd like to turn the call over to Paul Frantz, Head of Investor Relations. Mr. Frantz, you may begin.

Paul Frantz: Thanks everyone for joining us. With me today are Jeff Clarke, Yvonne McGill and Tyler Johnson. Our earnings materials are available on our IR website and I encourage you to review these materials. Also, please take some time to view the presentation which includes additional content to complement our discussion this afternoon. Guidance will be covered on today's call. During this call, unless otherwise indicated, all references to financial measures refer to non-GAAP financial measures, including non-GAAP gross margin, operating expenses, operating income, net income, diluted earnings per share, free cash flow, and adjusted free cash flow. A reconciliation of these measures to their most directly comparable GAAP measures can be found in our webdeck and our press release. Growth percentages refer to year-over-year change unless otherwise specified. Statements made during this call that relate to future results and events are forward-looking statements, based on current expectations. Actual results and events could differ materially from those projected due to a number of risks and uncertainties, which are discussed in our webdeck and our SEC filings. We assume no obligation to update our forward-looking statements. Now, I'll turn it over to Jeff.

Jeff Clarke: Thanks, Paul, and thanks, everyone, for joining us. We continued to execute well and delivered strong performance in Q3. Revenue was \$24.4 billion, up 10%, on strong ISG revenue growth. Diluted EPS was \$2.15, up 14%, and cash flow from operations was \$1.6 billion. Our focus on balanced profitability paid off, with improved operating income, earnings per share and cash flow. We're proud of the fact that year-to-date, our team has delivered 10% revenue growth in ISG and CSG combined and overall EPS of \$5.31. In ISG, our AI server momentum continued, and we saw substantial expansion in our five quarter pipeline. Orders demand was a record \$3.6 billion, up 11% sequentially, primarily driven by Tier 2 cloud service providers with continued growth in Enterprise customers. We have seen sequential orders demand growth every quarter since the launch of the XE9680. We shipped \$2.9 billion of AI servers in Q3, resulting in AI server backlog of \$4.5 billion. Our five quarter pipeline grew more than 50% sequentially with growth across all customer types. We continue to gain traction with Enterprise customers, large and small, with over 2,000 unique Enterprise customers since launch. Increasingly, Enterprises see the disruptive nature and the innovation opportunities with GenAI resulting in growing GenAI experimentation and proof-of-concepts. Underpinning our success is broad enterprise coverage, professional services, engineering leadership, and large-scale system design, all done with incredible speed and time to market. These highly specialized workloads and deployments require high-value engineering and a solutions mindset where custom designs and fast deployments are the norm. Beyond the AI servers, we like the profit pools that surround them, like power management and distribution, cooling solutions, network switches, network cables, optics, storage, deployment, maintenance, professional services and financial services. We

have accelerated the speed of innovation to respond to our customers' GenAI needs over the past year. A few highlights from the past two months. We launched our 21-inch ORv3 Integrated Rack 7000, in both a 44 and 50 OU rack design with integrated cooling, power and networking that is multi-generational and future-proofed up to 480 kilowatts per rack. This rack falls in our Integrated Rack Scaleable Solutions, which are focused on at-scale deployment ensuring Dell's AI Factories can meet the demands of foundational training at the Data Center scale. We are shipping the industry's first enterprise ready GB200 NVL72 server racks with our new XE9712, with direct liquid cooling that holds up to 72 GPUs per rack. We also announced at SuperCompute 24, a new AI server supporting NVL-4, also with liquid cooling, supporting up to 144 GPUs per rack, one of the industry's most dense designs. And we have the M7725, a dense compute design, which supports up to 27,000 CPU cores per rack to meet high-performance computing demands. Our IR5000 can achieve up to 96 GPUs per rack with a more traditional 19-inch rack design. Within the IR5000, which includes the XE9680L, we introduced the XE7740 and XE7745, designed for Enterprise customers focused on inferencing. We've increased the density and performance of our PowerScale products, updating the PowerScale F710 with 61TB of QLC SSDs and capacity density of 614TB per rack unit. And with a new software release that, when combined with our 200Gb Ethernet support, increases our streaming write performance by 163% and streaming read performance by over 2x versus the previous generation. An application of this technology that I'm excited about is a biomedical research center using the PowerScale F710 and the XE9680 to take their own unstructured data and train AI models that will advance personalized treatment plans resulting in better prevention, diagnosis and treatment of various types of cancers. The opportunity in AI is enormous and we are pushing the boundaries with our engineering, and all of this innovation is making its way into the Enterprise. In Traditional servers, demand improved double digits in Q3, making four consecutive quarters of year-over-year growth, driven by growing units and ASPs with denser core counts, memory, and storage per server. Customers are focusing on consolidation and power efficiency by modernizing their data centers with more efficient and dense 16G servers, freeing up valuable floor space and power that will support their AI infrastructure. In Storage, the demand environment continues to trail Traditional servers. PowerStore and PowerFlex demand grew double digits in the quarter, indicating that our focus on the midrange market, with product and software enhancements, will drive future growth and margin over time. The PowerScale F910 and F710 continue to ramp nicely, driving double-digit growth in the Unstructured All Flash portfolio. In CSG, we saw continued stability with demand growth in Commercial PCs for the third quarter in a row. More Enterprise customers are beginning to refresh, albeit modest and in a more price competitive environment. We are seeing an indication that customers are lining up their upgrade cycles with new AI PCs in the first half of next year, a clear signal that enterprises are balancing their need to refresh and their desire to future-proof their purchases. Consumer demand and profitability continues to be challenged. We are optimistic about the coming PC refresh cycle as the install base continues to age, and with Windows 10 reaching end-of-life in 46 weeks. The significant advancements in AI-enabled architectures and application development are welcomed tailwinds. In closing, AI is a robust opportunity for us in both ISG and CSG and interest in our portfolio is at an all-time high, with no signs of slowing down. That said, this business will not be linear, especially as customers navigate an underlying silicon roadmap that is changing. Regardless, we are winning deals at a price premium to our competition, delivering value across a much broader AI ecosystem. We have the balance sheet to grow the business and the credibility to deliver on our commitments to our customers. When you pair this with a recovering Traditional server market, impending PC refresh, and our updated storage portfolio, we are extremely well-positioned. Now over to Yvonne for the financials.

Yvonne McGill: Thanks, Jeff. Let me begin with an overview of our performance then I'll jump into the specifics of ISG, CSG, cash and guidance. In the third quarter, we delivered strong results including double-digit revenue growth while maintaining our focus on profitability. This resulted in EPS growth that outpaced revenue. Additionally, the business delivered strong cash performance. Our total revenue was up 10%, to \$24.4 billion. This was driven by continued strength in AI and Traditional servers. And our combined ISG and CSG business grew 13%. Gross margin was \$5.4 billion or 22.3% of revenue. This is down 140 basis points due to an increase in our AI optimized server mix and a more competitive pricing environment, specifically in CSG. Operating expense was down 2% to \$3.2 billion or 13.3% of

revenue. We remain focused on unlocking efficiencies within our cost structure and leveraging AI to reimagine our business and drive higher productivity. All of this while continuing to grow the business. Now, let's look at operating income. We delivered a 12% increase to \$2.2 billion or 9% of revenue. This was driven by higher revenue and lower operating expenses, which were partially offset by a decline in our gross margin rate. Finally, Q3 net income was up 11% to \$1.5 billion primarily driven by stronger operating income. And our diluted EPS was up 14% to \$2.15. Now let's move to ISG, where we delivered another quarter of strong performance. ISG revenue was \$11.4 billion up 34%. Servers and networking revenue was \$7.4 billion up 58%. This was a Q3 record for us. And we continue to see strong demand across both Traditional and AI servers. Storage revenue was up 4% to \$4 billion. The overall demand environment in Storage continues to lag that of Traditional server. However, as Jeff mentioned, we saw double-digit demand growth in PowerStore and PowerFlex, two of our marquee midrange offerings. We had strong ISG operating income of \$1.5 billion up 41%. This was driven primarily by higher revenue. Our ISG operating income rate was up again sequentially to 13.3% of revenue. The rate improvement of 230 basis points was the result of improving gross margins, especially in servers and reduced operating expense. Now, let's turn to CSG. CSG revenue was down 1% to \$12.1 billion. Commercial revenue was up 3% to \$10.1 billion, while Consumer revenue was down 18% to \$2 billion. CSG operating income was \$694 million or 5.7% of revenue. This is 50 basis points lower sequentially due to a more competitive pricing environment, primarily in the Consumer space. Enterprise demand was promising, though less-than-expected as we saw some demand push into future quarters. Profitability in the Commercial space held up well sequentially as customers continue to purchase more richly configured devices. Our Consumer business was weaker-than-expected as demand and profitability remain challenged. The PC refresh cycle is pushing into next year, but has significant tailwinds around an aging install base, AI-driven hardware enhancements like battery life and Windows 10 end-of-life. Shifting gears, Dell Financial Services originations were down 11% to \$1.6 billion in the third quarter. When you normalize for the exit of our VMware resale business and the sale of our Consumer revolving portfolio, DFS originations were up 4%, driven by strength in AI-optimized servers. Now, let's move to cash flow and the balance sheet. Q3 cash flow from operations was \$1.6 billion. This was primarily driven by profitability, partially offset by working capital. Our cash conversion cycle was negative 38 days with \$6.7 billion in inventory, which is elevated compared to historical levels due to our AI server business and AI order linearity. We ended the quarter with \$6.6 billion in cash and investments, up \$0.6 billion sequentially. This is the result of free cash flow generation partially offset by capital returns of \$0.7 billion. We repurchased 3.7 million shares of stock at an average price of \$107.53 and paid a dividend of \$0.45 cents per share. Since our capital return program began at the beginning of FY'23, we've returned \$9.8 billion to shareholders through stock repurchases and dividends. Finally, our core leverage ratio was flat sequentially at 1.4x. Now turning to guidance. The IT spending environment continues to be dynamic with certain areas experiencing faster growth than others. That said, our business model is advantaged. We have a broad portfolio of solutions and a go-to-market engine that's poised to capture these opportunities. We expect the strength in our server business to continue. However, enterprise and large customers are being more mindful of their PC and Storage IT spend in the short-term. Against that backdrop, we expect Q4 revenues to be between \$24 billion and \$25 billion up 10% at the midpoint of \$24.5 billion with ISG and CSG combined growing 13% at the midpoint. We expect the ISG revenue growth rate to be up mid-20s year-over-year driven by AI and Traditional servers and we expect the CSG revenue growth rate to be up low-single-digits year-over-year. OpEx is expected to decline mid-single-digits as we continue to drive efficiencies in the business. We expect the operating income rate to be up sequentially with continued improvement in ISG. We expect our Q4 diluted share count to be between 715 million and 719 million shares and our diluted non-GAAP EPS is expected to be \$2.50 plus or minus \$0.10 up 14% at the midpoint. At the midpoint of our full year guidance, revenue is expected to grow 9% and diluted non-GAAP EPS is expected to be \$7.81 up 10% year-over-year. Now turning to fiscal year 2026. We'll provide formal guidance in our Q4 call, but I know that many of you have questions about our view on next year. So let me walk you through our current thinking. We expect multiple tailwinds going into next year, including more robust AI demand supported by our strong five quarter pipeline. There's also an aging install base in both PCs and Traditional servers that are primed for a refresh. We expect ISG

growth to be driven primarily by AI servers followed by Traditional servers and storage. We expect CSG to grow as enterprise customers refresh a large and aging install base. We'll continue to balance growth and profitability, managing price within the competitive landscape and input cost environment. And we'll continue to invest in the business while driving cost efficiencies across the portfolio. In closing, I'm incredibly proud of the team and how we executed operationally. We remain focused on growth and profitability, always looking to optimize our margin rates and dollars. We have a unique operating model and strategy, where we leverage our strengths to extend our leadership positions and capture new growth. We have a strong control environment, differentiated cash conversion cycle, a strong balance sheet, financing capabilities with DFS, and investment grade ratings. All of these elements position us to win incremental business. Our confidence is underpinned by our 40 year history of success and durable cash flow generation. I'll echo Jeff. We are extremely well-positioned as we look to Q4, FY26 and beyond. We are winning in AI, Traditional servers are recovering, we are positioning Storage for the long-term, and Commercial PCs are stabilizing and on the cusp of a refresh. As always, we'll continue to provide differentiated solutions for our customers and drive shareholder value. Now I'll turn it back to Paul to begin Q&A.;

Paul Frantz: Thanks, Yvonne. Let's go to Q&A.; In order to ensure we get to as many of you as possible, please ask one concise question. Let's go to the first question.

Operator: We'll take our first question from Amit Daryanani with Evercore ISI.

Amit Daryanani: Good afternoon, everyone. Thanks for taking my question. I guess I just wanted to focus on the ISG margins. They came in at 13.3%, much better than I think what most folks were expecting. So can you just kind of dig into what the puts and takes where they enable this ISG margin performance that seems more impressed than what people were apprehensive about? And perhaps you just talked about the durability of these ISG margins as you go forward as well? Thank you.

Yvonne McGill: Thanks, Amit. I'll start on that. So ISG - let me start with revenue. Revenue grew 34% year-over-year, so an impressive performance with server and networking up 58% and storage up 4%. ISG operating income, like you said, was up 230 basis points quarter-over-quarter, which is up 530 basis points since Q1, so a nice recovery there as we had indicated. The Q3 quarter-on-quarter change was driven primarily by improving gross margins, so especially in servers, and a reduced operating expense percentage. Traditional servers demand improved and customers are buying more richly configured servers, which we've seen over the last few quarters. Units and ASPs grew with denser core counts and more memory and storage per server. So again, great to see. Customers focused on consolidation and power efficiency by modernizing their data centers, which frees up some volume -- or space for AI infrastructure. From an AI standpoint, AI servers as well as storage margins were flat to slightly up quarter-on-quarter, so that was again good to see. And we also realized more cost efficiencies quarter-on-quarter helping our operating income. Looking into Q4, expect ISG operating income rates to continue to improve.

Paul Frantz: Okay. Amit, thanks for the question.

Operator: Our next question comes from Toni Sacconaghi with Bernstein.

Toni Sacconaghi: Yes, thank you. I'm just wondering I think last quarter you were the midpoint of your guidance for the year was \$97 billion. Now the midpoint is \$96.1 billion. And I'm wondering what has changed in terms of coming up weaker than your expectations, especially in light of the fact that it sounds like the AI pipeline is really strong. I think last quarter you said it was multiples of backlog, which would have put it at \$11 billion to \$13 billion. And you said it grew 50% sequentially. So are we looking at a backlog that's a pipeline, excuse me, that's like \$16 billion or \$17 billion. And why is that not informing a more optimistic view for Q4 than what you had 90 days ago? Thank you.

Jeff Clarke: Toni, maybe a couple of things to think about what's changed in our thinking over the past 90 days and Yvonne going to build on this. First and foremost, as we tried to outlay in our remarks is the PC refresh is continues to move out. It was the source of our under run in Q3, and we're kind of reflecting that through the balance of the year that this refresh that we're all anticipating that we know is there with a large installed base and a Windows 10 end-of-life is coming, but the triggers have been slower to move. In many cases, late in starting and slow in refreshing, and we have a pent-up demand, we believe, that's pushing into calendar 2025 fiscal 2026. And the second element that we reflect in that annual guidance or specifically in Q4 is the shift in demand in AI to Blackwell. And it's the best reflection

we have is that demand is in our backlog now when we can ship it, and those are the two levers of why you see a change in our fiscal year and Q4 guidance. Yvonne?

Yvonne McGill: Yes. I'd say we did bring Q4 revenue guidance down, as you mentioned, Toni, and it's basically those two elements. PCs, it's not a matter of if the repurchase is going to happen, it's when, and we're seeing that move more into next year and then the unpredictability of the AI shipments. And so putting those two things together, we feel strong about the overall performance in Q4, but there's some more timing differences than what we were anticipating when we gave the guide the last quarter.

Paul Frantz: Thanks for the question, Toni.

Operator: Our next question comes from Aaron Rakers with Wells Fargo.

Aaron Rakers: Yes. Thank you for taking the question. Kind of sticking on the theme of the AI server dynamics. Jeff, I'm curious as you think about this several multiples AI pipeline, I think defined as five quarters in front of you, how do you think about the cadence of shipping Blackwell in volume? And do you have any thoughts around the GB200 NVL? How much of that would be as a percentage of your mix and the margin profile in those relative to maybe the Traditional or the AI service you ship to date? Thank you.

Jeff Clarke: Sure. And maybe a couple of things as we think about the five quarter backlog and bridging from Toni's question to yours. Look, it grew. We talked about several multiples. Without telling you the exact number, that multiple expanded quarter-over-quarter and grew in excess of 50%. The enterprise customer portion of that grew even faster. And remember, that's an unweighted pipeline. That's the full opportunity. Obviously, it's for us to convert in the marketplace, but it's a reflection of the growing demand. We saw it grow. We believe it continues to grow. We saw in Q3 a shift and a pretty rapid shift of the orders moving towards our Blackwell design, the specific one that you called out, the GB200. That's now on backlog. So the backlog in orders for the quarter, without giving you a specific percentage was a significant portion of the demand and now is a significant portion of the backlog. That's what we're trying to reflect in our guidance and went out the ability to get that material and then put it in the hands of our customers. You would have seen us 9 or 10 days ago make a post that we were the first to ship a GB200 design in the marketplace. We have them designed. We're ready to go. Blackwell is in production. It will continue to ramp and we'll be ready.

Paul Frantz: Thanks, Aaron.

Operator: Our next question comes from David Vogt with UBS.

David Vogt: Great. Thanks guys. Thanks for taking my question. So not to beat a dead horse in AI, I want to try a different topic. I don't think I heard you guys talk about federal spending or sovereign spending. Can you guys give us an update on kind of the impact that you're seeing in sort of the federal vertical? Is that having an impact on Traditional servers, storage and maybe on the margin commercial PC if it's captured in that category? And then how do you think about that getting resolved into 2025 as we move through a new administration handoff and what that might mean for the business? Thank you.

Jeff Clarke: Well, for the quarter, and Yvonne can build on this, as we saw the government, the federal government business was strong in Q3. The demand was strong across all types of products. With respect to the AI pipeline that I referred to, it does not have a sizable sovereign component to it. We're well aware of the opportunities. We're actively engaged across the globe in those opportunities, but that number does not reflect a significant sovereign number in it, which speaks to the opportunity that's still in front of us there. Yvonne?

Yvonne McGill: No, I think you hit on it, Jeff. The federal business, we did have a strong quarter within the third quarter and are continuing to see the pipeline build specifically there. So excited about that opportunity.

Paul Frantz: Thanks, David.

David Vogt: Thanks, guys.

Operator: Our next question comes from Erik Woodring with Morgan Stanley. Mr. Woodring, please remember to unmute.

Erik Woodring: Sorry about that. Thank you, guys. So, Jeff, earlier this year, there was a belief that AI server sales would pull forward some storage and services demand. And I'm just curious as you look at kind of the needs and evolution of your AI server customer base, whether this is still kind of an untapped storage and services opportunity that just hasn't materialized yet? And if it hasn't materialized

yet kind of why that would be the case? We'd love for you to double click on that. Thanks so much.

Jeff Clarke: Of course, Erik. Look, we're pretty excited about the opportunity in AI to expand beyond the individual node. I think we've tried to reflect that in our previous comments, but to maybe try to be very specific here, the opportunity is beyond the node into full rack scale integration. And in full rack scale integration, it's the networking opportunity, the storage opportunity, mundane things like cooling and how you actually build very efficient cooling subsystems to take the energy density out, how do we do power distribution, power management, putting telemetry in, doing power management, all of those are opportunities for us to expand our margins and why we believe we have a differentiated solution and ultimately are at a premium to our competitors. Then you add the service component to it, our ability to do L11 and L12, our ability to take these racks these integrated rack scale solutions and be able to design, support and service those in the field, deploy at scale, the financial services our DFS team can provide our customers in their early times and growing at incredible rates. That combination of opportunity exists for us in each and every one of these large opportunities and we're increasingly seeing the opportunity. Again, whether it's storage, whether it's networking, cabling, all of those matter, the design matters, and then being able to deploy it at speed is another characteristic of the Dell scale. Hope that helps, Erik.

Erik Woodring: Awesome. Thanks so much, Jeff. Yes, yes, no, that's perfect.

Jeff Clarke: Of course.

Paul Frantz: Thanks, Eric.

Operator: Our next question comes from Wamsi Mohan with Bank of America.

Wamsi Mohan: Hi. Yes, thank you. Last quarter just as a quick clarification, last quarter you noted AI revenues would be down quarter-on-quarter in 3Q. Can you talk about how we should think about Q4 as well? And my main question is, Yvonne, you mentioned an early look into fiscal '26 on CSG, where you're expecting revenues to be up. Can you talk about how you're thinking about margins also in that context and potentially any impact that you're contemplating from tariffs? Thank you.

Yvonne McGill: So I can talk a bit about the Q4 guide and we have taken into account in that guide the mix within AI and expecting, but expecting those margins, which I know we don't talk about holistically, but expect them to stay relatively consistent. So the mix would be what's driving there. We've talked about AI revenue and our offerings being margin dollar accretive and margin rate dilutive. So we'll continue to see that play out as we grow. But you've seen our results and we're continuing we had record or very high operating income in ISG, while again growing AI and growing across the portfolio. If I move into FY'26, which I think you're heading to, I think, it's a -- we're a bit early in the process to talk fully about FY'26. We kick off that process in full over the next few weeks, but very optimistic about where next year will head. We don't we do expect multiple tailwinds that support growth next year, so we feel positive about that. So more robust AI demand portfolio or opportunity with strong five quarter pipeline that Jeff already talked about going into the year, an aging install base for both PCs and Traditional servers, again opportunity there. We expect ISG growth driven primarily by AI servers followed then by Traditional servers and storage and we expect CSG to grow. And enterprise expansion refresh cycle that is coming and coming soon will help drive those the profitability and that revenue. So looking forward to next year, a little too soon for us to give a guide, but we will follow that up at the next earnings call for everyone and announce what we intend for next year.

Jeff Clarke: And then, Wamsi, your specific question about AI shipments in Q4, look, we've -- in our guidance, we've given the best reflection of what of that backlog that's on the GB200 clears in the quarter as it's now Blackwell is in production and ramping on our ability to translate that into shipments as well as the existing backlog on Hopper, it's our best indication of what we believe will shift in the quarter and that's what's reflected.

Wamsi Mohan: Yes. Great. Somewhat down sequentially.

Jeff Clarke: Pardon?

Wamsi Mohan: Somewhat down sequentially.

Yvonne McGill: Quietly down sequentially.

Jeff Clarke: Yes.

Paul Frantz: Thanks, Wamsi.

Operator: Our next question comes from Ben Reitzes with Melius Research.

Ben Reitzes: Hey, guys. Thanks. Question, I just wanted you to clarify with regard to your pipeline in AI. One of your competitors has had some well chronicled issues in the press and with their financials. And have you seen is that the reason that the pipeline surged and/or was that pipeline already baked? And are you seeing in November an uptick in interest in your solutions in AI? And then I was hoping, Yvonne, I don't think you clarified for next year what storage would grow or if it grows and you have a bunch of new products there. So some clarity on your highest margin product there would be appreciated among the other things you discussed about '26? Thanks so much.

Jeff Clarke: Sure, Ben. Let me give a start and then Yvonne can pick up where appropriate. Look, the opportunity in AI is immense. We see a formidable stack of competitors each and every opportunity. Again, each one of these opportunities are enormous in size and everybody's working to be competitive in those opportunities. Clearly, we are. We're showing up. We're showing up, we believe, with a very, very differentiated offer. That's why our pipeline is growing. That's why we believe we're having the success we have. That differentiation starts with co-architecting, innovation, the engineering, the technical stuff, these are highly complex solutions and being able to engineer in record speed. In the case of some of this, we're turning a design to a full scale deployment in less than five months, putting 50,000 GPUs in our customers' hands in five weeks. This is about rack scale solutions, driving networking, compute and storage altogether, doing the deployment and installation services that I described at scale. Our people, they'll badge people on-site 24/7 deploying, the professional service that go along with that, the financial service that goes along with that, which is also ultimately driving. We were the first to market with the new solution. Speed does matter. It is a differentiator. The pace of innovation is accelerating. We continue to see that. We're continuing to drive that and we're responding with more platforms. Recently, we broadened the portfolio. We continue to invest in capability. That includes our supply chain service capabilities. What we're doing on the go-to-market side with two different pursuit teams for the largest digital native opportunities as well as the enterprise and then, obviously, the R&D; side of that. Your second question, FY'26 while we're not giving guidance, we expect the storage marketplace to grow next year and we expect to take share with our Dell IP storage portfolio as we've invested into new solutions, new capabilities, making it more competitive.

Yvonne McGill: Yes. We expect growth across ISG, AI servers, Traditional servers, storage.

Ben Reitzes: Thank you.

Paul Frantz: Super. Thanks, Ben.

Operator: And our next question comes from Asiya Merchant with Citigroup.

Asiya Merchant: Great. Thanks for taking the question. We get a lot of questions from investors about the customer concentration in the AI backlog, orders, revenues. Jeff, I know you talked about the enterprise and you commented on that. So if you could just double click on, you know, what constitutes or what's the customer concentration in this backlog and, you know, the orders that you're seeing right now, that would be great. And then related to that, you guys have identified a TAM, a pretty substantial TAM of AI into calendar '25. And just given your competitive moats here and the fact that you're pricing to a premium, how do you guys feel about your market share in that \$120 billion plus TAM for calendar 2025? Thank you.

Jeff Clarke: Sure. Let me try to pull that apart. First of all, I think what's important to note, I think we've said that to be really clear, we have now sold over 2,000 enterprise customers since we began this journey. Our enterprise customers grew, continue to grow. The portion of the five quarter pipeline enterprise customers grew at a greater rate than the other side. Is our business still weighted towards the Tier 2 CSPs in building out those digital native platforms without question. But as we've, I think, said the last three quarters, enterprise continues to grow. It continues to be a larger portion of the opportunity for us. It's a larger portion of the pipeline, and we're only in the very, very early innings of enterprises figuring out how to deploy AI. They understand that it's highly disruptive. It drives higher levels of innovation, higher levels of productivity. They're in their experimentation, some of advanced proof-of-concept. And as they migrate through that, the opportunity is immense, in fact, I think a larger enterprise over the course of the decade when we think about inference and putting AI in production. Given our scale, given our go-to-market footprint, I like our chances there. It's what we do, it's where we sell most of our other gear and we think that's the alignment opportunity as well as forcing us to learn

and drive new levels of innovation with our largest customers today and that innovation actually makes its way to our enterprise solutions, which will make them more efficient and more capable. I hope that helped.

Operator: And our next question comes from Mehdi Hosseini from SIG.

Mehdi Hosseini: Yes, thanks for taking my question. Jeff, I want to get your thoughts as your key customer migrates AI platform from Hopper to Blackwell. What is your underlying assumption for the mix between HGX, DGX and MGX? Any thoughts there will be appreciated.

Jeff Clarke: That question was choppy on our end. Could you repeat that for us please? I apologize.

Mehdi Hosseini: Sure. I'll try again.

Jeff Clarke: It's much clearer.

Mehdi Hosseini: Sure. It's my cell phone. Sorry about that. As your key partner migrates from Hopper platform to Blackwell, how do you envision the change of architecture changing like HGX, DGX and MGX?

Jeff Clarke: Well, I think and the simplest way to look at that is Hopper is in production, will continue to be in production next year. Blackwell is ramping in production, will be in production next year. We have demand on both architectures. We continue to see that in the five quarter pipeline. Nothing's changing in that regard. I don't know how better way to describe it without giving you the specifics of the backlog, which I won't.

Mehdi Hosseini: But I'm just trying to understand. I'm trying to understand your assumption how customization of a new platform is changing like from HGX to DGX and MGX?

Jeff Clarke: Well, the way I look at it is we've designed a very open-based architecture. So if you look at our new IR7000 rack and take the XE9712 as an example. That architecture and the headroom that we've designed to 480 kilowatts being able to put the thermal solution and thermal headroom in place, the power distribution system in place. We believe we can take any combination of new designs based on the unique needs of our customers and design in short order and provide that at scale and speed to meet their needs. Much of this is driven by the model type, the technologies available, optimization and being able to help our customers weigh through that and then provide an open solution that allows us to move very quickly is exactly what we've built with our new ORv3 chassis, its ability to take things like we talked NVL72 and NVL4 design and subsequent new technologies that will come down the pipeline, I'm sure, in the future.

Mehdi Hosseini: Thank you.

Paul Frantz: Okay. Thanks, Mehdi.

Operator: Our next question comes from Krish Sankar with TD Cowen.

Krish Sankar: Yes. Hi, thanks for taking my question. Just you mentioned in the past about the advantage of having a full stack of server and storage as a unique value proposition to your customers. The question always comes up on attach rate. But is it the last five quarters you shipped about \$9 billion in AI server revenues? Is there a way to look at it and say how much was the attach rate you saw for your storage or services business based on the trailing five quarters? Thank you.

Jeff Clarke: I'm not sure I understand the context of the question. Is that AI specific? Is it general to ISG and infrastructure?

Krish Sankar: No, it's AI specific. I think you've spoken about how having the server and storage solution as an advantage. Is there a way to figure out that attach rate?

Jeff Clarke: Absolutely. So let me try to take a run at it because I think this is very important for us. I mean, I think I've said publicly multiple times that the AI opportunity for storage is immense simply because GPUs devour data. I mean, you have to feed the beast, and they're not very effective without a lot of information. Storage, it typically tends to be unstructured information. So we think scale out, file and object capabilities for training, tuning and inference are essential. We think parallel file systems for this kind of end-use transient data and these large training environments is essential. And remember, 80% of the data is on-prem. So we think AI is driving new needs in the storage architecture, which really drive to a three tier architecture. So the ability to scale, the ability to drive efficient deployment of storage, the ability to be flexible and above all high performance are all things required to meet these high performance modern AI workloads and that's what our portfolio is. The Dell IP portfolio is a three tier architecture moving towards disaggregated that allows us to scale CPU and storage and

networking independently to optimize for performance. All that said, our PowerScale platform, we believe, is absolutely the right platform where we deliver high performance file, soon high performance object, the ability to take our Project Lightning that we announced at Dell Technology World, a parallel file system built exactly and purposeful for AI and then take our Dell Data Lakehouse, which then allows us to do some data management and have the data platform around this and specifically metadata management, that combination of capability, we believe, is absolutely the future state of storage for AI and we're very well positioned. It's why we made a couple of references in our remarks about PowerScale, the F710 and the F910. One has 614 terabytes of capacity. If I remember correctly, the 910 is 1.46 petabytes of capacity. Both of those double as we go to higher density drives in the first half of next year that will be first to market. It's why I made reference to the performance attributes of both of those products are dramatically improved with the latest revision of software. We're not done, more to come. That's the opportunity and we believe we have the platform IP to be able to ultimately optimize our footprint in AI going forward. I know that was a mouthful, but that's how I think of it.

Krish Sankar: Thanks, Jeff.

Paul Frantz: Thanks, Krish.

Operator: Our next question comes from Michael Ng with Goldman Sachs.

Michael Ng: Hi. Good afternoon. Thank you for the question. I wanted to ask about the comments around AI server revenue being down sequentially next quarter despite the backlog growth this quarter. First, is that just because of the growing mix of Blackwell and component availability? And then second, in your view, is the backlog unusually high? And will we see AI server backlog flush at some point or is \$4.5 billion in the backlog normal with the potential to grow larger? Thank you.

Jeff Clarke: Again, to be clear, as the order shift in Q3 shifted dramatically towards Blackwell, that product is in its ramp and production, and obviously, our backlog reflects when we believe we're going to get the parts and be able to ship them, and we've translated that into our Q4 guidance, which so I can be clear, I wasn't clear as last time I said probably should have been. It's down slightly quarter-over-quarter in shipments from Q3 to Q4. The five quarter pipeline grew in excess of 50%. We're going to continue to work to convert that pipeline and I would expect us to continue to grow. I can't reflect it any better than that or say it any clearer. The opportunity for us is bigger this quarter than it was last quarter, in excess of 50% bigger, more in enterprise, new technology, new platforms, scalable platforms that are open that can meet the needs of today as well as what we think happens all next year, the headroom in that design, the investments that we've made in internal capabilities to service and provide financial products for these offers, the ability to engineer and innovate at speed are all things that are differentiating us in the market and we'll continue to focus on those.

Michael Ng: Great. Thank you, Jeff.

Paul Frantz: Thanks Mike.

Operator: Our next question comes from Samik Chatterjee with JPMorgan.

Samik Chatterjee: Hi. Thanks for taking my question. And Jeff if I can just change gears here a bit and talk about enterprise servers or sort of the non-AI servers. You've had double-digit growth here for a couple of quarters, strong demand. But I think there's a general investor perception that this is enterprise servers are a bit more of a cyclical business. And as you go into next year after the replacement demand, this year you might see a growth moderation on that front. If you can just sort of highlight what you're seeing with your customers? How do you feel comfortable with sustainability of these growth levels? And what's driving that into next year? That would be great. Thank you.

Jeff Clarke: Sure. You might recall, I think it was five or so quarters ago, we thought that the recovery in this IT cycle would be server-led. And the reason we thought that is we were in the middle of the longest digestion period in the history of the server marketplace. And we believe that coming out of this would be a faster recovery on servers, and that's exactly what we've seen. We're now on our sixth consecutive quarter of quarter-over-quarter growth. And as Yvonne said in her remarks, we're now at four quarters in a row of year-over-year growth, which tells me that we're seeing recovery and that digestion period is now in a period of consuming technology. I think what's very important about that is maybe the catalyst. The catalyst is a modernization. We think there's a modernization happening in both the traditional data center as well as that we could talk a lot about in modern workloads and AI. But in terms of the traditional data center, what we're seeing is a consolidation, a consolidation, as Yvonne

said so well to free up floor space and cooling space for AI. And what you get with a 16G server is you can consolidate anywhere from 3 to 5 sometimes 7 of our 14G servers into a single server. It's why we're seeing TRU expansion in our 16G, which is reflected in why our business was better in Q3. On average, a 16G server has twice as much memory and twice as much storage as a 15G server and we're seeing more workloads run on them. So that consolidation is well in play and you've probably seen some of the research in the marketplace, more CIOs are considering repatriation of applications today than they were a year ago, about twice as much. Whether that's driven from TCO, the fundamental economics, security or the fundamental performance and flexibility that you get, we're seeing that as the undertones of what's driving this demand cycle. And we are seeing unit growth. So unit growth, TRU expansion, consolidation, for all these reasons, we believe this consumption cycle, we're in the very early part of it and it is fuel to actually provide resources and capability towards the other side of modernization, which is the modern workloads.

Samik Chatterjee: Thank you.

Paul Frantz: Thanks, Samik.

Operator: And our next question comes from Simon Leopold with Raymond James.

Simon Leopold: Thank you very much. One quick clarification and then a question. You talked about expectations for OpEx to decline mid-single-digits. I just wanted to see whether you were referencing year-over-year or sequentials on that guidance. And in terms of the bigger picture question, I know it's too early to understand exactly what the tariff implications might be, but could you remind us of where your global manufacturing is and sort of what proportions are coming from what countries or regions? Thank you.

Yvonne McGill: So on to your first question, it was year-over-year with the answer there.

Jeff Clarke: And on tariffs, look, we've been at this a long time. We have a global supply chain that's built on resiliency and meeting customer demands. As a multinational company, we encounter tariff requirements in almost every country we operate and consistently navigate the geopolitical and macroeconomic cycles as they arise. We'll see what happens here. I think we're well-schooled and well capable of working through and navigating the changes that may be in front of us. We're not overreacting. We have a resilient and global supply chain, and we'll respond in the best interest of our customers and shareholders.

Paul Frantz: Thanks, Simon.

Operator: And our next question comes from Steven Fox with Fox Advisors.

Steven Fox: Hi. Good afternoon. I was just wondering if you could flesh out the storage business a little bit more, the underlying dynamics of the 4% growth. You mentioned some of the mid-range was growing double-digits. But if you could just give us a better sense of the rest of the sales growth and sort of what you're seeing margin wise and cost wise in that business? Thanks.

Jeff Clarke: Sure. Maybe a couple of comments and then Yvonne can certainly jump on in. So we talked about from a demand point of view the demand that we saw around our PowerStore and PowerFlex products. We talked specifically about the new PowerScale F710 and F910 growing. The focus of our organization and I think we conveyed this last time, which led to the improvement from Q1 to Q2 with our margin expansion. And storage is around the Dell IP products. So PowerMax at the high end, PowerStore in the mid-range, our new PowerStore Prime product is ramping up quite well, adding QLC to it. PowerScale, as I just talked about is an essential asset and capability that we want to drive and will continue to drive around the unstructured data side and AI. PowerProtect Data Domain, the new products that we launched, the 9410 and 9910 the ability to get those out as the preferred target for protecting these valuable workloads continues to be a source of growth and PowerFlex our software defined storage asset. We continue to see success with our partner first channel program. We're driving differentiation in the marketplace with our market leading data reduction rate capabilities and our broad coverage and specialty model serves us well that we have the largest sales force selling storage to more customers and we're continuing to focus on that. Those are the sources and levers we're leaning on to grow storage.

Yvonne McGill: And I'd add in the fourth quarter embedded within our guide, right, it's normally our strongest quarter from a storage perspective, so expecting growth there. And as Jeff said, we're leaning into that strong performance and leaning into our Dell IP, so helping to drive even more profitability.

Steven Fox: Okay. Thank you.

Paul Frantz: Thanks, Steven. Let's take one more and then we'll turn it over to Jeff to close it out.

Operator: We'll now take our final question from Tim Long from Barclays.

Tim Long: Thank you. Thanks for squeezing me in. Just wanted to get back to the AI servers and profitability. Obviously, when we went from CPU servers to multi GPU servers, it was a big jump in profit dollars gross profit dollars, but down in percentage. Just curious now that you're getting more involved in these NVL systems with 72 GPUs instead of 8. How should we think about the gross profit dollar move per server or the gross profit percentage? Obviously, the ASPs on the overall servers are going to be a lot higher with a lot higher GPU component within the bill of material, which is harder to mark up. So any comments you can give us on kind of how we should think about dollars and percentages?

Jeff Clarke: Maybe a couple of macro comments and then you can build upon that. One, certainly the ASPs of these high performance nodes with 8 GPUs in them and then shoved in a rack and driving densities 96 and beyond are very large. On a relative rate, Yvonne and I have been consistent, the rate is dilutive. I think what we've communicated now for three consecutive quarters is the rate has improved. What's fueling that rate improvement are many of the things that I tried to describe along the call today around services, the installation, the L11 and L12. What we're doing around DFS and the financing side of this, the opportunity around networking and storage and beyond. That's the opportunity for us to continue to improve the rates. It is a very competitive marketplace given the size of the opportunities, which is why we're leaning on our differentiation and our engineering capabilities to allow us to, we believe, perform to a market premium and we'll continue to focus on that. Yvonne?

Yvonne McGill: Jeff, I think you hit all those key points. Certainly, we're driving working to drive incremental operating margin dollars holistically and we are doing that within the holistic portfolio. So I think with that we will move to close.

Jeff Clarke: Yes. I'll give a quick couple of parting comments. We are well positioned to win with a unique operating model defined by our leading end-to-end solutions portfolio, large go-to-market engine, industry leading supply chain and unmatched global services footprint. And within AI, between our extensive portfolio and engineering expertise, we are providing comprehensive AI solutions for our customers from the largest at scale data centers to AI PCs, for enterprise customers to medium and small businesses. That's something no other company in the industry can do as well as us. And as you can see from our recent product launches, our innovation engine is strong, and we are only getting started. We've gained tremendous traction in AI solutions over the past year or so and there are no signs of it slowing down. The business model we've executed over the past 40 years was built for this moment. We have a solid track record of driving growth, profitability and cash flow, which all shareholders value. Thanks for your time today.

Operator: And ladies and gentlemen, this concludes today's conference call. We appreciate your participation. You may now disconnect at this time.