

## Ciena (Leadership in Optical)

October 08, 2024

### Corporate Speakers:

- Gregg Lampf; Ciena; Vice President of Investor Relations
- Scott McFeely; Ciena; Executive Advisor
- David Rothenstein; Ciena; Senior Vice President and Chief Strategy Officer
- Brodie Gagel; Ciena; Senior Vice President of Global Products and Supply Chain

### Participants:

- Karl Ackerman; BNP Paribas; Analyst
- Simon Leopold; Raymond James; Analyst
- Ruben Roy; Stifel; Analyst
- Michael Fisher; Evercore ISI; Analyst
- Karan Juvekar; Morgan Stanley; Analyst
- Tim Savageaux; Northland Capital Markets; Analyst
- Adrienne Colby; Citi; Analyst
- Ryan Koontz; Needham; Analyst
- Michael Genovese; Rosenblatt; Analyst
- Cobb Sadler; Catamount; Analyst
- Dave Kang; B. Riley FBR; Analyst

## PRESENTATION

Operator^ Good day, and welcome to the Expanding Leadership in Optical: A Q&A on Ciena's Strategy for Growth in AI and Data Center Markets Conference Call. (Operator Instructions). Please note this event is being recorded. At this time, I'd like to turn the conference over to Gregg Lampf, Vice President of Investor Relations. Please go ahead.

Gregg Lampf^ Thank you, Allison. Welcome, everyone, and thanks for joining us today. We hope you had an opportunity to listen to the webcast we posted on Friday. If not, please go to our events page to do so at your convenience. We're excited to host this discussion about Ciena's strategy as we leverage our expanding optical leadership into market opportunities such as AI and data center.

Today, we'll focus on your questions for up to 45 minutes. Scott McFeely, Executive Advisor; David Rothenstein, Chief Strategy Officer; and Brodie Gage, SVP of Global Products and Supply Chain, will participate in the discussion.

Scott will take a couple of minutes to provide a high level review of the presentation before we kick off the Q&A. Before handing it over to Scott, I want to remind everyone that the focus of this call is on the aforementioned webinar entitled Expanding Ciena's

Optical Leadership: Our Strategy for Growth in AI and Data Center. We will not be discussing financial information, results, expectations, or economic conditions.

While we'll not be discussing these items, we do expect to talk directionally about the technology market. As such, please take note of our safe harbor language from our most recent financial results call on September 4th, 2024, as well as our Form 10-Q filed with the SEC on the same date. With that, I'll hand it over to Scott.

Scott McFeely^ Thanks, Gregg. Hello everyone, and thanks for your interest in Ciena today. As Gregg said, we hope everyone had a chance to go through the voice over presentation that we posted a couple of days ago, but let me just briefly emphasize some of the key points in that presentation, starting with our overall investment strategy. For some time now, our strategy has been to invest in our optical foundation and to expand our business from that point of strength.

That strategy hasn't changed for several seasons, and it's presenting us with some very exciting opportunities as the optical networking space becomes even more important to our customers.

Optical bandwidth demand, as many of you have heard us and others say, has been growing on average at approximately 25% to 30% for many years, driven by things like the adoption of video and cloud networking. Those drivers continue, and although it's still difficult to quantify, we are starting to see real signs of AI impact on the demand in that wide area network.

Now, as we think about our customer base that serves that demand, we think about them in two broad categories. Cloud providers, and we participate with them, of course, connecting their data centers in many ways; and the service providers, that are providing the conductivity from end users and enterprises to get to the cloud, in many cases providing capacity to cloud providers as part of their solution set to interconnect those data centers through managed fiber optical networks, or, as you've heard us refer to MOFN networks.

Now, our relationship and incumbencies with both of these customer bases have never been stronger.

But strength is anchored in our leadership in three key attributes of our optical foundation, in particular our optical modems, where we have seen recently that we've extended that leadership with the introduction of WaveLogic 6.

Secondly, our optical line systems. These are becoming increasingly important to our customers' infrastructure, and we have seen significant uptake in our Reconfigurable Line System designed from the ground up in collaboration with a number of large cloud providers. And that system is now the de facto standard in all new line system deployments. And our market leading multi-layered domain controller to assist our customers in getting the most out of their optical and converged IP optical infrastructures.

This portfolio strength and the incumbency position, we feel very confident that we will continue to grow our optical system business faster than our historical market growth rates and will continue to take share. In addition, growth in our optical systems we see an additive opportunity for coherent plugs, particularly for deployments in and around the campus, and metro DCI use cases as a place where we historically haven't had much footprint.

Recently we've had a great win rate with our plugs, and we expect that success to continue, and we are starting to see meaningful revenue from these wins in our second half 2024 fiscal year.

Finally, as we think about expanding beyond our traditional strength in the optical transport market, we're looking for opportunities that will have a large market spend, preferably with market players, which we already have strong relationships with, that presents a discontinuity to allow the challenger to take share, and that discontinuity is focused around optical, and that allows us to bring the strength of our optical foundation as the basis to compete in those opportunities.

There are three such opportunities that we see in front of us right now that meet that criteria, two of which we've been speaking of for some time and having part of our expanded TAM discussions that we have with this community and are factored into any dialogues that we've had with you around long-term growth rates. Those are metro routing and next generation broadband access.

The third is a more recent topic, but getting a lot of attention, and that is the opportunity for coherent technologies to find their way inside the data center. The potential here is very exciting, but it's early days. Accordingly, this opportunity hasn't been factored into any of the long-term growth targets that we talked to you about. With that as a brief intro, operator, we'll now move to the Q&A part of the call.

## QUESTIONS AND ANSWERS

Operator^ (Operator Instructions). Our first question today will come from Karl Ackerman of BNP Paribas. Please go ahead.

Karl Ackerman^ Yes, thank you, gentlemen. I was hoping you could discuss what is the typical cadence for installed base upgrades for your WaveLogic offerings? I asked because you had previously indicated you still ship a fair amount of WaveLogic 3 today. But clearly, WaveLogic 6 Extreme and Nano appear differentiated to address current and future needs of higher networking bandwidth.

So if you could address perhaps the upgrade path or cycle you see of WaveLogic 6 relative to previous WaveLogic offerings, that would be very helpful.

David Rothenstein^ Brodie, do you want to take that one?

Brodie Gage^ So just to confirm the question, Karl, it's the installed base of platforms that WaveLogic 6 can take advantage of?

Karl Ackerman^ That is correct.

Brodie Gage^ Yes. So there's the installed base of systems that WaveLogic 6 can take advantage of. And probably more importantly, is the installed base of photonic line systems that are out in the market today that can handle WaveLogic 6, meaning that they're flexible grid and can support the higher baud rate, the 200 gig baud rate. That installed base, I would say, is approximately \$6 billion to \$7 billion at this point in time.

So there's a substantial installed base of photonic line systems that can handle new coherent modems, WaveLogic 5e, WaveLogic 5n, WaveLogic 6e and WaveLogic 6n going forward. And that installed base would be in the largest service providers in the world as well as the largest cloud providers in the world.

Karl Ackerman^ Yes. Yes. Thank you for that. If I may sneak in another one, one of the main investor questions we've received recently is the adoption of MOFN or managed optical fiber networks.

And so I guess the question is, if data centers are increasingly dictating the architectures of managed optical fiber networks, is this an opportunity or a risk for Ciena, given the company's rather limited exposure today to metro DCI applications, but higher exposure to metro applications and telecom networks? If you could highlight the opportunities and challenges that you see both in metro as well as long haul from the adoption of managed optical fiber networks, that would be very helpful. Thank you.

David Rothenstein^ Yes, Karl, this is David. Let me start. And Scott or Brodie can jump in. Overall, as you heard me talk about a little bit on the recording, we look at MOFN or the managed optical fiber, and its twin I suppose, the dark fiber conduits, as a real interesting opportunity for us, and we don't necessarily delineate between long haul and metro in that regard, right?

The reality is, as the connectivity demand increases, whether it's in the WAN or inside, you start seeing cloud providers choosing to outsource the WAN connectivity, either because they can't own and operate the fiber due to regulatory limitations or from a pure capacity or time to market standpoint, they don't want to.

And so we absolutely have seen, and I use Lumen as a couple of examples, there you've seen recently where service providers are really now providing and operating managed optical fiber networks on the cloud provider's behalf, or even selling long-term rights to the unlit fiber. So it's a really interesting example of how the cloud provider connectivity requirements, you're absolutely right, it is their connectivity requirements being driven by DCI architectures are really starting and continuing to impact service provider behaviors.

So we look at this as a really interesting, important opportunity, even at a time where overall, especially internationally, service provider spend is obviously not returned to historical levels. Scott, Brodie, do you want to add anything on that?

Brodie Gage^ Yes. The other thing I would add is traditionally the hyperscalers would either build their own network or they would rely on a service provider where they could not build out their network, places like India or Asia, et cetera, where they might not have the regulatory ability to build out the network or the desire to be able to build it out.

What we're seeing is not only are those large cloud providers, the Googles, the Metas, the Microsofts, the Amazons of the world requiring MOFN connectivity, but we're also seeing Oracle, xAI, Akamai, Tesla, what we refer to the tier two cloud providers also looking for connectivity. And where the service providers used to use a portion of their network to build out the network for these hyperscalers, they're now building out dedicated networks for those cloud providers and it's absolutely a net upside opportunity to Ciena.

Karl Ackerman^ Thank you.

Gregg Lampf^ Thanks, Karl. We'll take the next question, Allison [ph].

Operator^ Our next question today will come from Simon Leopold of Raymond James. Please go ahead.

Simon Leopold^ Thank you very much for doing this and taking the question. I wanted to refer to your Slide 22, I think, which is the one in the deck that compares pluggables to performance systems. And what I was hoping you could help us with is understanding some of the underlying assumptions for either cost or architecture of how you arrive at what's presented in this slide, as well as what's the distribution of how much of the network is at the different lengths you describe in that particular slide. And then I've got a follow up.

Brodie Gage^ I'll take the first part of it, Scott and David. So they're in the chart package. There were -- it was a chart titled All Links are Not Equal. And there were three buckets or three graphs, a metro graph at 500 kilometers, regional graph at 1,300, and a long haul graph at 2,700 kilometers. And each of the graphs has a pluggable capacity and cost and a performance capacity and cost.

And what we were comparing is our vertically integrated cost for a full solution to build out a pluggable network inclusive of fiber, photonics, and transponder shelves with those pluggable optics versus our next generation WaveLogic 6e performance modem. So like for like next generation technology.

And in metro applications, you can see that there's a 30% cost benefit to pluggable coherent type architectures, but there's a 40% negative impact on fiber capacity associated with that pluggable solution.

In the middle of the regional network, the cost is basically the same between a coherent pluggable and a high performance network or a high performance modem. But the capacity is a 30% degradation on pluggable, so you get less utilization of your fiber. And in long haul, the high performance optics were about 40% less expensive than pluggable optics. And the capacity that you can get on the fiber is 2x.

And so the voice over I would give is that we break our business down into cloud providers, the Metas, the Googles, Amazon, Microsoft, Oracle, et cetera of the world. Over 90% of our business is on the regional and long haul portion of these graphs. And every one of those large cloud providers are taking and standardizing on our next generation high performance modem because they want to maximize the capacity of the fiber and the cost delta for regional, long haul and submarine applications and such that it makes sense for them to deploy performance modems.

On the left hand side of this chart where pluggables do make a lot of sense, if you have the fiber in the metro or metro DCI type use cases. This is a pretty new business for Ciena and less than 10% of our cloud provider business is in this segment and that's why we have a net upside opportunity associated with coherent pluggables.

Simon Leopold^ Great. And then just as the follow up, I wanted to ask about this newer opportunity around coherent technology being used inside data centers. I wonder what's your current thinking or best guess as to the timing of when this technology becomes cost effective and necessary.

And then how are you thinking about really the economics because today coherent is much more expensive than direct detect. So how do we get to that crossover where the economics make sense? Where do the prices need to be? Thank you.

Brodie Gage^ Yes. So maybe I'll take it and then David and Scott, you guys can jump in. We see the in and around the data center opportunity kind of in three applications, and you can see it in the first chart in Dino's section of that presentation. There's the Metro DCI, which is sub 100 kilometers. Today it's 400 gig ZR moving to 800 gig ZR.

We are just starting to deploy coherent in that space and we have, I think three of the top four cloud providers choose us for 400 gig, and we got the only win on 800 gig for that use case. So that is absolutely going to coherent moving away from COLORZ and the PAM-4 implementations that were there just a couple of years ago.

With regards to inside the data center, we really see two applications. There's a campus application which is typically 2 kilometers to 20 kilometers, and then actually inside the data center where you're connecting from rack to rack or rack to optical cross connects.

With regards to the campus application, we believe that is going to start at 1.6T. And if you looked at what we announced at ECOC, we announced a coherent light 1.6T

solution. It's really going after that campus application and maybe some of the intra data center applications that require optical cross connects.

Google deploys optical cross connects. That adds loss and latency in the network and coherent is a good implementation for that. So number one, we see it starting in the campus first and then moving down deeper inside the data center. There is a cost and power equation to it, but there's also just, we believe IMDD is going to run out of gas. So you can debate the distance. Some say 6 kilometers, some say 8 kilometers. But 1.6T IMDD runs out of gas and that's where coherent light will have a play.

And then with regards to the power, we believe we can get the power of a coherent light solution down to the same power as a PAM-4 1.6T solution. And with volume, we believe we can get the cost points down into that level as well.

With regards to timing, I think it's going to take time. There's going to be a battle in front of us. I think the campus will happen sometime, maybe in 2026, and then inside the data center a little bit later than that. I'll ask David and Scott to comment as well.

David Rothenstein^ Yes, I think that was good. Brodie. I would just maybe, Simon, add one or two other things on top of what Brodie just articulated, the first of which, when relevant to Karl's earlier point, we are still selling a fair number of WaveLogic 3 cards, and if you apply that same dynamic inside and around the data center, this is going to take time.

It's not as though a switch is going to flip on a certain date and coherent will completely subsume direct detect in and around the data center. It will depend upon a number of different things. Obviously, the particular operators network architecture as well as obviously their end user needs, right? So our view is that there will be over time, a market for both direct detect and coherent.

But as Brodie points out, at the higher data rates, the higher flow rates, we do believe that the coherent optical does have significant benefits in terms of performance specs over direct detect.

Simon Leopold^ Thank you very much.

David Rothenstein^ Scott said it is that this is a net incremental opportunity for us. This is not in our 6% to 8% long term growth rate that we've talked about.

Gregg Lampf^ Great. Thank you. Operator, take the next question.

Operator^ Our next question today. Sure. And our next question is from Ruben Roy of Stifel. Please go ahead.

Ruben Roy^ Thank you. Dave, I guess just to follow up on that last point you made and thinking through what's incremental here, the slide presentations, which by the way are

very helpful, you did talk about starting to hear about potential new applications, AI inference, potentially moving closer to the end user, closer to the edge, with clusters moving closer, and obviously the disaggregation of the data center.

Just so I get this right, are you saying today that the incremental opportunity to sort of the way you guys think about the longer term CAGR is relegated to inside and around the data center? Or are you starting to hear about potential incremental opportunities with some of these other WAN type of opportunities? You did mention outsourced WAN as an opportunity. Just trying to get my arms around what's incremental and what's not.

David Rothenstein^ Yes. Thanks, Ruben. It's a fair question. Let me give it a shot. So one of the things that I talk about in the presentation is the hypothesis that if you kind of take cloud adoption and AI writ large and the relative potential impacts on network traffic flows, you could see a world where it actually takes the annualized bandwidth growth above what Scott started out by saying, the annualized 25% to 30% that we've seen now here for quite some time.

To the extent that it bolsters that 25% to 30% bandwidth growth, I think you incredibly said that that would be encompassed within our current view about long-term growth rates. However, if you ascribe to the notion that it could take bandwidth growth and the underlying data generation fueling it could take it above those historical, that would be incremental on top for both optical systems and other aspects of our business. That is still an unknown, right? It's a hypothesis that is going to have to be tested.

In terms of what elements of the portfolio are included in that kind of assumptive growth rate, Ruben, really the way to think about it, if you think about, for example, in terms of plugs, right, the plugs that we sell into our own systems or into third party switches and routers, as we've been doing for a while, that you can credibly assume that that is encompassed within the growth rate.

However, some of the applications that Brodie's speaking of here around the metro campus, 2 kilometer to 20 kilometer, and certainly taking any of the component technology inside kind of the less than 2 kilometers, interact, interpod and so on and so forth, those would be net incremental opportunities on top of that historical growth rate.

Ruben Roy^ That's really helpful, Dave. Thank you for that. And then as a follow up, Dino also mentioned an expansion opportunity being potential component sales. So standalone DSPs, maybe some mixed signal high speed analog sales into the marketplace. Are any of those types of discussions happening today with potential end customers? Or when you say expansion, is that sort of further out?

David Rothenstein^ Let me give it a shot and ask Brodie or Scott to jump in. I would say the discussions are absolutely happening right now with a combination of players in the overall ecosystem, and that would include, obviously the cloud provider customers who we have deep customer engagements with on a variety of different elements of their



networks as well as other players, whether it be the GPU players, systems, integrators, and others.

Everyone is kind of looking around at the potential opportunities, applications and use cases here, and we are very much involved in those discussions. What I would say is we have not obviously publicly announced anything yet, and wouldn't and won't until something really becomes crystallized.

However, in terms of the applicability of those foundational optical technologies, as you disaggregate the modem into an ASIC, or even potentially into composable chiplets, or even the underlying intellectual property we believe does have legs, we're just not exactly sure how that's going to instantiate at this time.

Ruben Roy^ Thank you very much.

Scott McFeely^ The only thing I'd add, David, is if you look at how our customers consume our capabilities, the deeper you get into the WAN, the more likely it is that they're going to consume end to end systems, the closer you get into that metro DCI. Think of it as package plugs. We're fully aware as we step inside the data center, the consumption model is there, and the architectures are going to dictate that we're going to have to open up and self-components are IPR, and we're absolutely willing to go there.

Gregg Lampf^ Great. We'll take the next question, Allison [ph].

Operator^ Thank you. Our next question will come from Amit Daryanani of Evercore ISI. Please go ahead.

Michael Fisher^ Thanks. This is Michael Fisher on for Amit. I just wanted to talk a little bit more on the Lumen Microsoft deal. And knowing we can't probably get into any specific details, maybe just broadly about MOFN deals. I'm trying to understand something like this or something of this size, what's the incremental amount of optical equipment or types of optical equipment that Lumen will need to purchase to support this kind of deal, or just anything you can help that can help us understand the magnitude that these sort of deals can have in terms of PAM expansion.

David Rothenstein^ Brodie, you want to start?

Brodie Gage^ Yes, sure. I mean, it's very difficult to say, to be honest with you. It's dependent on reach, and it's dependent on how much capacity gets put onto these networks.

So we have MOFN deals that could be as small as \$5 million for each opportunity. We have MOFN deals that could be across the entire United States with hundreds of 400 gig wavelengths, which would be \$40 million, \$50 million for that one MOFN opportunity. So the range of the different MOFN opportunities, again, is dependent on the reach and dependent on how much capacity that they're going to be putting on that network.

David Rothenstein^ And what we have, Michael, what we have said with respect to MOFN, we haven't specifically quantified the size of that opportunity.

But we have said today, at least, if you think about what cloud providers are driving indirectly, beyond, of course, the direct DCI business, which is about 30%, 35% of our current revenues, you'd look at MOFN combined with subsea, and we've said the total of those indirect is another 10% to 15% of revenue, which is why we've said the combined direct and indirect impact of cloud providers on our revenues is around 50% of total. So that's what we've quantified so far.

Michael Fisher^ And then just to follow up on that kind of 10% to 15% of revenue, is that if we go back, call it three, four years, is that a much smaller figure there for MOFN and subsea?

David Rothenstein^ It's hard to do an apples to apples. I know where you're going, Michael. It's hard to do an apples to apples comparison. In terms of MOFN, the way we define it today, it obviously is significantly greater. However, there was kind of an earlier version of that which we used to call CMS or carrier managed services, where you had a number of tier one service providers doing an element of CMS into other end users.

So it's hard to do an apples to apples comparison. But MOFN, as we talk about it today, really didn't exist three, four years ago.

Scott McFeely^ I think there's a chart in the deck where we walk through the growth of the cloud provider spend in direct terms as a portion of the broader industry spend. And that's been increasing substantially year-over-year for the last decade. I think there's a corresponding growth rate in the indirect piece as well.

And you can probably see it most directly in the submarine market where if you go back ten years, that market was dominated by service provider consortiums that were providing bandwidth to various different constituents. The weight in that market is absolutely shifted to the cloud providers as well, and they are either controlling the new cables or consuming most of the capacity on those cables.

David Rothenstein^ Great. Thanks for taking my questions.

Operator^ Our next question will come from Meta Marshall of Morgan Stanley. Please go ahead.

Karan Juvekar^ Hi, this is Karan Juvekar on for Meta. Thanks for the question. I guess just generally how much variability are you seeing in the approaches that cloud vendors are taking? And when you're looking at the adoption of coherent at shorter distances, do you imagine some of that adoption will be highly variable by the cloud vendor?

David Rothenstein^ Let me give a high level view and Scott or Brodie can follow up. I think you've heard me say this before. I do think we, and I said the royal we, Ciena included, as well as the Street, tends to lump customer segments in as a homogeneous unit, right?

And that's true of the cloud providers, even the big four that we talk about are you very different businesses with different business models, different network architecture approaches, different allocation of CapEx at any given time. And it does result in ebb and flows and different applications. So in terms of variability, higher than you might expect even within some of these domains, like shorter reach, metro DCI.

That said, we are continuing to engage with all of them very heavily. You've seen us take down three of the four with 400 ZR wins over the past couple of quarters. And we also have, to our knowledge, the only 800 ZR win with a major cloud provider. So there is variability across them. At the same time, make no mistake, they are all looking at what each other is doing. It is resulting in a hyper competitive arms race.

And you're seeing it, for example, with each of them taking up their CapEx not once but twice materially within even this year, I think the collective is over \$15 billion from the already high CapEx baseline that they announced at the beginning of the year. So it's variable, but at the same time there's some degree of consistency in terms of watching what others are doing and their approaches.

Karan Juvekar^ Got it. Thank you.

Operator^ Our next question today will come from Tim Savageaux of Northland Capital Markets. Please go ahead.

Tim Savageaux^ Hey, thanks a lot. Kind of want to go back to this Lumen TAM discussion, and you've mentioned it a couple times here in the slide. They're purchasing 10% of Corning's fiber capacity every year. That's more than \$500 million.

I wonder how you would view an equipment opportunity relative to that number. But I really wanted to get to a broader competitive discussion. Lumen historically been a big Infinera account. Clearly the merger with Nokia may have competitive implications. They just announced what looks like a pretty big deal with CoreWeave. That appears significant.

So a couple of things on that. Maybe to go a little bit deeper on the TAM and other Lumens out there, I think is the question most folks want to know. And then how Ciena is positioned both specifically there and elsewhere, given the changing competitive environment. Thanks.

David Rothenstein^ Yes. Hey, Tim, thanks. Let me give it a shot. And again, Scott or Brodie, please jump in.

Let me start with Lumen, right? Yes. Obviously, a significant long-term agreement with Corning for next gen fiber optic cable supply. And then as we talked about a significant long-term construct with Microsoft to provide the backbone connectivity. So super, super big deal for Lumen. I think they're really leaning hard into not just those opportunities, but others.

We feel very good about our positioning with Lumen. We've been a longtime supplier to them and we continue to engage strongly with them. We've not quantified that opportunity for us and won't today.

But from a competitive, dynamic standpoint, we believe we will get our fair share, or more than our fair share of the resulting opportunities arising from those MOFN and dark fiber deals from Lumen.

In terms of broader competitive point that you were raising, look, what we've said about Nokia and Infinera is a couple of things, right? We've been competing and competing effectively, we think, against each of them independently for a long period of time.

And certainly we are seeing, and believe we will continue to see near-term opportunities, certainly before that deal closes and after it, as they assuming it closes, they look at things like rationalizing portfolio, handling account support, customer engagement and things like that, and making some decisions in that regard. And we have seen already opportunities coming our way as a result of that.

That said, if they're able to execute successfully on that deal, that's a big if. But if they are, they will be an incrementally better resourced competitor from layer zero to layer three in the network. And so we don't take that lightly and we intend to continue competing aggressively against them going forward.

Last point you asked is, are there other Lumens out there? I think, look, in terms of direct applicability to the earlier question about variability, maybe not direct, but we absolutely see large-scale MOFN opportunities. India for example, where it's one of those places that from a regulatory standpoint, the cloud providers cannot own and operate fiber there.

So we absolutely see MOFN opportunities through Jio and Bharti being significant going forward. Other locations, ASEAN, Southeast Asia, similar dynamics. We see those being real, meaningful MOFN opportunities for us. I think that covers that.

Brodie Gage^ We see all. Yes, we see on the MOFN opportunities, we see all major operators within a country have visibility to those. And the cloud providers are looking for diversity both in their fiber routes and the service providers that they do business with.

Tim Savageaux^ Great. Thanks very much. Appreciate it.

Operator^ Our next question will come from Adrienne Colby of Citi. Please go ahead.

Adrienne Colby^ Hi, it's Adrienne for Atif. Thanks for taking the question. On the pluggable side, you mentioned having over 100 customers now for your 400 ZR, ZR+ solutions and the first 800 ZR win. I was hoping you could update us on the demand you're seeing for 400 ZR evolving with the availability of 800 ZR, ZR+.

David Rothenstein^ Brodie, Scott.

Brodie Gage^ I believe -- yes, I'll take that one. With regards to 400 gig ZR and ZR+, I think service provider customers and cloud provider customers will be deploying that for the next 5 to 10 years at a minimum.

The cloud providers typically have a quicker turnaround on technology than the service providers, but expect to have a long tail with regards to that deployment, even with the introduction of 800 gig ZR ZR+.

Scott McFeely^ Yes, it's important to also note that to get to 800 gig ZR, they do require the upgrade of their switching infrastructure, which will take some time. So 400 gig will have a long tail to it, if for only that reason alone, to help dimension it a little bit. As we sit here in year now in '24, we'll have a few tens of millions of dollars of revenue and sort of pluggables around that 400 gig ZR5 nano solution, and we expect that to grow going into '25.

Gregg Lampf^ Okay, thank you. We'll take the next question please, Allison [ph].

Operator^ Our next question will come from Ryan Koontz of Needham. Please go ahead.

Ryan Koontz^ Great. Thanks for the question. Just a quick clarification on your comments around the edge inferencing opportunity certainly makes sense logically. But I wondered, in your discussions with customers, are you seeing any hard requirements for that yet, or is this something that is still many years out?

David Rothenstein^ Hey Ryan, it's David. I think it's a little of both, right? Are we in a position to quantify in terms of volume or timing of that? The answer is no. Again, my hypothesis is, with 80% of the GPUs being sold for inferencing, which is done largely in a centralized way, given the need -- I'm not even going to say the need for low latency. It's going to have to distribute those inference GPUs closer to the end users, right?

When and how quickly, that's a big open question. So these are part of discussions that we are having, but is it resulting right now in the near term with hard near term demand as a result? The answer is no. I think that will be a function of the ongoing data center infrastructure buildouts that we're seeing today. And I do expect that to be instantiating over time, but I can't quantify it for you right now.

Ryan Koontz^ That's fair. In your discussions, do you hear any different requirements than you would normally see from cloud providers in terms of optical new products needed or anything?

David Rothenstein^ I don't. Ask Brodie and Scott that question. I certainly haven't heard of any. Again, it's the idea of a more distributed data center infrastructure, whether it's because of the monetization to the edge or data sovereignty or power consumption, would drive a need for incremental optical DCI connectivity. I don't believe there are very specific new requirements for optical as a result of that.

Ryan Koontz^ Okay, that's fair.

David Rothenstein^ Yes. The customer conversations -- sorry, go ahead.

Ryan Koontz^ No, go ahead, please.

David Rothenstein^ No. The customer conversations that we're having right now is the vast majority of the traffic that's going to get pushed onto the network from an AI perspective, has more to do with training and getting the large data models into those core data centers. So that's what's driving the bandwidth across the backbone, submarine, portions of their network, et cetera.

Because of how important fiber is to them, and that, going back to my chart on capacity versus cost, fiber assets are very important to them. So they have been pushing for more C&L van deployments to trying to double the capacity of their fiber by going to the L-band. But outside of that, there's no real changes at this point.

Ryan Koontz^ Okay, that's great. And we haven't touched on switching and routing much today. I know that's an important segment for you for growth, and you've got certainly a place at the table there, but the revenue growth has been somewhat stalled out. I wonder what you see your insertion point there looking like.

I assume this is dominantly in the communication, the CSP side of the world, not cloud, that you see an opportunity for switching and routing and where you see, you can differentiate and break in there.

David Rothenstein^ Yes, Ryan, thanks. You're right. I mean, this is -- the particular focus of this presentation, this topic was really leaning hardware into optical systems and interconnect. But you're absolutely right. Routing and switching is and remains a really important strategic initiative and TAM expansion opportunity for us. You're right. It is today in the near term, largely a service provider segment.

So a combination of us being the challenger and softness driven by various recovery dynamics globally that we and others have talked about, I think is why you're seeing a slightly down year on overall revenue. But we absolutely see really logical intercept

points for the various elements of routing and switching. So broadband access or enterprise and residential PON applications.

Clearly, the government subsidy funding kicking in from BEAD in '25 will be a really significant catalyst, we believe. Funding is taking longer than expected, as you might have seen, but we believe that we are incredibly well positioned with our solution in the space and look forward to getting our fair share of that.

In terms of metro or converged routing, again, we do see the need for high speed optical fiber connectivity driving a convergence of the optical transport and IP layers, where again, leading with our optical technology and our optical technology leverage, we believe is and is going to be a really significant catalyst for us in terms of driving our share in that space as well.

Ryan Koontz^ Got it. And one last, if I could, in the software domain, you talk about your capabilities there, and I wondered if that is that a different set of competitors? Do your peers stand up kind of head to head with you there in terms of their software portfolios, or do you run up against kind of smaller private players to compete in that space or maybe homegrown inside the telco?

David Rothenstein^ Ryan, it's a little bit of both. Kind of, you think about kind of network transformation driven by software, right? Again, driven primarily by service providers.

They're the ones who have to rethink and modernize their complex businesses and operational structures, and AI is only increasing the importance of that. If you think about it, with Blue Planet, with a federated inventory and assurance solutions and I and navigator that Scott talked about at the outset, the multilayer SDN controller, that single pane of glass that allows an operator to manage all the platforms across them.

It's a combination of some of the legacy players, the incumbent OSS vendors like Nokia and Ericsson, also in Amdocs, who typically are the most popular among service providers.

But there's also some kind of new cloud-based entrants because really, especially with Blue Planet but also navigator, really looking at cloud native, cloud-centric tools that automate service delivery, you can see some cloud-based entrants that are there as well, like a service now or even a salesforce.com. So it is a very fluid and evolving competitive landscape in the space.

But our goal is to have a competitive advantage by building a profitable automation business and deliver flexible, targeted solutions for the tier one telcos.

Ryan Koontz^ That's helpful. Thanks so much.

David Rothenstein^ Of course.

Gregg Lampf^ Thanks Ryan.

Operator^ Our next question today will come from Michael Genovese of Rosenblatt. Please go ahead.

Michael Genovese^ Great, thanks. I guess I have another timing question on this idea of more distributed data centers because the cluster sizes are getting bigger and I guess power restrictions and things like that, so that we're going to get more data centers and instead of having mega huge, massive ones, there'll be still be huge, but there'll be more of them.

And this idea that drive and accelerate the DCI market, can you talk about basically what you're thinking there and the timing of that? Is that something that we would see in '25? Is that more of a '26 event? What do you guys think?

David Rothenstein^ Brodie, want to give it a shot?

Brodie Gage^ Yes. So the data center buildups have been happening. Certainly power and the deployment to these GPUs and their AI clusters is forcing them to build out more data centers across the world.

So, I mean, they are going to have to build out more data centers, and when they build out those data centers, they're going to have to connect them through optical systems and through metro DCI type implementations. So it's kind of factored into our view going forward.

David Rothenstein^ And it does, Mike, and you know this, right? It does take a while to build out a new data center, right? Again, depending upon where you're at, given permitting and other items, it can take anywhere from two to three years even to build out a data center. So even if those dynamics we're seeing play out right now, what you're asking is, you know, when is the downstream benefit flowing to a provider like Ciena? I think it's hard to say. It depends really upon the particular operator, the timing of their own data center infrastructure buildups.

Michael Genovese^ Okay, thanks for that. And then my other question would be just, can you comment on your -- how you're thinking about your own supplier strategy or any changes that you have made recently and how those play into this discussion?

David Rothenstein^ Could you repeat that? I'm sorry, you cut out on my line.

Michael Genovese^ Yes. It's about your supplier strategy, the component vendors you use, the EMF providers you use. How are you thinking about aligning that with the opportunity and the kind of changes that you're making in your supplier strategy?



David Rothenstein^ So no real -- from a systems perspective, no real change. We work with all major suppliers across the globe, both from a semiconductor perspective and an optical perspective, and we work with all the major contract manufacturers as well. I think where things will become different is when we go inside the data center both for the campus and the interact type applications.

So we are looking, there is a large ecosystem for IMDD gray optic deployments there with a set of ODMs that play very well into that space. As we move in there, that's where we're starting to see some of the changes and some of the different partners that will have to turn up. Obviously InnoLight, the Optilink, CloudLite.

There's a bunch of them that have Fabrinet that have relationships into the cloud providers that we need to make sure we have as well.

Michael Genovese^ Thanks, again.

David Rothenstein^ Yes.

Operator^ Our next question today will come from Cobb Sadler of Catamount. Please go ahead.

Cobb Sadler^ Hey guys, thanks a lot for taking the question. Just kind of 10,000 feet, how did -- so you're talking about disaggregating your systems to a pretty granular level. You're going to sell DSPs, you're going to sell drivers. It sounds like you may license your technology. How did that come about? And I guess is the answer -- you do have at the three nanometer process, it sounds like your competitors probably aren't.

And so you do have some sort -- you have a large lead, 18 month plus lead for systems. So do you think you might have the same sort of lead when you start selling DSPs or drivers versus the existing competition? Thanks a lot.

David Rothenstein^ Let me give it a shot. Go ahead, Brodie.

Brodie Gage^ Yes, go ahead. Go ahead, David.

David Rothenstein^ So, hey, Cobb, let me take it from a high level and let Brodie do the more specifics. So at a high level, first of all, this type of disaggregation, Cobb, reminder, is not new to us, right? We saw it initially with cloud providers and open line systems and wave server and alien wavelengths. So desegregation is not just happening for the first time now. We saw it in the WAN, and we were obviously incredibly successful in leveraging that opportunity.

Here with cloud providers, they're looking at, as Scott alluded to earlier, really fundamentally different consumption models and different commercial models where optical line systems could be consumed independent of optical interfaces and coherent plugs could be consumed as merchant models, or as we talked about, even disaggregated

the optical technologies themselves in support of the AI fabric connection. So this, we believe, is a logical extension of our foundational optical technologies being really well suited for the cloud provider's desire to disaggregate the consumption inside.

Cobb Sadler^ Got it. But for example, if you're going to sell DSPs, I mean, do you have the sales infrastructure? Could you use your existing module sales infrastructure, or I'm sure you'd be glad to scale as sales ramp, but you're prepared for this model. And I'm assuming what's the most recent catalyst and kind of catalyst for the conference call is the cloud providers are kind of pushing you to speed things up. And that's the question.

Brodie Gage^ Yes. So for the optical systems that we talked about earlier, so regional long haul into the cloud providers, they're going to buy systems from us and our competitors. So there's very little disaggregation outside of it. Maybe they'll buy some line systems separate from transponders, et cetera, but there's very limited disaggregation.

As you go into metro DCI and selling pluggable optics, in some cases, our customers are looking for either a dual source for the full pluggable where they're asking us to arm an ODM or another provider with our DSP.

And as we go inside the data center, those cloud providers, yes, they're willing to buy 1.6T pluggable optics from us, but at minimum they're going to want a second source to that full pluggable. So they're asking us, are we willing to disaggregate further, either at the DSP level or even further down with some of our certis, ADC, and DAC capabilities? So that's where it's coming from.

Are we set up to support it? Yes, we are. We have set up an interconnects team that can support at the system coherent module as well as at the component levels.

Cobb Sadler^ Okay, great. Thank you very much.

Gregg Lampf^ Thanks. You're welcome. Allison [ph], we'll take one more question, please.

Operator^ Our next question will come from Dave Kang of B. Riley FBR. Please go ahead.

Dave Kang^ Thank you. Just one question regarding coherent pluggables versus LPO. Clearly, you're trying to push your coherent pluggables into the data centers, whereas other companies trying to push their LPO. What's your take on LPO? What would be the pushback on LPO?

David Rothenstein^ There's a lot of different techniques to connect within the data center. There's co-packaged optics, there's linear drive optics, there's near drive optics, and then there's transceivers.

And all of them are trying to solve the same thing, come up with the lowest cost and lowest power way to connect the racks together, or connect their campus data centers together. That, to me, is orthogonal to whether or not IMDD or coherent is going to have a play long term. So even if you go linear drive optics or you go co-packaged optics, IMDD is going to run out of gas and there's going to be a need for coherent.

So, as we set up ourselves in this disaggregated fashion, we're open to sell the full transceivers, and there absolutely will be deployments for the full transceivers. But we may also sell chiplets and intellectual property to enable some of those other consumption models, LPO, near drive optics, CPO, et cetera. Make sense?

Dave Kang^ And to follow up to that. Yes. Earlier you talked about power. I guess I'm hearing that, like, once black walls, you go from 8 GPU s to like 72 GPUs. I mean, power is going to increase exponentially. So even current solutions, I mean, that's going to be a big problem. So you said pluggable power will go down to the current transceivers, but then, sounds like that may not be good enough that it needs to go beyond, lower than current transceivers. Can coherent technology get there? Is that possible?

David Rothenstein^ Yes. Yes. And the other thing to remember is Ciena is the only one that's implementing their coherent designs in 3 nanometer technology. We're the only one that have a 224 gig certis and a 400 gig certis in 3 nanometer technology. And having that leadership allows you to drive lower power type implementations.

And the other thing that's important to note is when we talk about coherent inside the data center, it's not a full -- it's a different implementation than what you'd put in the LAN portion of your network. You might not need the same amount of dispatch -- anyway there's a lot of technical things you can take out of the chip to reduce power. That's what we've done with our coherent light implementation.

Dave Kang^ And I may have missed it, but then when do you say the time frame would be like in the next two, three years or...

David Rothenstein^ Yes, we're going to be engaging with 1.6T coherent light mid next year. I expect -- I don't expect any type of volume deployments until 2026.

Dave Kang^ '26. Got it. Thank you.

Gregg Lampf^ Thank you, Dave.

Operator^ That will conclude our question and answer session. And I'd like to turn back to Gregg Lampf for closing.

Gregg Lampf^ Thanks, Allison [ph]. And thanks, everyone, for joining us today. Again, please listen to the voice track slides on our website. There's really a lot of good information there, a lot of details that we go into, as well as what we did today. As

always, if you have questions, please reach out to us and we'll address those as quickly as possible. Thanks again, and have a great day.

Operator^ The conference has now concluded. Thank you for attending today's presentation. And you may now disconnect.