

TRANSCRIPT

DELL - Dell Technologies Inc at Goldman Sachs Communacopia &
Technology Conference

EVENT DATE/TIME: September 10, 2024 / 5:45 CST

CORPORATE PARTICIPANTS

Jeffrey Clarke Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

CONFERENCE CALL PARTICIPANTS

Mike Ng Goldman Sachs & Company, Inc. - Analyst

PRESENTATION

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

Great. Thank you, everybody, for joining the session. Welcome to the Dell Technologies keynote fireside chat at the Goldman Sachs Communacopia and Technology Conference. I have the privilege of introducing Jeff Clarke, who is the Vice Chairman and Chief Operating Officer at Dell. Jeff has been at Dell since 1987 and is responsible for running the company's day-to-day business operations and setting long-term strategy.

My name is Mike Ng, and I cover Dell and hardware here at Goldman. We have about 35 minutes for today's discussion.

So first, thank you so much for being here, Jeff. It's a pleasure to have you here. So I've heard that last month, you celebrated your 37th anniversary at Dell. So congratulations on that.

QUESTIONS AND ANSWERS

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

And I think as a result, you've had a front seat to some of the biggest technological changes that have happened, not only at the company, but also in the industry more broadly. I was wondering if you could first just talk about the technological shifts that we're seeing today, particularly as it relates to AI and whether you think that it's different, transformative, relative to what we may have seen in terms of technological shifts in the past?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

Sure. Maybe just a quick context. If you think about what's happened over the past many years, it took -- think about adoption rates, 50 years for all US households to get electricity. It was roughly 30 years to get 90% adoption rate of the Internet. It's roughly 35 years to get adoption of PCs to 90%, and roughly 25 years for phones to make it to 90%.

This one is really different. I've not seen anything move this fast. A data point, at least for us, a year ago in Q2, essentially zero percent revenue, and in the most recently reported Q2, 40% of our server and networking revenues were AI revenues. This thing is moving incredibly fast. In the last 12 months or four quarters, if you prefer, we've sold nearly \$9.5 billion of AI infrastructure, shipped \$6.5 billion of AI infrastructure. It is extraordinarily fast.

And if you think about it, I have made some bold predictions that by 2026, over half of data center demand will be AI. We look at two orders of magnitude in computational intensity growth over the remainder of the decade, so '27 quarter of ops. You got token growth at 151 times through the next handful of years. We haven't even talked about inference, and inference will be 90% of the AI workload by the end of the decade. There's nothing that's moved this fast.

And I get asked, why? And this is an inflection point. There's a fundamental change in the technology stack. This notion of accelerated computing is replacing a lot of manual work and human work with computers. It's driving new levels of productivity that haven't been seen since probably

the industrial revolution. There's not many gifts of 20% and 30% productivity that I can think of certainly in my working lifetime, and as a result, this is a conversation in every boardroom.

It's going to change the basis and it is changing the basis of competition. And if you don't have a strategy here, it's an existential threat. You will be out of business. That's why this is different.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

That's great. Could you talk a little bit about how Dell plays into these AI infrastructure investments that we're seeing across the industry today? Naturally, there's been a tremendous amount of investments by hyperscalers, but also, I'll call them, Tier-2 cloud, AI CSPs, enterprises, and sovereigns. If we were to think about those customer segments or whatever segment you think would be appropriate to address, where does Dell play best?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

Well, I think where Dell plays best is helping all of our customers, large and small, large cloud service providers, sovereign wealth funds, small businesses, deploy AI at scale, helping them develop a strategy, implement, ultimately adopt, and scale that. And the way we've communicated that broadly to the marketplace is through AI factories -- big ones, small ones.

Clearly, when we talk about the largest opportunities where there's a considerable amount of consumption of infrastructure and the largest CSPs at our customers, these are very bespoke designs. They're very custom designs, and what we've been doing is really driving differentiation in our offer. It's beyond the server. It's easy to say, I have a GPU; I put a GPU in a server, and I have one of these things.

These are very complex systems, and where we're distinguishing ourselves in the marketplace today and getting a premium for it is the engineering. This is a technical pursuit. It's not a sales pursuit; it's a technical pursuit. It's an architectural pursuit. And we are winning with, whether it's density, whether it's power, whether it's cooling, ultimately RAC integration and then the ability to deploy those RACs at scale in very short periods of time, I think if there's anything, these opportunities have really helped us with this -- really sharpen our edge. That we are much quicker than we used to be, more responsive than maybe you remember Dell being.

We are unbelievably responsive in these very large custom designs, and we're winning for it. So it's beyond the box. It's networking; it's storage. It's the services that encompass the solution. That is why we're winning, and that's very scalable to enterprise.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

That's great. And just on that point around the customization, the engineering value add. What is that really driven by? Is there some sort of gap in reference designs that are out there? Or are there specific workloads that your customers are seeking to address with AI infrastructure and Dell helps to accomplish?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

The reference designs that exist are very good starting points, but our customers want more density. So whether that started at 64 GPUs in Iraq, now at 72. We were the first to put 72, the first to 96. We're working on designs that are well beyond that, looking at what used to be a data center tile, if you will, of 10 to 15 kilowatts of power, now at 50 kilowatts of power, 75 kilowatts of power, now designing over 100 into next year at 200 and beyond that at 400 and beyond.

It's that engineering confidence that the reference designs allow us the latitude to design with as a basis. You think about what we've done with whether it's our 6U compared to our competitors. 8U, they're moving to 6U; we're moving to a 4U. We're getting more I/O slots in it, so we get

better performance. We get the GPUs in there. We're able to now put together a high-bandwidth fabric with low latency for the networking. That's where we're winning on the engineering side.

We started with air cooling. Others will talk about other forms of cooling. You don't need it until you get to certain energy densities. Typically, it's 1,000 watts, where you have enough energy density to need direct liquid cooling. The ability to engineer that solution at scale, that's what we're doing. As a former engineer -- they don't want me to call myself engineer anymore -- but as a former engineer, the ability to differentiate and to help these customers solve their specific problem is immense. And it's why, again, I think we're getting a premium in the marketplace for it and winning.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

Yeah. I want to ask about earnings, which just happened two weeks ago for Dell. Dell posted a quarter where EBIT margins for the ISG segment were able to improve nearly 300 basis points sequentially. AI server margins also grew. I was wondering if you could just give a post mortem on that piece. What's been driving the margin expansion within ISG? What's driving the margin expansion in AI servers?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

I mean for us, it was a very solid quarter, as you mentioned that highlights \$25 billion revenue, 9% growth, EPS growth of 9% of \$1.89. Free cash flow, \$1.3 billion, gave back \$1 billion of shareholder returns. And within that, you had ISG grow at 38%. And then that 38%, you had certainly what we talked about AI being a huge component of that. And it's great to see, and we improved margins.

Part of it is we're selling at a premium in the marketplace and extracting more value than just out of the box. Again, our ability to sell the services and deploy at scale, networking, and the storage subsystems that go around these integrated RAC solutions is incrementally allowing us to make better margins than we did the previous quarter.

On the storage side, it's a great story, at least from my seat. We sold more Dell IP storage. So our PowerMax product, our PowerStore product in the mid-range, our PowerScale, our file and object portfolio, our PowerProtect data domain product, all grew in demand double digits, all expanded their margins. So you take our most profitable products, they grow; they expand their margins in the most important region, in North America.

We do a little better managing discounting, and we find new value streams around data rate reduction, and you have a recipe of why our storage margins expanded. When storage margins do well, ISG does well. And then the third component of the ISG business is our server business, which continue -- or what we call the traditional server business, that continues to grow quite nicely. Five consecutive quarters of sequential growth, three consecutive quarters of year-over-year growth. We've got things hot over there in ISG.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

And I don't want to understate the strengthened storage because two quarters ago, storage was in a slightly tougher place -- more third party IP mix that was a bit of a drag on margins. Maybe you can talk about what changed in those few months.

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

Well, clearly, the challenges we had in Q1, you referenced correctly, is we had a mix towards partner IP at the expense of Dell IP, and that ratio changed from Q1 to Q2, which obviously helped. Look, I think we are -- a couple of things that we are navigating. One is a dynamic relationship with our partner IP, specifically VMware, and working through the changes in how the go-to-market is going to work.

There's no question, as we've navigated through that, it slowed the business a little bit. But I'd also tell you that I think the real driver is modern workloads are really favoring, what we call, three-tier architecture. For its performance, its scalability and efficiency, and we're seeing that in these

new modern workloads take off. So we're optimistic. I think in our guidance we reflected, we think our storage business is going to grow in the second half of the year. Standing by it.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

Great. That's great. And you talked about the traditional server strength, right? Five consecutive quarters of quarter-on-quarter growth. It's great to see in what still feels like a somewhat cautious IT spending environment. Are we at the point where you're ready to call an inflection? Are traditional servers set to return to more normalized growth from here? And is there any disruption from the AI investments that are happening in the market?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

It's a good question. I've been asked that all day today. I'm certainly not ready to sit here and say the markets recovered, and it's back to where it was pre this digestion period. But I think there are three signals from what we see that help communicate the markets recovering.

The first one would be, when we think about it is we went through the longest digestion period in the history of the server marketplace -- eight quarters. Data centers are full of older product. The older product, while capable when it was bought, on a relative basis today are very, very different in productivity capability.

So you have an aging of the data center that didn't get refreshed for a two-year cycle; that's never happened. And we're beginning to see large customers with large bids begin to refresh. And again, not calling recovered, but recovering, and we find that as a very encouraging signal.

The second component, which I think is somewhat related, as customers are looking at their AI plans or desires, they're quickly coming to the conclusion -- I need space and I need power. So we think a consolidation is occurring that you can consolidate older servers with newer servers and consume less space and consume less power. And we think that consolidation is, I think, a very important consideration with AI.

And here's why we think there is consolidation going. Our TRUs continue to expand. They're driven by more cores, more DRAM, and more NAND per servers sold.

I gave three quick references earlier today. I think it maybe -- for the broader audience, that would be helpful. Today, a 16G server, which is what we ship today, versus what we would have shipped four and five years ago, has three times more -- as many cores in it as that old product. It's 25% to 35% more power efficient, and it can roughly replace three to five old servers with a new server. So that's space, that's power that is now available for AI.

And then the third component that we think is driving traditional server demand is the continued repatriation of workloads back to on-prem into on-prem private clouds. So it's that backdrop that we think lends itself to the traditional server market recovering and continuing to do well.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

That's great. And as a follow-up to that, there certainly is server demand because of growing and new workloads, repatriation of workloads from the cloud. How do you think about this whole consolidation concept, one new server replacing three to five old ones? Are you indifferent to that, because the content for the new server is that much higher, so from a revenue and profitability perspective -- ?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

Yeah, Mike, I think the way that at least we look at it is the unit -- the number of units may not be what once ever it was, but the value of the unit is going up considerably higher. If you look at our ASP growth over the past handful of years, it's considerable in servers. And it's all driven by the

three things that I just mentioned: more cores, some more powerful PCs in every -- or more powerful microprocessors in every server, more DRAM around it, and more storage.

I'm fine with that. It's a good business. It's helping our customers consolidate workloads, be more efficient. It opens up more storage opportunities and provides the space and power for what is coming AI workloads to the enterprise.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

And if I could shift back to AI servers for a moment, Dell has clearly found a tremendous amount of momentum among AI CSPs, large cloud companies. Could you talk a little bit about where we are in enterprise and also sovereign investments in AI infrastructure? I think there's a bit of a debate whether or not those customers will materialize in a way to continue growth in the segment to the extent that there's any slowdown among like AI CSPs or Tier-2 cloud.

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

Yeah. I don't have the sentiment that it's not going to materialize in enterprise. In fact, just the opposite. There's no question AI is coming to the enterprise. One, it's where all the data is. And data is very expensive to move. And in many cases, data is proprietary, is unique, it's part of your business model, your value add, your secret sauce. It's not going to be transferred into other things. So data gravity is clearly driving AI, over time, to the enterprise.

And maybe another way to look at it is if you look at the most -- the five foundational models, their current rev, they've been trained on -- you pick your number 30, 40, 50 terabytes of data. The Dell company has hundreds of petabytes of data. And we're not unique. Let's just say it's all value-added data, for the moment. We're going to want to train, more importantly, we're going to do fine tuning and run inference on our data to make us serve our customers better. And every customer is going to go through that same sort of calculus.

They're going to try to understand what part of their data allows them to serve their customers better, produce their products and services better, serve their customers better in terms of services or end user services, and that's what's in front of us. So when we talk about enterprise demand -- I think I made reference in the earnings call -- the number of enterprise customers grew from Q1 to Q2. It was bigger than it was in Q4. The amount of revenue dollars coming in, in Q2 from enterprise customers was greater than Q3, which was greater than Q4.

And what we're trying to do with this five-quarter pipeline comment that we make is a reference of the opportunities we see in front of us. And that, I think I made reference to the number of customers, and that is growing, and the number of revenue dollars is growing.

So continuing to signal that we think enterprises are moving nowhere near as fast as these handful of customers buying large clusters, and they're not going to deploy it as large clusters. They're going to deploy it as a use module here and a usage model here and a use model there. We're seeing customers experiment. In some cases, experiment going to proof-of-concept and in many cases, going from proof-of-concept into production.

Five primary use cases really are driving what we see in enterprise. One is around cogeneration. Two is around agents and sales assistance. Three is around content creation, content automation. Four is around customer service, and the fifth is around supply chain. Those use cases are fairly universal across most companies, and it's where we see AI being deployed in time and where data is. That was a lot of them, so I hope that was helpful.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

Yeah, it was helpful. I love it. And in terms of enterprise interest, are there any industry verticals that are more front-footed in making these AI infrastructure investments than others?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

So if we look at customers and enterprise, there's a subset of this community that is very aggressive in pursuing AI, the quant traders. Well, they love this stuff. Every second matters, right? Every millisecond or nanosecond matters.

So I think about what's happening in national labs and deployment in national labs solving significant science problems. I think of pharmaceutical. I think of what's happening in healthcare, whether that's looking at radiology film and interpreting that, or back to pharmaceuticals looking at proteins and breaking down proteins to help solve new drug opportunities or healthcare opportunities. Industrials, manufacturing, clearly, are customers buying today, and then oil and gas. And then the institutions of higher learning, universities across the globe.

So a pretty good smattering. One not bigger than the other, but all moving in that proof of concept and finding use cases where they can deploy AI gear. And again, for us, it's not the box. These are very complicated systems to build what -- we use the word systems of helping them with the network fabric, the storage subsystem and the surrounding services to help them deploy. We use a reference of L11 to L13, which is essentially RAC scale integration testing in our facilities, transportation to our customer site, out of the truck, test it, roll it in, install it and then service it. That's, we think, very repeatable all the way through the enterprise into the small companies.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

And on that installation servicing maintenance piece, could you talk about how important services might be for an AI system or an AI server installation and contrast that with traditional? Is there a notable difference?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

Yeah. Maybe I'll give you a real life example. So we just introduced this notion of an integrated RAC scale system. And in that, let's take -- we talked about 96 GPUs per RAC. And you deploy that, and let's take a hypothetical 50,000 GPU cluster. It's 52 megawatts, 6,144 nodes, 4,000 switches, 150,000 network connections, 200,000 opticals, a couple of hundred miles of network cable, 1.5 miles of water pipe, 30 miles or so of rubber hose and other types of dissipation of power, and someone's got to hook that up. It's kind of hard.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

I'll take your word on that.

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

And so the reason I say this -- because you can't do one part of it or if you do, you need a lot of help. And what we've been doing is building our capabilities to be able to deploy, whether that is a 500 GPU cluster or a 50,000 GPU cluster, everything in between, the technical skills to do it. In fact, the pursuits themselves are very technical engineering-oriented.

I'm involved in all of those. It's fun again. It's really cool talking to customers about architecture, workload, design and then building it for them. And again, I think that's very applicable to what we'll see in the enterprise, helping it. And our job is to build AI factories that are easy to deploy to make AI easy for our customers.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

I want to go back to something that you mentioned earlier when we were talking about enterprise investments, which was storage, data gravity, the cost to move data. And I was wondering if you could just talk about whether there is some sort of difference in storage solution required for

AI workloads and AI infrastructure relative to traditional enterprise storage, and whether Dell has to do anything from a new product innovation perspective to accommodate that need?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

So if you look at the attributes of what's happening -- what we'll loosely define as modern workload, specific example being AI. These modern workloads AI are very demanding, and typically -- our bare metal deployments are typically container-based and/or multi-hypervisor based. So if you think about it, those are the requirements.

And then within that, there's requirements. Performance is a given. It's just -- these have to be high-performance storage subsystems They want flexibility. They want scalability. They want efficiency and ultimately want to be able to drive cost.

And when we look at that, for us, we believe that maps to our three-tier architecture. And the reason for that is because you can independently scale compute, networking, and storage, where a alternative architecture like HCI, really great for ease of use and ease of scale, but you can't independently scale networking, compute, and storage. They all scale together. In other words, if you need more storage, you have to buy more compute. If you need more compute, you have to buy more storage and so on. That's not the case with the independent nature of three-tier architecture.

Then specifically, what's happening is it's file and object. Most of what's occurring or what's being trained or what we see being trained in the future and ultimately run inference on is file and object -- scale-out file and object subsystems, like our PowerScale product. The other one in these very high-performance foundational model training is around parallel file systems, which is why in May, we made an announcement of Project Lightning that we're developing our own grounds up AI-based parallel file system.

Many that exist today came from out of the realm of HPC. I understand why that's the case, and they're solving a considerable part of the problem. We think the ability to design that with low latency, persistent memory, the ability to bring high throughput in there and then have many concurrent users to really deal with the transient end use that is done in parallel file systems is an advantage we'll bring to the table. So we're developing that. And then our file and object portfolio is very robust, and that's helping us win in the storage space inside AI. Does that help?

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

It's incredibly helpful. Thank you for that. I'd like to move on to the Client Solutions Group part of Dell's business. As we move towards the end of 2024, we're, I guess, three years from peak PC shipments, which was at the end of -- or is 2021. I guess where are we on the PC refresh at this stage? What's gone better than expected? What may have fallen a little short relative to what we may have thought at the beginning of the year?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

Well, at the beginning of the year, I thought we'd be talking about a PC refresh. And today, it's -- we're struggling to even use the word refresh. It hasn't happened. And I don't see it happening in the next weeks. I think it's pushing. Clearly, we've reflected on our guidance that it's pushing into the end of the year and into next year.

I remain very bullish about the opportunities around a PC refresh. You might ask why. Well, first of all, the installed base is old. It's never been older. It's large and old. The first COVID systems are four years and five months old now. And they were all Notebooks, because we sent everybody home that had to be mobile, and the average Notebook is a 3- to 3.5-year replacement cycle.

We're four-plus. So they're not designed to work four plus. So you have this aging component. It's more Notebooks, and it has to move.

The second thing is you have this forcing function called Windows 10 end of life. That's not moving. October 25, Windows 10 no longer. And we're one more quarter closer to it. So you have this wave that's mounting, if you will, so to speak, that has to transition. History tells us the more you

compress against the Windows retirement, the more that rebound happens almost instantaneously, because you have to clear to move out those systems.

I read some research the other day that said 61% of all Windows licenses in North America are still Windows 10. That's considerably more at the same point in time, 14 months ahead of the Windows 7 end of life, which tells you it's pushing, but it's also mounting.

And then the third component, again, as a former PC designer, man, there's some cool stuff coming. We're going to be NPUs in all of these things. We're going to see an application base develop here in short order that is underway in developing the greatest general-purpose device -- productivity device on the planet.

The PC is only going to get better. They're all going to have capable NPUs in them. When I'm sitting here doing something you can't do, you're going to go, what you have? And go mine has an NPU and yours doesn't, You're going to want what I have.

From a historical perspective in the PC, hardware has always led software. So it's not uncommon for us to have hardware capability in the product before the software wave happens. And I think that's what we're set up for here.

With a purpose-built accelerator and NPU in every PC, again, I've been bold, I said by the end of the decade, I think the installed base flips over and everything is going to be an AI PC. Got a little work to do between now and the end of the decade, but I'm excited about it.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

Well, there's a coiled spring of PC demand with October 2025 as well, right?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

Well, there's announcements last week by Intel of the next-generation processor coming with a more capable NPU. So it's continuing.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

Just while we're on the topic of PCs, Dell has leading ASPs across the industry, if the third-party industry market research is to be believed. What is driving that relative to the broader industry? Is it the premium consumer mix? Is it the commercial mix? Function of both?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

I think it's a little bit of both. So the way we look at it, our ASPs are roughly twice the industry average. Industry is 600 and change, and you do our math, it's 1,200 and change. So roughly 2x. What drives that? Clearly, a greater mix to commercial. Our commercial to consumer mix is 85-15, significantly higher than the rest of the industry.

In general, the way we sell with our direct sales force, we sell a higher mix. So in general, our CPU is a more capable CPU, higher resolution screen, more DRAM, more storage than the average PC, driving that higher -- or that delta between the two. I think our direct model of driving attach, docs, displays, services drives a differentiated margin in a differentiated ASP.

So that opportunity around business model, mix, selling technique, and a commitment to drive the things around the box is very important to us. And it's why you've probably seen recently, we've talked about, in forums like this, extending that PC reach to more peripherals. There's a \$200 opportunity around every unit for speakers, keyboard, mice, cameras, webcams, et cetera. We're trying to tap into that increasingly more so, because it drives greater value.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

In the last couple of minutes here, maybe I'll ask you to make a few closing remarks on where you see Dell heading over the next few years? Are there any specific themes or factors that you want investors to keep on top of their mind?

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

Well, again, I think if -- I'd like the audience to leave, look, we're disciplined operators. We're committed to consistent revenue growth, consistent profit growth. We run the company on free cash flow generation, and we've made a commitment of capital distribution back to our shareholders. We have a long-term value framework that we believe we can operate within, and we're committed to that. Again, we're disciplined operators.

I think what's lost sometimes is the fact that we can't apply what we do to other sectors. And I think AI is actually demonstrating that the Dell disciplined model, the operators that we believe we are, can absolutely apply to new growth categories. And it's -- and not only is it this -- you may think of as Dell being now 40 years old and maybe slow -- not at all. We're architecting significant GPU cluster design in handfults of weeks and responding to customers in days and winning, and winning with superior design, superior services.

I think about the four things that make us special is our large go-to-market presence, our supply chain, our R&D and innovation engine and our services, and we can apply that to broader categories, AI being an example. We're having success in 5G telco. The edge is another opportunity for us, and we're going to continue to do that. That's what we do.

Mike Ng - Goldman Sachs & Company, Inc. - Analyst

Jeff, that's a great way to cap off this session. Thank you so much. It's been a privilege to have you on stage.

Jeffrey Clarke - Dell Technologies Inc - Vice Chairman of the Board, Chief Operating Officer

Thanks for having me.
