

Ciena Investor Forum 2026
Transcript of the Day's Events

Gregg Lampf, VP of Investor Relations

Introduction

Welcome. I'm Gregg Lampf, Ciena's VP of Investor Relations. Thank you for joining this webinar, where our team will discuss Ciena's strategy, innovation leadership, and our financial goals during this pivotal period for our company and the industry. With decades of experience and innovation, Ciena has established itself as the clear leader in high-speed connectivity. Since last year, when we hosted a similar event, we've continued to advance our industry-leading optical technologies that underpin today's network infrastructure for AI workloads, and that set a strong foundation for future demands.

Backed by focus and leverageable R&D, deep expertise, and proven scale, we've extended our leadership in optical systems while increasingly applying our coherent optical capabilities to a growing portfolio of inside the data center connectivity solutions, where rising bandwidth demand is pushing existing technologies to their limits. In today's constrained supply environment, we're working proactively to secure multi-year customer demand commitments and long-term supply agreements, strengthening our ability to support growing demand.

These efforts are delivering results, driving operating leverage and accelerating earnings growth. We encourage you to listen to the full webinar, including important financial updates on how we plan to continue converting our strong market position into long-term value for our customers and our shareholders.

As per this slide, indicates and directs accordingly, please note that will be discussing a number of topics that will be forward-looking and measures that will be non-GAAP. Thank you.

David Rothenstein, Executive Vice President and Chief Strategy Officer

Industry Dynamics and the Expanding Market Opportunity

Hi. This is David Rothenstein, Chief Strategy Officer with Ciena. In connection with the other presentations that will be given at our Ottawa Investor Forum this year, and as a predicate for that content, I'd like to provide some detail and context around current and evolving industry dynamics, particularly those being driven by AI, our view about categorizing the key markets in which Ciena plays, and the architectural shifts going on within those markets, and how all of that translates into an expanding market

39 opportunity for Ciena, particularly reflected in our view, about the total addressable
40 market for our business.

41 At the highest level, AI is creating a new era in communications that is being built on top
42 of previous eras. We can refer to the era from 1995 to 2010 as having been
43 fundamentally enabled by the internet, which oversaw the growth of connectivity. The
44 internet's global connectivity platform gave rise to the subsequent cloud era, with the
45 introduction of the digital cloud economy, and that cloud era continues to grow.

46 And now, in much the same way as cloud applications were built on internet
47 connectivity, AI is building interactive intelligence on top of the cloud and is
48 springboarding off of large data center infrastructure investment, largely driven by
49 enterprise adoption of cloud service platforms. One of the key characteristics of these
50 subsequent eras is the definition and redefinition of scale, with the scale of all aspects
51 of AI dwarfing that of previous eras in communications. Another key characteristic is
52 that optical technologies have enabled each successive era. In short, AI has created an
53 inflection point for our industry.

54 Whereas the internet era was driven by the telecom service providers, the drivers of
55 industry change today are the cloud providers. Hyperscalers are funding this AI
56 investment wave, and they've recently been joined by new entrants called the
57 neoscalers, or neoclouds, who are building new GPU-as-a-service or AI-as-a-service
58 business models. Since ChatGPT was launched three and a half years ago, this
59 customer segment has been investing and continues to invest unprecedented sums
60 into the AI market. Analysts now expect their Capex spend this year on everything
61 relating to AI, including data centers, compute, power, and network infrastructure to
62 reach \$1.2 trillion, up over \$100 billion from the beginning of this year.

63 That number is expected to grow to at least \$1.6 trillion next year, and probably higher,
64 and grow into some significantly larger number by 2030. Indeed, it's worth pointing out
65 that this market growth has been very difficult to forecast. As we illustrate on the right-
66 hand chart. Actual spend has consistently been higher than the previous year's forecast
67 of various industry analysts.

68 While it's not shown on this slide, there's an important point to note. There's a lot of
69 industry and media focus, understandably, on the cloud providers, but service providers
70 are playing and will continue to play an important role around the world. They've largely
71 underinvested their networks in recent years. Their access to fiber means that they will
72 benefit from an expanding set of managed wholesale opportunities.

73 Data center build outs will increase to support data sovereignty rules and regulations,
74 and traffic will ultimately need to leave the data center and traverse the service provider
75 network over the WAN in order for AI investments to be monetized.

The large-scale investment in AI infrastructure is showing up in data and bandwidth forecasts, and this slide highlights some examples. The amount of global data being generated continues to increase. IDC's latest model highlights the continued acceleration of data growth, with hundreds of zettabytes expected to be generated by 2029. Omdia has forecast the amount of data service providers are expected to move across their networks, which is a proxy for how much AI is influencing the bandwidth that end users consume. The data that's related to AI is growing at over 35%, indicating that AI will be a significant driver of new growth across the metro. In terms of bandwidth, we're running out of adjectives to describe the growth rates.

Inside the data center, the traffic demand and growth rate is not only enormous, but also industry analysts have been unable to forecast it successfully. Signal has needed to double their forecast in each of the past two years, and network bandwidth is growing at a similar rate, albeit at a much smaller capacity. One of the reasons for network forecast uncertainty is the introduction of scale-across. This special class of data center traffic interconnects, back-end fabric of training compute clusters located in geographically distributed data centers, with extraordinarily high amounts of bandwidth required. Our internal projections, which are very preliminary, estimate that scale-across bandwidth will account for one third of total global network demand by 2029.

But AI training isn't the only driver of new growth. The final chart illustrates how the energy consumed by inference workloads is forecasted by McKinsey to grow faster than training over the next few years. This indicates a significant shift from building out the foundational AI training engine towards using and monetizing the power of AI.

Building on these positive growth signals, we believe we are only in the early stages of a substantial multi-year ramp in optical connectivity demand. While this chart is based solely on our internal assessment, it provides a view of directionality for the expansion in optical communications bandwidth in the network and breakdown in workloads over the network by the end of the decade.

Starting with non-AI. The cloud is not going away. Cloud backlog figures continue to rise, and traffic associated with network cloud applications are growing at a steady pace. We expect that dynamic will continue going forward. By 2030, we estimate that the proportion of Capex associated with non-AI workloads will still account for approximately 30% of the total. On top of this, we've seen a surge in inter-data center traffic relating to scale-across and other high-capacity bandwidth flows associated with the distribution of AI training workloads.

Models are getting bigger and will need to be trained over greater distances, and ongoing power, space, and cooling limitations will continue this trend. We estimate AI training workloads to consume approximately 30% of Capex by 2030.

113 And then turning to AI inference. While the last three years have been primarily spent
114 building out infrastructure to train large language models, the coming years will be
115 spent building out the infrastructure, both to distribute these models across the
116 network and to monetize the investment.

117 In so doing, inferencing will become the dominant driver of global bandwidth. And we're
118 just at the very beginning of that inflection point. By 2030, we anticipate inference
119 workloads to assume the dominant share of Capex at approximately 40%, which brings
120 the total spend on AI workloads to 70%. Inference has already started in the hyperscale
121 domain, but it will quickly spread out to geographically distributed regions and metros.

122 It's important to note that we're projecting all three drivers of connectivity demand to
123 continue growing, albeit at different rates, well into the future. As an important example
124 of this growth, the scale of AI data center infrastructure is evolving quickly. The growth
125 of massive GPU clusters over multiple data centers is redefining the units we use to
126 describe the scale cluster size, energy consumption, and bandwidth capacity.

127 Over a few short years, the number of GPUs in an AI cluster has grown from being
128 measured in thousands to millions, and corresponding cluster energy consumption has
129 grown from megawatts to gigawatts. As a result, AI is redefining the fundamental units
130 of scale in both the data center and the network. The expansion of cluster size is
131 redefining the AI unit of compute, evolving from the single processor, to a rack of GPUs,
132 to a data center comprising thousands of racks. This effectively multiplies the unit of
133 compute by 100,000 times. It's also having a direct impact on the cross-sectional
134 bandwidth carried between data centers and in the network.

135 Historically, DCI connections would be required to carry hundreds of gigabits to a few
136 terabytes per second. This could be achieved using one or a few wavelengths within a
137 single optical fiber. Today, fabric expansion between adjacent data centers is
138 demanding petabits per second of cross-sectional capacity, and this is significantly
139 more than a single fiber can handle.

140 Moving from a single wavelength that can carry one terabit to a single fiber carrying 100
141 wavelengths, to a conduit that can carry a thousand fibers. In redefining the new unit of
142 capacity, this results in the ability to ultimately support hundreds of petabits per second
143 per conduit. And in the case of both compute and bandwidth capacity, all of these units
144 will exist and coexist for a long time to come.

145 Putting it all together, it's the sum total of these bandwidth and other AI-driven
146 dynamics that give us such strong confidence in the durability of industry demand for
147 optical connectivity, and the criticality of network infrastructure. It's also serving to
148 underwrite our views about the expanding market opportunity and specifically market
149 growth. As we've said recently, we project the total addressable market, or TAM, for our
150 business to effectively double over the next three years. More specifically, we see a

151 three-year growth rate of 24%, where the market grows from approximately \$27 billion
152 today to approximately \$52 billion in 2029.

153 I now want to turn to breaking down this TAM a bit further. First, by going through some
154 of the key architectural shifts within our key markets, we'll then show you our view of the
155 TAM by those markets, and then by looking through the lens of our product portfolio,
156 after which we'll show a view of the TAM broken down by product groups.

157 With respect to all of these views on total addressable market, it's important to keep in
158 mind that our growth projections are not just based on a rising tide of market growth,
159 but rather the assumption that we will continue to take market share in this
160 environment. And the other presentations at our Investor Forum will explain exactly why.

161 With the introduction of AI over the past three years, we witnessed significant growth in
162 high-bandwidth connectivity by the cloud providers, introducing new markets for high-
163 performance optical connectivity. Within the traditional WAN markets, our business has
164 provided bandwidth capacity that's enabled the circulatory system of the internet, first
165 for long-distance networking and then expanding to metro backbone and regional
166 markets, and then access and aggregation at the network edge.

167 Over the past decade, we've seen the build-out of data centers in support of globally
168 distributed cloud applications. This drove new investments into front-end data center
169 fabrics and the DCI services that connected those data centers together across the
170 WAN.

171 With the introduction of AI, a new back-end fabric has evolved to support AI's training
172 function. Massive bandwidth growth in the fabric has stimulated rapid growth of
173 markets for 400G, 800G, 1.6T and soon 3.2T optical interconnects. And it's also created
174 new high-bandwidth connectivity between different data centers in the form of
175 distributed training, or scale-across.

176 And now today, cloud's front-end fabric is expanding to accommodate increased cloud
177 and AI adoption via inference that's also accelerating DCI connectivity to support the
178 interconnection of many new data centers. We categorize these trends in the context of
179 the evolution of optical connectivity inside the data center, and in the context of a new
180 market category we're introducing today that we're calling AI-WAN. In this market, we
181 include the different DCI applications associated with cloud networking, AI training, and
182 AI inference.

183 Turning to some of the key dynamics and architectural shifts within each of these
184 markets. As we know, the traditional WAN markets have been historically over-indexed
185 to the service providers. Now, with AI in the AI investment era, we can see how these
186 markets and how we serve them are being significantly impacted. In the network
187 backbone, as enterprises increase their use of AI, and cloud providers look to build out

their DCI connectivity reach around the world, bandwidth and fiber-scale services are on the rise in the form of managed wholesale services such as MOFN, dark fiber, and other infrastructure services, creating new growth markets for the service providers. At the network edge, internet-class networks are not expected to meet the bandwidth and quality-of-service demands of AI, such as latency or determinism, which require a rearchitecture of the edge, opening a new market for inferencing across routing and switching interconnect infrastructure. And in terms of network operations, extreme scale and complexity is impacting the service providers' ability to transform the network operations capabilities, which is creating new markets for the application of AI intelligence and AI ops to traditional OSS, management, and control domains.

With respect to the market that we're calling the AI-WAN, here the AI-driven shifts in architecture are defining a new set of WAN-related applications. First, network DCI or Data Center Interconnect, continues to expand in support of both cloud and AI.

Cloud providers' need for operational simplicity at high volume with lower price and power, has challenged the traditional modem consumption model in metro and regional areas, introducing a new market for ZR and ZR+ class coherent pluggable transceivers. And a new DCI mesh is evolving to interconnect new colocation data centers that are being built out in support of private enterprise hosting of AI models or neocloud data centers offering GPU as a service. And with data center scale-across, as we've noted, a limited supply of energy to data centers is forcing AI training workloads to expand across multiple data center locations, creating a new market for lossless, ultra-high capacity coherent optical connectivity over many fibers, forcing a rethink of traditional amplified fiber line systems. And the data center markets, which are of course, heavily over-indexed to the cloud providers, are also being significantly impacted by AI.

Starting with the scale-out fabric. Today's existing approach that combines packet switches and IMDD plugs is too large and power hungry to meet next-generation specs. As a result, the market is now exploring new ultra-high capacity co-packaged and near package optics solutions to overcome this challenge inside the data center. And between buildings across the data center campus, IMDD modulation is hitting bandwidth and distance limitations, making this the likely first deployment scenario for Coherent-Lite.

With respect to data center scale-up, future generations of XPU's require higher-bandwidth interconnectivity, but the distance that existing copper interconnects can transmit is reducing as data rates increase. This is opening markets for new copper cable solutions, but will ultimately lead to the further opticalization of the data center via ultra-dense optical input/outputs based on silicon photonic interposers and co-packaged optics.

225 And finally, with respect to data center operations. The bandwidth and power required
226 to manage millions of devices is straining traditional out-of-band management
227 communication solutions, which are based on older technologies and has created a
228 new market for DCOM, a PON-based out-of-band network.

229 Stepping back and looking at all of these architectural shifts, they are not only
230 impacting the industry but are also directly operating in favor of Ciena. These shifts are
231 coming to us, and our strategy is purposely designed and aligned to take advantage of
232 them. They are creating increased opportunities for our business, and Brodie and Dino
233 will talk about how we're addressing them from a portfolio standpoint and why we win.
234 And they are at the core of what's driving the expanding addressable market.

235 Now that we have further context of key dynamics within these three primary markets,
236 we can return to our TAM chart and break it down by those markets within the \$52 billion
237 overall market in 2029. First, in the traditional WAN markets, we see very strong growth
238 from \$21 billion to \$26 billion, representing an 8% growth rate.

239 This is important for two reasons. First, that growth rate is much larger than the
240 historical low-to-mid-single digits growth rate. And second, that growth is off of a very
241 large base. Second, in the AI-WAN market we see very strong growth, from a market
242 going from \$7 billion to \$14 billion over the next three years, representing a 30% growth
243 rate. And finally, on the inside the data center market, we see extreme growth, growing
244 from a base of almost nothing today for Ciena to \$12 billion in 2029.

245 Turning now to Ciena and its products, we can see specifically how our strategic
246 portfolio investments map across these markets. Across the top is an illustration of
247 where our primary customer segments play in relation to these three markets and their
248 sub-domains. With the service providers spending and operating in the WAN markets,
249 as well as in the DCI portion of the AI-WAN, for for all of the reasons that I mentioned at
250 the outset of the presentation. And with the cloud providers spending and operating all
251 the way from the network backbone and the WAN through the AI-WAN, into the data
252 center and data center operations.

253 Perhaps more helpfully, our portfolio can be categorized into product groupings. Across
254 the WAN and the AI-WAN at the top of the page, you can see the Optical Systems
255 portion of the portfolio. This consists of our WDM optical networking products,
256 including modems and photonic line systems, including the new RLS Hyper-Rail
257 product supported by Navigator, our multi-layer domain controller for network
258 operations.

259 Below is a grouping of our Interconnects portfolio, consisting of plugs, modules, and
260 components. Starting with coherent modules and the ZR class of coherent pluggables
261 for both WAN and AI-WAN applications, it also includes Coherent-Lite for Metro
262 campus data center applications, co-packaged and near-package optics solutions, and

263 disaggregated components for scale-up and scale-out applications, and DCOM for data
264 center operations management.

265 Separately and more easily illustrated on a stand-alone basis are the WAN routing and
266 switching products, including the coherent routing and access and aggregation
267 products, as well as our Blue Planet software OSS solutions. All of which is supported
268 and underpinned by our global Services portfolio of implement, maintain, and
269 advanced services.

270 An important takeaway is how, consistent with the evolution of network architectures
271 and applications, several of our product solutions span across multiple domains and
272 even the markets themselves.

273 Now, against the backdrop of the product portfolio and those product groupings, we'll
274 turn one more time to the TAM growth chart. Again, within the \$52 billion overall market
275 in 2029, there are a few key takeaways.

276 First, we see a 2% growth rate combined for the Routing and Switching portion of the
277 network in which Ciena plays, as well as Network Operations. These areas are heavily
278 directed to the service providers and the growth rates indicative of the role being played
279 by those solutions in networks and network applications going forward. Second, there is
280 strong 15% growth in the Services portfolio, which is highly correlated to the projected
281 growth of the installed base of our hardware products. That's both Optical Systems and
282 Routing and Switching.

283 Third, there is 18% growth rate in the Optical Systems portion of the portfolio, which
284 again is significantly greater than the historical growth rate. Fourth, and finally, we see
285 almost 100% growth in Interconnects, again off of a small base, but still very significant
286 given the role being played by plugs, modules, and components in both the data center
287 and the AI-WAN. This clearly represents the fastest growing opportunity in the industry.

288 Taken together, Ciena has a product portfolio that is designed to maximize the growth
289 opportunities in the context of the industry dynamics and architectural shifts that we've
290 discussed. And you'll hear more about the product portfolio in other presentations at
291 the Investor Forum.

292 In summary, we have a tremendous opportunity ahead of us, with an almost
293 generational set of AI-driven industry dynamics and architectural shifts. In that context,
294 our vision and our strategy haven't changed and are incredibly well aligned with those
295 market opportunities. So, our priority is on scaling all aspects of Ciena in order to turn
296 that strategy into execution and take full advantage of the opportunities ahead of us.

297 Turning to separate presentations from the Investor Forum. Brodie Gage, our EVP of
298 Products and Technology, and Dino DiPerna, our EVP of Research and Development,
299 will focus on how we're scaling our product portfolio, and specifically our Optical

300 Systems and Interconnects products and technologies to address the expanded market
301 opportunity that I've discussed. They'll go through how we're addressing and will be
302 addressing the industry dynamics and architectural shifts, and specifically why we will
303 win in each of those cases.

304 Grant Hoffman, our new EVP and Chief Supply Chain Officer, will discuss how we're
305 framing supply chain as a strategic growth enabler and specifically how we're scaling
306 our supply across a number of dimensions to service the accelerated customer
307 demand we're seeing today and what we're projecting over the next several years.

308 And Mark Graff, our EVP and Chief Financial Officer, will discuss how we're leveraging
309 our industry leadership to scale the value of and from our business with the goals of
310 delivering strong revenue and profit growth, strong cash generation, and meaningful
311 long-term returns for our shareholders. These presentations contain a great deal of very
312 relevant content, and I strongly encourage everyone to listen to them all.

313 Thank you very much for listening, and we look forward to continue engaging with the
314 financial community on Ciena's journey going forward.

315

316 **Brodie Gage, Executive Vice President and Chief Product and Technology Officer**

317 **Technology Leadership Across Ciena's Core Markets**

318 Hello. My name is Brodie Gage, Senior Vice President and Chief Product and Technology
319 Officer, together with Dino DiPerna, who's our Senior Vice President and Chief Research
320 and Development Officer, we're going to talk about the technology leadership Ciena has
321 across our core markets.

322 Ciena's addressable market covers a vast array of network applications. On the left-
323 hand side of the chart, you can see our traditional network markets, typically service
324 provider WAN networks from network edge through network backbone and network
325 operations. In the middle, what we call the AI-WAN, which is inclusive of network, DCI
326 and scale-across fabrics. And finally, on the right-hand side of the chart Inside the Data
327 Center, from scale-up to scale-out, as well as our DCOM solution for data center
328 operations.

329 Not only does Ciena play in each one of those network applications, they are the only
330 supplier that breeds leading technology across the full Interconnect and Optical
331 Systems continuum. We provide optical solutions inside the data center for scale-up
332 and scale-out, across the AI-WAN, and traditional network markets from metro, long
333 haul terrestrial, as well as submarine. We are the only ones that can deliver optical
334 solutions all the way from scale-up to submarine transport.

335 What we're going to focus on in this portion of the presentation is the traditional network
336 markets on the left and the AI-WAN applications inclusive of network DCI and fabric
337 scale-across. There's two major sections. One, we'll talk about the architectural shifts,
338 which is expanding Ciena's addressable market, and two, why does Ciena win across
339 technology and systems leadership, scale and deployment advantage, and our
340 innovation engine?

341 As I just mentioned, there are architectural shifts in the network that are expanding
342 Ciena's addressable market. And one of the lenses that you can put on our business is a
343 customer lens. And we really break it down into two major customer units, the cloud
344 providers and then the service providers. On the cloud provider side of the equation, the
345 unit of currency used to be wavelengths. And we used to talk about 100G, 400G, 800G,
346 etc. With the amount of demand and capacity growth that is going onto those cloud
347 provider networks, whether it's for scale-across, data ingestion, content distribution, or
348 inferencing, the unit of currency is shifting from wavelengths to fibers. The other thing I'll
349 state is that scale-across is still in the early innings.

350 We've spoken multiple times about synchronous and asynchronous scale-across
351 applications. What we're seeing is there are more and more deployments within each of
352 the hyperscalers, and there are more and more hyperscalers that are starting to do
353 training across distance. Inference is going to have an impact on both the cloud provider
354 network and the service provider network. Inference drives distributed connectivity and
355 expands connectivity beyond centralized AI infrastructure.

356 It will impact the cloud provider networks, of course, but the service providers are going
357 to have a larger and larger role to play as inference becomes more pervasive. On the
358 service provider side, these applications are driving new service models. We've talked
359 about MOFN, or managed optical fiber networks. There's applications for dedicated
360 fiber, and even on-demand connectivity, and new services.

361 The net of all this is that these architectural shifts are increasing the strategic value of
362 optical systems. With the capacity growth and the unit of currency shifting from
363 wavelengths to fibers, it drives a new generation of optical platforms. On the left-hand
364 side of the chart, you can see how AI capacity would impact amplifier huts along a long-
365 haul route using traditional technology.

366 And if you wanted to support 20PB across a network, you would need twenty-two huts
367 at each ILA site. That is not sustainable. And that's what's driving new photonic
368 architectures like multi-rail, where we can drastically increase the capacity and the
369 density of our amplifier solution.

370 It's not only having an impact on line systems, though. If you're going to light up full
371 systems or full fibers, it doesn't make a lot of sense to plug in your transponders one at a
372 time in a pluggable form factor. What makes more sense is to have very low-cost, low-

373 power connectivity from your routing layer into your transport layer and light up the full
374 spectrum one at a time, versus one channel or one wavelength at a time. That drives
375 much lower cost, much better density, and lower power and cost per bit. These new
376 architectures, and the amount of traffic that's going on the network, requires innovation
377 in how capacity is packaged, deployed, and managed.

378 It also goes back to our philosophy that the winning hand is to not have one type of
379 solution, just pluggable optics or just high-performance modems. The right solution
380 requires both.

381 Scale-across is one of the first use cases that starts to put this new architecture in
382 place, and in fact, is one of the applications that's moved the unit of currency from
383 wavelengths to fibers. We've talked about one hyperscaler deploying our solution
384 across about 80km of distance and carrying over 20PB per second of bandwidth. For
385 that span and that type of capacity, it requires over 400 fiber pairs just for that one
386 application.

387 I want to highlight that each hyperscaler project are going to be different – the number of
388 GPUs, the number of fibers, unprotected or protected. But for very illustrative purposes
389 – again, each hyperscaler will be different – but for illustrative purposes, deployments
390 say for a 500,000-GPU cluster, can reach to about \$500 million. That's not Ciena
391 revenue, but that's the TAM of the opportunity, or about \$1,000 per GPU in that example.
392 And that's without amplifiers.

393 Once distances grow and reach a 1,000km span length, for example, you now need
394 inline amplifiers. And in that case, the project size can double. If it was the exact same
395 design, it could be doubling to \$1 billion. And if you went to 2,000km of reach, that
396 project – if it was exactly the same, which it won't be, but if it was – can jump to a 4x
397 multiplier, or \$2 billion, because you also need to change the end-point transponders,
398 from pluggables to high-performance modems.

399 Again, this is a generalized example, and the specifics – the number of fiber pairs, etc. –
400 will be different depending on that project.

401 Because of constraints on power and the need to locate data centers across regions to
402 access that power, we expect every hyperscaler will want to do scale-across – both
403 synchronous and asynchronous – over extended distances.

404 And we co-developed both our next-generation RLS line system – our Hyper-Rail
405 product - with that customer for terminal configurations with redundancy and
406 protection switching built in, and we developed a set of 800G coherent optics that are
407 deployed in that use case.

408 Going forward, we expect that same customer to deploy in more regions – it will take
409 time – and we see other customers coming on board to deploy both synchronous and

410 asynchronous scale-across applications. When they do that, they require more dense
411 amplifiers. That's where Hyper-Rail comes in. And in some cases, as they increase the
412 distance, our high-performance platforms and modems are required along with the
413 plugs.

414 As I mentioned earlier, we see scale-across as entering the early innings of a multi-year
415 investment cycle. There will be more deployments, and there will be more deployments
416 over distance. Our customers cannot get GPUs in one data center, and as they
417 distribute those GPUs, the network becomes more and more critical.

418 This AI infrastructure build out will create opportunities for the service providers as well.
419 If you start on the right-hand side of this chart, with distributed training, the service
420 providers have a role to play as well, whether it provides conduits or fiber or managed
421 optical fiber networking, there are certain regions where the cloud providers do not have
422 the regulatory rights to build out a network. They have to rely on the service providers,
423 and the service providers will provide that network infrastructure on behalf of the cloud
424 provider.

425 As we move more and more to inferencing, the service providers have the real estate
426 assets and the fiber connectivity into major data centers, central offices, and even the
427 end consumer and enterprise customers. They're going to have a role to play in
428 inferencing fabrics and providing the connectivity for inferencing traffic.

429 If you go to the middle of the chart and AI cloud distribution, the reality is that for data
430 sovereignty, there's going to be more and more distributed AI cloud infrastructure.
431 Certain regions are going to want to have their data stay within their continent or within
432 their country. That's going to drive more and more cloud or AI cloud distribution in the
433 network. And the service providers will have a role to provide that connectivity.

434 As we move more and more towards the edge of the AI infrastructure, there's a need to
435 enable enterprises to connect to the AI cloud or interconnect the private enterprise AI
436 infrastructure as well. And finally, there's all this discussion around building, operating,
437 and connecting AI infrastructure at the edge of the network. These service providers
438 have the data centers, the offices, the fiber infrastructure to provide that edge AI
439 infrastructure. And you see more and more partnerships. I believe you'll see more and
440 more partnerships between the XPU providers, as well as the hyperscalers or cloud
441 providers, together with the service providers.

442 AI is expanding where and how service providers can participate in the connectivity
443 value chain. And again, we believe there's a significant opportunity for the service
444 providers going forward.

445 Now that we've covered the key trends in the industry and what's driving the capacity,
446 the next question is why does Ciena win? And it really comes down to three major items.

447 Number one is technology and systems leadership. We've driven best-in-class
448 performance over multiple generations, and we have systems-level understanding of
449 how technology is used and the challenges it must solve. We've been in this market for
450 decades, and we leverage that experience to build best-in-class products.

451 Number two is scale and deployment. We output from a supply chain perspective, more
452 photonics than anybody else in this industry. And we have the breadth across coherent
453 modems, photonics as well as routing and switching. We also have end-to-end service
454 and support capabilities that we'll talk to in more detail, both for the service providers
455 and the cloud providers.

456 And finally, we have a best-in-class innovation engine. We get faster to market than
457 anybody else, and we drive better cost and power per bit than anybody else in this
458 industry. The other thing that we do is we build reusable technology building blocks.
459 That allows us to leverage these technologies across multiple use cases and multiple
460 applications. We're going to cover these in more detail.

461 On the technology and system side, I'm going to start with our coherent modems, which
462 drives traffic down the fiber. We've been the first for every generation of coherent
463 modem, whether that was 100G with Wave Logic 3 or now 1.6T with Wave Logic 6
464 Extreme.

465 We've led in coming out with the next generation of technology every single time. And
466 that, of course, drives the lowest cost and power per bit. The other thing, though, that
467 we've been doing over the last decade is we've been co-developing our products with
468 key customers, especially with the hyperscalers. And we built those platforms to suit
469 their needs, whether it's a data model perspective, APIs, streaming telemetry,
470 operational aspects, etc. We really co-developed.

471 And what that's allowed us to do is to sell our coherent modems, yes, and our
472 Waveserver system products. Yes, and our pluggable form factors, but also as coherent
473 modems that can be integrated into their platforms as well. So, we repeatedly
474 translated our core technology leadership into new platforms, new form factors, and
475 deployment models as these cloud architectures have evolved.

476 If you take our latest generation WaveLogic 6e 200Gbaud 1.6T-capable modem, and you
477 compare it to all of our system competitors, 140Gbaud, 1.2T-capable modem on a
478 typical long-haul network across the United States, what type of advantages do our
479 customers get by deploying WaveLogic 6?

480 Number one, they get more bits per wavelength. So, within the same footprint, a
481 200Gbaud modem reaches 1.6T, where 140Gbaud typically falls back to 800G. So, you
482 get double the capacity per modem. For 800G services, two by 400G-type services, a
483 1.6T modem on that typical network supports 96% of the routes with no OEO

484 regeneration. A 1.2T modem requires on half of the routes a regeneration point, which
485 adds cost, complexity, and power.

486 And on a capacity-per-capacity basis, a 200Gbaud modem supports 25% fewer
487 modems to carry the same set of traffic across that network. It's a massive advantage
488 for Ciena and the customers that deploy this technology, and our two-year technology
489 lead delivers meaningful network advantage and has driven many competitive wins.

490 Shifting gears now to photonics, Ciena believes that the photonic layer is the foundation
491 for AI-scale connectivity. Some of our routing competitors talk about the Routed Optical
492 Network and the elimination of photonics completely from the network. That is the
493 wrong way to build an infrastructure.

494 And for AI scale, photonics are going to become more and more important. We've seen
495 that in the metro network, we've seen it in the long-haul and submarine network, and
496 we're now actually starting to see it in the campus scale-across and even inside the
497 data center for multiple use cases, whether that's front-end switching or for XPU
498 interconnect. Photonics are the lowest-power, lowest-latency, and lowest-cost way to
499 move bits across an infrastructure.

500 Investing in the photonic layer, like Ciena has done, also drives system knowledge. And
501 because you understand what's happening at the photonic layer, Ciena is able to
502 customize its coherent modems and other solutions with that wide expertise and that
503 breadth of knowledge across both the coherent modems and the photonics. The most
504 efficient network is one that keeps traffic optical for as long as possible.

505 Similar to our coherent modems, we're now on our sixth generation of photonic system
506 innovation. We've been working and co-developing with multiple hyperscaler
507 customers now for over eight years on our next-generation open line system, which we
508 call RLS. We started with a C-band optimized CDC solution. Then we went to C-and-L,
509 integrated C-and-L. And then we went to what we call colorless direct attach and
510 passive point-to-point-type use cases, subsea use cases at step three. Realizing that
511 our customers have multiple deployment models and multiple huts and solutions they
512 need to build into, we developed 300 millimeter depth configurations. We developed
513 high-density scale-across use cases and platforms. And now our high-density Hyper-
514 Rail solution is that sixth generation. We've built these systems and co-developed these
515 solutions again for over eight years.

516 We've built the APIs, NBI, declarative configuration, streaming telemetry, auto
517 calibration, ASC noise loading, Layer Zero control plane with spectrum sharing, optical
518 protection, direct-attach, security solutions together with these hyperscale customers.
519 Our competitors must overcome years of accumulated photonic systems expertise,
520 automation, and deployment know-how. Hyper-Rail is not a new start. It's the evolution
521 of what we've been doing now for over a decade.

522 So Ciena has led in photonic technology development, and we've co-developed with
523 multiple hyperscaler and service provider customers. What has that done from a share
524 perspective? Well, if you take the Dell'Oro Group Optical Transport report up to Q2
525 2026, you can see that Ciena has recognized 4.5x more photonics revenue than our
526 nearest systems competitor. We've driven 5x more photonic revenue than all others
527 combined, and we've driven 25 points of market share gain over the last four years, and
528 we now have 70% market share in open line systems.

529 A critical component of our strategy going forward to continue to lead in photonic line
530 systems and continue that generational advantage. Our next step is what we call RLS
531 Hyper-Rail. And as I mentioned earlier, as scale-across goes across larger distance and
532 inference starts to impact the network in a larger way, you can't rely on the existing
533 solution. The number of huts required for the bandwidth is just, is just too much.

534 So there's a massive densification of amplifiers. We're going to try to get a 32 x density
535 improvement in our first release coming out at the end of this year, but that's not where
536 it will end. We're targeting even further density improvements and more vertical
537 integration as part of our solution. We will leverage all the co-development that we've
538 already done with all the hyperscalers. We'll leverage all the same software that we've
539 developed over the last eight years. And because we have incumbency and integration
540 into operational systems, we have a massive advantage as we move to this next
541 generation of Hyper-Rail technology.

542 On the final chart of our photonic line systems, I just want to kind of differentiate what
543 Ciena has done from a strategy perspective versus our nearest system competitor in the
544 middle, and the merchant or component suppliers that are starting to get into the game
545 as well.

546 When we first started our photonic systems way back when, prior to 2010 even, we
547 relied on component suppliers to build our full line cards and our full solutions. We
548 developed the software, but the hardware was developed by merchant suppliers, the
549 gain blocks, etc. From 2018, when we started the co-development with the
550 hyperscalers, we decided we were going to take back control and have the design
551 authority on both the hardware and software.

552 So yes, we buy components from merchant suppliers, but we build our own line cards
553 and we build our own software on top of that. And what that allowed us to do was
554 customize form factors, customize density, customize where we put Raman, where we
555 put EDFAs, etc. And it allowed us to control the software architecture and the software
556 solution on top of it and differentiate on multiple levers.

557 We're going to continue to do that into Hyper-Rail and beyond, and leverage all the co-
558 development work that we've already done. Our nearest competitor on the system side
559 continues to have full line cards built by component suppliers, so their ability to

560 customize for use cases and for different customers is much less than what Ciena is
561 been able to achieve.

562 Now they're just getting going, we believe, with Hyper-Rail and trying to develop some of
563 their own solutions, but they got at least eight years to catch up and all the work that
564 we've done.

565 And then from a component supplier perspective, component suppliers really don't
566 have that system knowledge or that software around the APIs, the data models, the
567 operational integration into our hyperscale, our service provider customers. And I think
568 they're really lagging versus what we've been able to achieve over the last few years.

569 The second major reason why we win is because our scale and deployment advantage.
570 If you look at our services, of course, we offer all the traditional network services that
571 many of our system competitors offer as well. And we do that on a global basis, whether
572 it's to manage, maintain, support, build, etc. We've always offered those services. But I
573 think really where we've differentiated is we've really gone after some hyperscale-
574 oriented services as well.

575 I talked about co-development of technology. But we also co-develop and help them
576 build large-scale optical networks around the world. We support flexible operating
577 models from a system perspective, a plug perspective, or even from a component
578 perspective. We help them accelerate their deployments. We have over 80 engineers,
579 for example, that are helping one hyperscaler customer deploy their optical network at
580 scale.

581 We maintain and support their network around the world, and we even offer proactive
582 optimization, migration, and evolution services to our customers as well. And really,
583 what that's led to is a doubling of our hyperscalers services revenue from '24 to '26. Our
584 largest services customer is now a hyperscaler. And we really differentiate versus all of
585 our competitors. And having over 1,000 deployment optical domain experts, we've done
586 over 5,000 projects or 20,000 systems per year, and we have over 400 global sparring
587 depots around the world.

588 So along with our technology differentiation, our services differentiation complements
589 that, but also provides a large moat that our nearest competitors are going to have to
590 cross.

591 The third major reason why we win in this market is a best-in-class innovation engine.
592 We own and become vertically integrated where we can differentiate. We design for
593 supply flexibility. I spoke about photonics and how we've taken over the design authority
594 on all our line cards. That allows us to choose multiple suppliers and be able to output
595 more than anybody else.

596 We optimize across materials systems. We are fabless, but we leverage CMOS,
597 BiCMOS, gallium arsenide, silicon photonics, and indium phosphide materials. And, of
598 course, we invest, and we've increased our investment to continue to lead in this
599 market, the AI-WAN and the traditional WAN networks.

600

601 **Dino DiPerna, Executive Vice President and Chief Research & Development Officer**

602 **Technology Leadership Across Ciena's Core Markets (Cont'd)**

603

604 Here at Ciena, we are definitely proud of our science and advanced technology. But at
605 the end of the day, the key metrics that drive our technology programs are power per bit,
606 space per bit, and cost per bit. In fact, Ciena were the pioneers of coherent DSP-
607 assisted optical communications technology, and that dates back some 25 years ago.
608 In the late 1990s, while the industry was busy deploying 10G wavelengths, we were
609 pushing the technology towards 80G and even 160G.

610 The physics and tech associated with that work made it very clear that the traditional
611 on-off light modulation schemes of the past would simply not scale. This drove us to
612 adapt, of course, and advance the coherent techniques long used in mobile radio, but
613 not yet applied in the domain of optical communications. Looking back over the past 15
614 years and five generations of our wave logic coherent modems, we have consistently
615 been first-to-market in delivering that next step in capacity, and connectivity, and reach.

616 We've increased individual wavelength channel capacities by a factor of 16, from 100G
617 on a wave to an industry-leading 1.6TB on a wave. Today, these adaptive WaveLogic
618 modems also greatly simplified the deployment engineering required to turn up and
619 operate these global massive-scale networks. And all the while, we've reduced the
620 power and space by over 95% on a per-bit basis.

621 A coherent modem or transceiver is much more than a single chip. The large CMOS DSP,
622 or digital signal processing engine, is definitely one of the critical components, but the
623 coherent modem also depends on high-speed analog ICs. These include drivers and
624 transimpedance amplifiers, or TIAs, along with the high-performance electro-optic
625 devices, such as the transmit modulator, the received coherent mixer, photodetectors,
626 and of course, a tunable laser.

627 Even the DSP device is not simply a homogeneous digital CMOS piece. It includes a
628 digital core that performs the advanced signal processing used to encode and recover
629 the signal, as well as specialized mixed signal blocks such as a high-speed digital-to-
630 analog, and analog-to-digital converters, and the client-side SerDes. These high-speed
631 converters can ultimately limit the transceiver performance, regardless of the strength
632 and ingenuity of the digital signal processing core.

633 This applies to modems that are high-performance sub-assemblies, like the one on the
634 bottom right, as well as to the coherent pluggable modules – 400G, 800G, and emerging
635 1600G. Indeed, Ciena owns the critical intellectual property and design control over all
636 of those blocks within a coherent modem.

637 Our expertise across device types, material systems, and expertise in the optical
638 propagation allows us to make the right trade-offs across the functions, ultimately
639 improving both performance and cost. This is far more advantageous than simply
640 assembling a collection of separately sourced components. Each of those components
641 comes with its own margins and nuances, and they have to be treated as standalone
642 black boxes.

643 For example, we can use DSP functions in our math processor to compensate for
644 analog device implementation artifacts. We can make it a little easier on the analog
645 parts. We can also tune the design to specific characteristics of the different electro
646 optic materials that I'll speak about in a little bit, and we can also shape the individual
647 specs to deliver the best propagation performance per watt, which of course is the
648 ultimate objective.

649 Even companies that claim to build their own DSPs may still need to buy critical mixed
650 signal blocks from third parties. That not only adds cost, but more importantly, it can
651 often gate them in terms of time-to-market and the ultimate performance that they can
652 achieve from the modem.

653 Exploring the theme of vertical integration and various material systems a bit further, the
654 Ciena strategy is vertically integrated by design while remaining fabless. Some would
655 argue that owning the fab itself can offer an advantage, and this is certainly true in
656 specific cases, for example for cost. But that's only true if the fab runs at near-full
657 capacity, and more importantly, only if that particular material system remains the right
658 choice over time and over different applications. The reality, of course, is far more
659 nuanced.

660 Indeed, optimized solutions depend on multiple technologies and material systems,
661 each selected for the needs of that application and adapted as those requirements may
662 evolve. Returning to the key modem building blocks that are shown in the previous
663 charts, the electro-optic functions already rely on a various range of material systems.
664 Indium phosphide, of course, is essential for generating the light that we use in optical
665 communications, and it can also support functions such as modulation.

666 But material systems such as silicon photonics can offer better manufacturability and
667 cost in some of those applications. Looking to the future, there are newer, emerging
668 high-speed generations that will be deploying technologies such as thin film lithium
669 niobate, and even newer materials in crystals such as barium titanate for certain critical

670 functions. The high-speed analog devices are typically implemented in silicon
671 germanium bipolar, or BiCMOS technologies.

672 And lastly, the DSP itself is built on state-of-the-art CMOS. Ciena's WaveLogic 5
673 generation, for example, was the very first to use seven nanometer CMOS. WaveLogic 6
674 again was the very first to exploit the three nanometer CMOS technology. And, of
675 course, with an eye to the future, we have a new suite of two nanometer devices in
676 development for our next generation of WaveLogic.

677 The key point is that broad-based material system expertise, combined with the
678 flexibility to choose the right technology for each application, is what's required to
679 consistently deliver the best system-level solution.

680 It's worth spending a moment on R&D leverage and reuse. At Ciena, we view these
681 functions electro-optics, electronics, DSP algorithms together with all the associated
682 material systems, as a suite of reusable technologies. In fact, a toolbox.

683 We evolve and optimize those building blocks to meet the requirements of each target
684 network application. Examples range from data center interconnect, including scale-
685 across through the metro, long haul, and transoceanic networks. We then instantiate
686 those technologies across multiple form factors, including our Waveserver product line,
687 the 6500 platform, a suite of coherent modules, and a broad pluggables portfolio that
688 includes QSFP, OSFP, and the emerging XPO formats. The Ciena advantage comes from
689 combining these broad technology ownership, deep system expertise, field expertise,
690 and the efficient R&D leverage.

691 To summarize this section, Brodie and I covered both our traditional WAN and the
692 emerging AI-WAN market segment opportunities, along with the reasons why Ciena
693 continues to win and sustain its market leadership.

694 Firstly, Ciena possesses a distinctive combination of deep multi-technology expertise,
695 extensive system-level knowledge and most definitely a culture of innovation that's
696 always focused on redefining what comes next.

697 Second, we have a proven ability to turn this innovation into market-leading products
698 and real-world customer solutions.

699 Third, multi-generational field deployment experience that converts this technology and
700 products into operational advantages for our customers.

701 And lastly, the financial scale and strength to continue investing in the innovation and
702 technology development that's required for this fast-changing, demanding market. And
703 of course, the results will speak for themselves. Ciena's estimated market share of 25%
704 in service providers and 36% in the cloud market segment most certainly reinforces that
705 leadership. Thank you.

706

707 **Brodie Gage, Executive Vice President and Chief Product and Technology Officer**

708 **Expanding In and Around the Data Center**

709 Dino and I are now going to talk about how Ciena is going to expand in and around the
710 data center.

711 In this section, as I mentioned, we're going to focus on inside the data center markets.
712 We're going to have two major sections. Number one is why is AI changing connectivity
713 requirements in and around the data center. And section two, we're going to focus on
714 four use cases: data center out-of-band management, or DCOM, scale-up, scale-out,
715 scale-across, inclusive of Coherent-Lite technology. We'll talk to the trends, what our
716 solution is and why we differentiate.

717 So what are the key challenges that we see with growth and scale inside the data center
718 driven by this AI connectivity? Number one is just the data center infrastructure
719 management. That has traditionally been built with copper cabling, Ethernet leaf-and-
720 spine switches. It takes up rack space, it takes up power, and is really not set up to
721 provide management at AI-scale. And a new solution is required there.

722 On scale-across, we've talked about the trends, but our hyperscale customers can't get
723 GPU and the density and quantities they require in one data center. They require
724 multiple data centers to house those GPUs, and now they need lossless, high-capacity
725 connectivity over distance.

726 On scale-out, there are new density requirements. There's insufficient front panel
727 density for traditional transceivers. They're running into power and scaling issues. And
728 as rate gets increased, the ability to support with traditional technologies goes away.
729 And for shorter distance scale-up, passive copper is getting shorter and shorter, again,
730 as rates increase. We'll talk to all four of these use cases and what our solution is. We'll
731 start with data center out-of-band management, go to our Nitro Linear Redriver
732 solution, talk to our Vesta CPX portfolio. And then finally for scale the cross and campus
733 use cases, we'll talk to Coherent-Lite.

734 As data centers expand and the number of devices continues to grow within a data
735 center, the out-of-band management solution needs to change. As I mentioned
736 previously, the traditional solution consists of Ethernet switches and console servers
737 that sit within the rack. Those are connected with multiple copper cables in a redundant
738 fashion to an aggregation layer. And then that aggregation layer is connected up to a
739 spine or router layer, again through copper cabling.

740 The challenge with that is that that copper cabling is very expensive and very complex.
741 Hard to track where those cables are going and why. The Ethernet switches and console
742 servers take up rack space that could be utilized for TPUs or GPUs within the rack.

743 They take up a lot of power. They take up a lot of space, and it's an extremely complex
744 and static solution. This out-of-band management solution was ripe for new innovation.

745 Recognizing that a new solution was required, Ciena decided to try to leverage its PON
746 technology to provide a much lower-cost, lower-power, lower-space solution, leveraging
747 three major components.

748 One is an 8112 converged router. That's a routing platform that we've built. Two is a PON
749 OLT and ONU that you can plug in to that converged router or at the edge of the network.
750 And three, is a data center-optimized ONU that can sit now on top of the rack or
751 underneath the rack, but not take up the rack space within it.

752 It's passive optical networking. It requires zero rack units within the rack, and it's much
753 higher density than existing solutions. So, we took that idea and that technology – we'd
754 co-developed it with a leading hyperscaler, which was Meta – and we've been in
755 deployment now since Q4 of 2024. Think we've shipped over 300,000 ONUs on top of
756 racks now year to date. So we're scaling within that customer. But we also have multiple
757 other customers that were engaged with now as well.

758 We're obviously not standing still with our DCOM solution. We're going to continue to
759 lead and scale our existing solution, drive security improvements, speed of deployment,
760 scale of deployment. We're going to go after new hyperscaler and neoscaler demand
761 and opportunities with new ONUs and OLT hardware.

762 We've containerized our DCOM controller that allows us to deploy not only in our
763 solutions or our routers, as well as third-party routing platforms that may already be
764 inside the data center. We're leveraging our own software, but we're also working and
765 contributing to have SONiC support for our solution as well. And finally, this market, of
766 course, is applicable to multiple hyperscalers, but it's also applicable to the enterprise
767 and small data centers, as well.

768 Now, each opportunity won't be as large as a hyperscaler, but they'll add up over time.
769 So, we're building general data center market applications, inclusive of having software
770 off-box software, management, and support capabilities. So, we're reusing and
771 repackaging the technology to address a broadening opportunity within the hyperscale
772 and neoscalers, as well as with the general market.

773 Why is Ciena going to win in this DCOM market? Number one, we have great
774 relationships and deep customer engagements with Meta, and we're going to continue
775 to evolve this solution with that customer. We're engaged with multiple other
776 hyperscaler and neoscaler customers with this DCOM solution. We're focused on
777 DCOM, we're not focused on the residential broadband access market. We're really
778 focused on this DCOM market. We have a proven platform. We have outputted, as I
779 mentioned, hundreds of thousands of units already.

And we have flexible routes to market. I mentioned the containerized DCOM solution, that allows us to deploy again in our products, but also third-party routers and switches. What's not mentioned here and is a key differentiator again is our services. When you're deploying hundreds of thousands of units, the ability to commission, install, deploy, and support that type of scale globally is critical for our customers. So, along with the technology and product, our services are a key differentiator here as well.

Dino DiPerna, Executive Vice President and Chief Research & Development Officer
Expanding In and Around the Data Center (Cont'd)

As the interconnect rates continue to climb, the same physics that drove the coherent inflection point in the WAN is now forcing changes in data center interconnect architectures and technologies. This is showing up in two major technology shifts: the first from copper to optical, and the second from IMDD signaling optics to coherent DSP-assisted optics.

Inside the data center, interconnect rates are typically paced by the I/O or so-called lane speeds of the switch ASICs. The graphic on the left shows how these devices have evolved over time, with I/O lane speeds shown on the Y axis. While most deployments today still use 100G lanes, 200G lane devices have emerged in 2026.

Each switch ASIC lane rate generation maps to a typical interconnect rate that's shown on the Y axis of the graphic on the right. For example, eight by 100G lanes produce an 800G interconnect, with higher rates following the same progression as you move up the Y axis. The same chart on the right shows which technologies apply across different reaches, from meters to kilometers that are shown along the X axis. It also roughly maps the familiar application categories such as scale-up, scale-out, and scale-across against those reaches.

Although exact distances always vary, the technology transitions are clear as interconnect rates increase, copper is useful over progressively shorter and shorter reaches, eventually shrinking to only a few meters. And that space is being addressed by traditional IMDD or gray optics, and the emerging near package and co-packaged optical approaches. Similarly, as rates continue to climb, traditional IMDD optics also become reach-limited, creating the need to transition to coherent-enabled optical solutions.

Before turning to optical technologies, it's worth spending a few minutes on copper. Even the strong proponents of optics would agree that staying in the electrical domain as long as possible generally delivers the lowest cost and power. Not surprisingly, there are several copper-based solution options. The simplest is shown at the top, which is

817 direct attached copper or DAC. It is by far the simplest and lowest-power interconnect
818 approach. And as the previous chart shows, it quickly reaches its performance limits as
819 the bit rates go up.

820 To extend the copper reach, the industry adopted active electrical cables or AECs.
821 These cable assemblies include a full transmit and receive electrical retiming at each
822 end. While they improve the performance, they also add significant cost and power
823 consumption.

824 As these interconnect rates continue to rise, however, that performance boost becomes
825 less meaningful, and this leads to a third option which has emerged the active copper
826 cable or ACC. Based on a small, highly specialized Ciena linear device, ACCs provide a
827 meaningful performance improvement over a purely passive cable, while using only a
828 fraction of the cost and power of a full AEC. The table on the right compares ACC with
829 AEC at 200G per lane. Most notable ACC consumes barely one sixth the power of an
830 AEC, while giving up only a very small amount of reach and performance.

831 The Ciena solution for active copper cables is a custom device known as the Nitro 2004.
832 It is implemented in high-speed silicon germanium bipolar technology with four devices
833 integrated into each ACC cable. It delivers stronger link performance than competing
834 solutions while maintaining very low power consumption. It also preserves the
835 operational simplicity of the passive DAC approach. Plug and play operation. No
836 configuration. No special power feeds required. We are currently qualifying the solution
837 with multiple customers, including hyperscalers and ODMs.

838 As discussed, copper-based solutions are attractive for their simplicity, cost, and power
839 efficiency, but they eventually reach their performance limits. And at that point, the
840 industry has to shift to optical IMDD solutions. For several generations now, the
841 dominant model has been to connect the switch ASIC lanes directly to the front panel
842 pluggable optical modules themselves. But as those interconnect rates increase, the
843 signal propagation even within the box becomes more and more difficult, and that often
844 drives these costly, power-hungry electrical retimers similar to the ones that I just
845 described in the AEC on the previous page.

846 The answer, of course, is to move the optics much closer to the switch ASIC, reducing
847 the signal propagation channels and eliminating the need for expensive, power-hungry
848 retimers. This approach is known as co-packaged optics, or CPO. Some single-vendor
849 implementations have emerged, but that ecosystem or that closed ecosystem model
850 really undermines the supply chain advantages of today's system.

851 Specifically, how do you maintain separate supply chains for the switch device and the
852 optics, while preserving that customer choice across multiple suppliers? The second
853 equally important is decoupling the switch and optics innovation cycles. How do the

854 data center operators upgrade each independently, rather than being forced to change
855 everything at the same time?

856 The optimal approach preserves the advantages of both models, of course. Use
857 modular, pluggable optical modules that can be assembled directly beside the switch
858 ASIC, much like the front panel pluggable but connected to an internal interconnect
859 assembly. Each module is individually pretested and confirmed as a known good part
860 before being attached to that subassembly. And then ultimately support an open
861 ecosystem through this open CPX. This is an MSA that was launched over the past six
862 months. Ciena is a founding member and holds both the chair and editor positions. The
863 group now includes nearly 50 members. It should be noted that open CPX and the Vesta
864 200 both support the CPO, as well as the related variant NPO or near package optics.

865 The Ciena solution is known as the Vesta 200 CPX. Each Vesta module provides a full
866 bidirectional 6.4TB interface incorporating 32 lanes of 200G optical IMDD. A set of 16
867 modules can deliver the 100T switch interconnect that's shown on the right-hand side
868 of the slide. Ciena's individual modules combine highly specialized photonic and
869 electronic ICs that deliver superior interconnect performance, very low power, and
870 proven high reliability. Vesta also uses a 2D optical connector array for maximum
871 density. This positions the solution very well for the next-generation 200T switches. They
872 will require 32 of these modules.

873 At present, we are sampling devices for multiple key customers, both hyperscaler as
874 well as switching OEMs. Vesta technology certainly extends optical IMDD to its practical
875 physical limits, but beyond that point, we have to move into the coherent optical
876 domain.

877 Now, data center use cases do vary, but generally fall into the three categories shown on
878 the left. The first is just reach. As data centers have grown, servers and switches no
879 longer fit within a single building, turning large facilities into multi-building campuses.
880 Combined with the need to find suitable space and power, this has pushed
881 interconnect distances from a few kilometers to as much as 10 to 20km.

882 The second use case is sheer fiber utilization. At shorter distances, operators have often
883 dedicated a single fiber pair for each interface. But as both scale and distance
884 increases, multiplexing multiple interfaces across different wavelengths becomes more
885 and more important, and this will enable fiber utilization gains of 20 times or even more.

886 The third case is the optical circuit switch, or OCS. As OCS is introduced, it inevitably
887 adds optical impairments such as loss and or noise, and therefore increasing the
888 performance required from the interconnect technologies.

889 Considering the 1.6T generation, these scenarios individually or in combination push
890 beyond the practical reach of traditional optical IMDD signaling. 1600 ZR technology

891 could certainly address this requirement, but it's likely overkill for these applications
892 and adds unnecessary cost and power. The optimal answer combines the best of both
893 approaches: power and cost, close to what an enhanced IMDD would require, with the
894 performance needed to support the longer reaches and emerging use cases. And that
895 solution is Coherent-Lite.

896 Returning to our reusable building block approach, Ciena can select and optimize
897 elements from our foundational coherent toolbox to target this Coherent-Lite
898 application, covering the inside the data center, the campus, and selected scale-out
899 environments. As a leader in coherent technology, Ciena can draw on a deep portfolio
900 of algorithms, techniques, and system expertise to bring the Coherent-Lite solution to
901 market.

902 Our next-generation development includes an optimized Coh-Lite engine, which is built
903 around a dedicated two nanometer CMOS DSP, combined with a new set of optimized
904 electro-optics. This technology will support a 1.6T and 3.2T Coherent-Lite plug in the
905 OSFP form factor, as well as an ultra-high density 12.8T plug for the next generation of
906 liquid-cooled switches. Ciena also supports multiple routes to market, from complete
907 pluggable assemblies down to individual components, including a standalone DSP.

908 Building on that multiple routes to market theme, much of our interconnect technology
909 portfolio is already available at multiple levels of the ecosystem. Our systems business
910 certainly uses the full breadth of these interconnect products, from high-performance
911 WaveLogic modules through the complete pluggable portfolio. The Vesta 200 is also
912 likely to be incorporated into a full spectrum transponder in the near future.

913 Pluggable optics are sold both to end user data centers as well as network providers,
914 and to networking platform OEMs. At the module level, we also sell the platform
915 partners, and this includes high-performance WaveLogic modems that have been
916 designed into traditional network provider platforms, as well as hyperscaler home-
917 grown boxes.

918 Finally, our ASICs, including the coherent DSPs and the Nitro redriver, are sold as
919 standalone components to module and cable manufacturers.

920 To sum up, Ciena delivers differentiated technologies that solve the critical scaling
921 challenges inside the data center. While the solutions differ by application. They all
922 build on decades of optical communications and high-speed transmission expertise.
923 Our close collaboration with data center operators and hyperscalers supports co-
924 development. It aligns our solutions to the application needs and enables rapid, high-
925 volume deployment.

926 Finally, our multiple routes to market reach customers across the ecosystem,
927 maximizing both exposure and our opportunity. Thank you.

928

929 **Grant Hoffman, Executive Vice President and Chief Supply Chain Officer**

930 **Built for the Scale Ahead**

931 Hello. My name is Grant Hoffman, and I manage the supply chain at Ciena. I've been
932 with Ciena for just over two months and I'm excited to review progress and the future
933 ahead.

934 An advantage of being new is that you can quickly identify what is truly unique about a
935 company. One thing that has become clear to me is that Ciena's Supply Chain is not
936 just an operational organization. It is a strategic differentiator that has been built over
937 decades.

938 Today, I would like to walk you through five topics outlining why I believe this to be true.
939 First, our global footprint and operating model. Second, how we invested ahead of
940 growth. Third, how we are securing supply in an increasingly constrained environment.
941 Fourth, how we're scaling through vertical integration and design for excellence, or DFX.
942 And finally, how we're delivering products and services to customers at the scale and
943 speed the market now demands.

944 Before I talk about where we're going, I want to start with a broader perspective. I've
945 spent much of my career in technology and manufacturing businesses that shipped
946 tens or hundreds of millions of units annually.

947 Throughout that experience, one thing became very clear: supply chains aren't stressed
948 by normal growth. They're stressed during inflection points. Many members of our
949 leadership team have been with Ciena for two or three decades. When I asked them
950 how often they've seen supply chain conditions like this one, the answer surprised me:
951 only three times.

952 The first was the telecom expansion of the late 1990s. The second was during the
953 pandemic, and now we're experiencing the third event. What's particularly interesting is
954 that there were roughly 20 years between the first and the second events, but less than
955 five years between the second and third. The causes are different, but the outcomes are
956 remarkably similar. Supply shortages, extended lead times, capacity constraints,
957 volatility, and the disruptions are becoming larger and more frequent.

958 What separates winners from losers is whether they can learn from the last disruption
959 before the next one arrives. And that's exactly what Ciena has done. Our partner
960 engagement and geographic footprint is the result of years of strategic decisions. Our
961 R&D and supply chain teams are co-located in Ottawa, accelerating product
962 introduction and problem resolution.

963 We operate through leading global manufacturing partners. We design our own test
964 processes and assets. We maintain global distribution capabilities, and we operate a
965 global reverse supply chain network. But perhaps more importantly, many of these
966 decisions were made before they became obvious. Over a decade ago, we consciously
967 focused manufacturing strategy based on resiliency. We invested heavily in USMCA-
968 based ecosystems across Canada, Mexico and the United States. We built regional
969 flexibility before regional flexibility became fashionable. Those decisions today are
970 paying dividends. That's not luck. That's preparation.

971 One of the most important lessons from prior disruptions is that capacity cannot be
972 built overnight. By the time demand arrives, it's already too late. And that's why we've
973 been investing aggressively ahead of growth. You can see that reflected here. We've
974 increased capital investments roughly five times historic levels. We've more than
975 doubled manufacturing capacity, and we've doubled distribution capacity.

976 What's more important is that these aren't theoretical investments. They are translating
977 directly into output capabilities across key product families. Interconnects, utilizing our
978 WaveLogic and RLS. And why are we doing this? Because demand is growing faster than
979 traditional planning assumptions anticipated. The challenge for Ciena is not creating
980 demand. The challenge is ensuring our supply chain can convert that demand into
981 revenue and customer success.

982 Our strategy is straightforward. Invest ahead of the need. Build capacity before it
983 becomes constrained, and maintain the ability to scale alongside our customers.

984 The next capability is one of our strongest differentiators, and that is vertical integration.
985 Historically, network equipment companies relied heavily on third-party suppliers for
986 critical technologies. That model works in stable environments, but becomes risky
987 when demand accelerates rapidly.

988 Our response has been to expand Ciena's development and ownership of critical
989 technologies in our products. To support our growing demand, we leverage both third-
990 party solutions and internally developed components in our WaveLogic products.

991 Why does this matter? Because ownership of a large portion of the components creates
992 control. Control over roadmaps. Control over supply assurance. And improved control
993 over cost and execution, with the ultimate objective to reduce dependencies on factors
994 outside of our control and increase our ability to deliver, regardless of market
995 conditions.

996 The next three elements of supply assurance after vertical integration our long-term
997 supply commitment, extended purchase agreements, and multi-vendor sourcing. Over
998 the past several years, one lesson has become abundantly clear. When supply
999 environments tighten, relationships matter. Planning horizons matter, and balance

1000 sheets matter. With that understanding, we have intentionally moved beyond traditional
1001 industry practices for many of our most critical technologies. We're extending purchase
1002 commitments and establishing multi-year agreements aligned to our growth outlook
1003 through the end of the decade.

1004 These commitments focus on some of the most constrained technologies in our
1005 industry, particularly across photonics, optical components, and modem-related
1006 technologies. In many cases, we've gone a step further by helping suppliers fund
1007 capacity expansion and technology investments.

1008 Why is that important? Because when demand outpaces supply, suppliers make
1009 choices. They prioritize customers who demonstrate long-term commitment, strategic
1010 alignment, and the financial strength to support growth. Ciena's balance sheet and
1011 disciplined approach have allowed us to become a preferred partner, creating access to
1012 future capacity and strengthening our ability to support customer demand.

1013 At the same time, we recognize that resiliency requires more than capacity. Every major
1014 industry disruption has reinforced a simple truth: single source dependency creates
1015 risk. That is why we have deliberately expanded and qualified multiple suppliers across
1016 products and future platforms. We continue to broaden participation across
1017 strategically important components.

1018 What's particularly encouraging is that many of the leading industry technology
1019 providers want to partner with Ciena. They see the same long-term market opportunity
1020 that we do. They recognize our growth trajectory, and they understand the importance of
1021 our position within the optical networking ecosystem. The result is a supply network
1022 that becomes stronger, more resilient, and more scalable as we grow, providing a
1023 significant competitive advantage to helping ensure we can continue to deliver for
1024 customers through every phase of the market cycle.

1025 One very important item immediately stood out to me when I joined Ciena. There's a
1026 very close connection between the customer, R&D, and the supply chain. In many
1027 companies, products are designed first and manufacturing figures out how to build
1028 them later. That doesn't scale. Instead, Ciena partners closely with our customers on
1029 new product specification and partners with Supply Chain on Design for Excellence
1030 principles.

1031 Those are: Design for manufacturability. Design for cost. Design for serviceability.
1032 Design for reliability, and design for test. This means products are engineered with scale
1033 in mind from day one. Supply chain engineering, operations, and manufacturing work
1034 together early in the development process, and the benefits are clear. Faster product
1035 ramps, higher quality, improved cost, and fewer surprises when demand increases. As
1036 Dino discussed earlier, we're building on a tremendous engineering foundation. Our role
1037 is to ensure that innovation can be delivered at scale, and that's what DFX enables.

1038 As we scale the business, our ability to execute is increasingly determined by the
1039 strength of our talent and the breadth of our ecosystem. Over the past year, we have
1040 deliberately invested in both through targeted hiring, acquisitions, and strategic
1041 partnerships. By doing so, we're expanding the capabilities, expertise, and capacity
1042 required to support our next phase of growth together.

1043 Through investments in strategic talent and the expanded ecosystem, we create a
1044 powerful multiplier effect. This strengthens our ability to retain customers, capture new
1045 opportunities, and scale efficiently, creating a durable foundation for sustained growth
1046 and long-term value creation.

1047 I'm going to close with the supply chain capability that is core to customer retention,
1048 and that is service and delivery.

1049 At the end of the day, customers buy more than a product. They buy outcomes.
1050 Products delivered on time, reliable deployments, strong operational support, and the
1051 ability to scale when their business scales. We combine traditional network expertise
1052 with hyperscaler-inspired agility to deliver a seamless customer experience across the
1053 entire lifecycle, from advisory and deployment through optimization and ongoing
1054 support.

1055 With one of the world's largest optical service networks supporting customers in more
1056 than 116 countries through extensive field, repair and logistics capabilities, we help
1057 customers deploy faster, operate efficiently, and resolve issues quickly. This global
1058 scale and operational excellence strengthens customer loyalty, supports long-term
1059 growth, and reinforce our leadership position in the market. The goal is simple: make it
1060 easier for customers to grow with Ciena.

1061 The environment we are operating in today is one of the most supply-constrained
1062 periods of our industry's history. But it is also one of the greatest growth opportunities.
1063 The strategy you've heard today is built on four capabilities: We plan ahead of demand.
1064 We secure supply through strategic partnerships, long term agreements, and vertical
1065 integration. We build for scale through close integration between engineering and
1066 operations. And we deliver through flexible global manufacturing, fulfillment and
1067 service networks.

1068 These capabilities did not happen overnight. They are the result of decades of deliberate
1069 investment and lessons learned from prior industry disruptions. We diversified our
1070 footprint before resiliency became a requirement. We invested in strategic inventory
1071 and supplier relationships before capacity became constrained. And we built an
1072 operating model designed to scale alongside our customers. While we've done the work
1073 we recognize in today's unprecedented hypergrowth environment, there's always more
1074 to do.

1075 We believe Ciena has one of the most resilient and scalable supply chains in the
1076 industry, positioning us to secure capacity, support our customers, and capture the
1077 growth opportunity ahead.

1078 We have planned for this moment. We have invested for this moment, and we're ready
1079 for the scale ahead. Thank you.

1080

1081 **Marc Graff, Executive Vice President and Chief Financial Officer**

1082 **Industry Leadership Scaling Value**

1083 Hello, everyone, and thank you for joining. This is Marc Graff, Ciena's Chief Financial
1084 Officer.

1085 Today I'll update you on the journey we've been on for the past few years, the amazing
1086 results we've achieved, and our plans and financial objectives through 2029.

1087 It was a little more than a year ago when Ciena laid out its near-term financial goals,
1088 namely structurally improving gross margins, driving towards world-class working
1089 capital, and having a focused capital allocation strategy. I'm happy to report we've made
1090 excellent progress against all three of these objectives, as they have been the core on
1091 delivering the fundamentals for Ciena's business model.

1092 You've seen [adjusted] gross margins improve by over 200 basis points year-on-year
1093 through an aggressive cost-reduction program, pricing actions, and accretive product
1094 mix. These dynamics are sustainable and have reset our [adjusted] gross margin base to
1095 solidly mid-40s. Our working capital management has been a focus, not only for
1096 financial hygiene, but also to power the growth of the company.

1097 We have the resources to fund our growth over the long-term, and to ensure that our
1098 capital is deployed to meet growing customer needs while reinforcing efficiencies such
1099 as cash conversion cycles.

1100 Lastly, we have a disciplined capital allocation framework that focuses on investing in
1101 our current leadership portfolio as its top priority, and maintaining an eye to inorganic
1102 opportunities as our second priority. Being the only pure-play optical systems and
1103 interconnect solution provider and a leader in the industry, is the result of decades of
1104 deliberate investment decisions, trade-offs, and execution to meet this extraordinary
1105 market opportunity.

1106 Lastly, when capital is generated beyond these needs, we are committed to returning it
1107 to our owners, which over the last five years has averaged over 100% of free cash flow.

1108 And the results speak for themselves. Since 2024, revenue has grown an average of 27%
1109 to our current \$6.4 billion 2026 guidance. [Adjusted] operating margins have expanded

1110 by 1,100 basis points over the same time frame, and [adjusted] Earnings per Share have
1111 more than tripled. And it's this earnings potential that we are unleashing. We have
1112 tripled our [adjusted] Earnings per Share with only, only a 27% revenue CAGR.

1113 We're able to deliver these results because we're the only pure-play optical systems and
1114 interconnect provider on the planet, and that allows us massive leverage for both
1115 research and development and go-to-market activities.

1116 These impressive results have been gated only by industry supply. As we've talked
1117 previously, we are seeing acceleration in the demand for optical systems and
1118 interconnects as customers realize the network is the unlock to monetizing AI
1119 investments. Over the past two years, we've seen orders – a proxy for unconstrained
1120 demand – double from '24 to '25 and increase at least another 50% into 2026.

1121 Yet revenue – proxy for supply in a constrained environment – has increased only 27%
1122 on average over the past two years. As a result, over the same period, backlog has more
1123 than doubled from '24 to '25, and again from '25 to '26 to an expected \$10 billion exiting
1124 2026. Clearly, there is not a demand issue. What is clear is the need for the industry to
1125 accelerate its capacity requirements for a number of years to catch up to demand.

1126 And while it's typically seen as a negative indicator, industry lead times need to
1127 compress in order to establish a healthier balance between supply and demand,
1128 effectively providing a cash unlock to a growing backlog. And based on our latest
1129 outlooks, we don't see a balanced supply-demand environment returning before 2028.

1130 So, as it applies to our directional outlook for 2027, we see a very similar dynamic
1131 occurring, namely demand outstripping supply. We continue to believe that we'll exit
1132 the fiscal year with at least a 50% growth in orders to yield a minimum \$10 billion in
1133 backlog, which covers our 2027 revenue outlook and starts to build the 2028 revenue
1134 base. As importantly, we have secured the supply to enable at least a 30% revenue
1135 growth going into 2027. Our supply chain team has pulled all the levers to ensure that
1136 we can deliver this floor and position Ciena for upside, as additional supply may
1137 become available.

1138 In fact, we believe that we have secured supply to meet demand to extend the 30%
1139 trend through 2029, yielding a 2029 revenue target of approximately \$14 billion. With
1140 customer commitments and long-term supply agreements, our revenue will exceed
1141 what we believe the top four hyperscalers' Capex growth rate to be. Additionally, the
1142 competitiveness of our portfolio will enable continued share gains and exceed the TAM
1143 growth rate, driven primarily by our optical systems, such as Hyper-Rail and WaveLogic
1144 6e, and accelerated by the massive growth in our Interconnect portfolio of plugs,
1145 DCOM, modules, and the beginning of the CPO/NPO product ramps. And we have
1146 requested supply commits for '28 and '29 to exceed these growth rates. The next three

1147 years will be an inflection point for the company as it cements its leadership position
1148 and the traditional and AI-WAN and establishes its footprint inside the data center.

1149 While our ambitions are great inside the data center, we have taken a relatively modest
1150 approach to 2029 revenue from newer product lines like CPO, NPO and Coherent-Lite.

1151 As Brody and Dino reinforced, the move to optical within the data center is accelerating,
1152 and Ciena leads the industry in optical connectivity expertise. So, these products could
1153 provide significant upside to our baseline assumptions as the markets and applications
1154 developed for them. Together, there is \$11 billion of 2029 TAM for these new markets in
1155 which we can participate.

1156 We are not relying only on growing revenue to drive profitability. We believe that we can
1157 achieve 50% [adjusted] gross margins by 2029, underwritten by our industry leadership.
1158 First, as I've talked about previously, we have made excellent progress on pricing and
1159 term discussions with our customers that extend over the next three years. We'll see
1160 these actions start to impact 2027 and annualized for the remainder of the three-year
1161 period.

1162 Second, we'll see improved unit costs, as new products like WaveLogic 6n – our 800G
1163 plug – continue to ramp to mature volumes. And our engineering efforts to reduce
1164 product costs, through both design and technology improvements, continue to be an
1165 active investment.

1166 Last and perhaps most exciting is our evolving product portfolio delivering increased
1167 value, allowing structural improvements to our [adjusted] margins. Products like RLS
1168 Hyper-Rail and DCOM that deliver meaningful TCO advantages, and modules and
1169 components that have improved [adjusted] margin profiles, underwrite structural
1170 improvements to our [adjusted] margins.

1171 With our ability to deliver more value to the market, we are also seeing the benefits of
1172 our business model take hold. On top of [adjusted] gross margin improvements that I've
1173 just discussed, we are able to get significant leverage from our R&D investments in
1174 optical technology across multiple product lines. This reusable engineering model
1175 proliferates leadership discoveries that can be productized for multiple use cases. This
1176 creates an enormous amount of operating leverage for the company. On top of that, we
1177 see expanding leverage from our go-to-market efforts as well. As our largest customers
1178 are multi-segment, multi-use case customers, we're able to deepen the co-creation
1179 and co-development efforts while growing revenue, creating additional leverage. Taken
1180 together, [adjusted] gross margin improvements, and increased operating leverage, we
1181 see a path to an [adjusted] operating margin of 32% to 35% by 2029.

1182 Our earnings strength is directly translating to improving what is an already strong
1183 balance sheet. With the completion of the 2026 convertible debt issuance, we have the

1184 strategic financial capability to continue to invest in the business while taking
1185 advantage of inorganic opportunities. And with the scale that we are seeing in the
1186 business, we will continue to generate significant amounts of cash from operations,
1187 with an expectation of approximately quintupling – or 5x-ing – cash generated from
1188 operations. In addition, with the growth of earnings, I expect that we'll take our already
1189 low net leverage ratio down to zero or even below.

1190 Which begs the question: What are we doing with all this capital? Ciena has a history of
1191 being a thoughtful and deliberate steward of our owners' capital, and that trend will
1192 continue. As I noted before, our first priority is to invest in the organic opportunities of
1193 the business, such as next generation of WaveLogic technologies, new material
1194 systems, and products that compete inside the data center. In fact, we are currently
1195 estimating we'll invest somewhere between \$2.5 billion and \$3 billion in R&D over the
1196 next three years.

1197 Second, we're constantly looking across the ecosystem for technologies and talent that
1198 will accelerate our market leadership. You've seen us take decisive action with our
1199 Nubis acquisition, which has positioned us well for expanding inside the data center.
1200 And with our recently announced \$200 million Ciena Venture Fund, we'll take an even
1201 more active stand to seek out those opportunities. In total, our balance sheet has given
1202 us the flexibility of a double BB+ rating and roughly \$20 billion of M&A capacity to
1203 deploy.

1204 Lastly, once our organic and inorganic opportunities have been captured, we'll return
1205 excess capital to our owners. Over the past five years, we've returned a cumulative
1206 108% of our free cash flow to our owners. Moving forward, we are committed to
1207 returning a minimum average of 70% of free cash flow back to our shareholders, barring
1208 any major M&A activity. Overall, we're excited about the opportunities to generate
1209 leading returns for our owners.

1210 It's against this backdrop that we are realigning our reporting segments. It has been
1211 seven years since we adjusted our reporting structure, and the business has
1212 dramatically changed over that time frame. First, we've seen the rise of the hyperscalers
1213 and AI as a major demand driver.

1214 Second, our product portfolio has become even more optically focused. Third, our
1215 opportunity inside the data center represents a new growth frontier with our
1216 Interconnects portfolio.

1217 And last but certainly not least, this new structure simplifies our ability to communicate
1218 with the investment community while increasing the transparency of our different
1219 segments. Specifically, we'll report four segments that align to our \$52 billion 2029 TAM.
1220 Optical Systems, with an estimated \$21 billion TAM growing at roughly an 18% CAGR

through 2029, will include our line systems, WaveLogic and Waveserver systems, as well as the Navigator Control Suite.

Our interconnects portfolio will include the coherent plugs, DCOM (or data center out-of-band management), our Nitro and Vesta products from the Nubis acquisition, as well as optical and electrical components. This segment is expected to grow at an average annual rate of 95% through 2029, to a TAM of \$17 billion. Our global services segment will now include the Platform Software and Services revenue, and is expected to grow 15%, with a 2029 TAM of \$2 billion.

Finally, we've combined our non-DCOM Routing and Switching business with Blue Planet into Routing and Other, with a 2029 TAM of \$12 billion and a CAGR of 2%. We're excited about this new structure that will take effect with fiscal '27 and be first reported with the Q1'27 results.

So, pulling it all together, we are committed to a step-function improvement in our financial model. Our three-year revenue target will be a[n approximate] 30% CAGR through 2029, achieving a \$14 billion revenue level at roughly 50% [adjusted] gross margins.

After 15% to 18% [adjusted] operating investments, our business model's leverage will accelerate our [adjusted] operating margins to the mid 30s, with a free cash flow margin expected to be approximately 20%. And finally, we expect our adjusted tax rate to remain in the 20% range. As you can see, we are translating revenue growth into massive earnings power.

To sum it all up, Ciena is in an excellent position to continue – and in fact, to increase – our earnings power, leading to sustained returns for our owners' investments. We are committed to building on and extending our market share in these growing optical markets. We believe the inevitability of the data center opticalization plays to our strengths, as both an optical systems expert and a trusted partner to global hyperscalers, neoscalers, and service provider customers.

The foundation of our pure-play Optical Systems and Interconnects business model will drive significant earnings expansion as the world moves from copper and electrons to fiber and photons. And we're absolutely committed to being good stewards of our owners' capital, by investing first in the business and returning excess capital to our owners. Put simply, Ciena and its owners are positioned to enjoy increasing profits and returns.

Thank you.

A note on forward-looking statements and non-GAAP measures:

You are encouraged to review the Investors section of our website, where we routinely post press releases, Securities and Exchange Commission (SEC) filings, recent news, financial results, supplemental financial information, and other announcements. From time to time, we exclusively post material information to this website along with other disclosure channels that we use. Information in this presentation and related comments of presenters contains certain forward-looking statements that involve risks and uncertainties. These statements are based on current expectations, forecasts, assumptions and other information available to the Company as of the date hereof. Forward-looking statements include statements regarding Ciena's expectations, beliefs, intentions or strategies regarding the future and can be identified by forward-looking words such as "anticipate," "believe," "could," "estimate," "expect," "intend," "may," "should," "will," and "would" or similar words.

Ciena's actual results, performance or events may differ materially from these forward-looking statements made or implied due to a number of risks and uncertainties relating to Ciena's business, including: the effect of broader economic and market conditions on our business and that of our customers, including their spending; the development and use of artificial intelligence and its impact on overall networking technology spending; our ability to execute our business and growth strategies; supply chain constraints or disruptions including increased costs and lead times; the introduction of new technologies by us or our competitors; the timing and size of customer orders, their delivery dates and our ability to fulfill and recognize revenue relating to such sales; the level of competitive pressure we encounter; the product, customer and geographic mix of sales within the period; changes in foreign currency exchange rates; factors beyond our control such as natural disasters, climate change, acts of war or terrorism, geopolitical tensions or events, and public health emergencies, epidemics, or pandemics; changes in tax or trade regulations, including the imposition of tariffs, duties or efforts to withdraw from or materially modify international trade agreements; cyberattacks, data breaches or other security incidents involving our enterprise network environment or our products; regulatory changes, litigation involving our intellectual property or government investigations; and the other risk factors disclosed in Ciena's periodic reports filed with the Securities and Exchange Commission (SEC) including its Annual Report on Form 10-K filed with the SEC on December 12, 2025 and included in its Quarterly Report on Form 10-Q for the third quarter of fiscal 2026 filed with the SEC on September 3, 2026.

All information, statements, and projections in this presentation and the related earnings call speak only as of the date of this presentation and related earnings call. Ciena assumes no obligation to update any forward-looking or other information included in this presentation or related earnings calls, whether as a result of new information, future events or otherwise.

1296 In addition, this presentation includes historical, and may include prospective, non-
1297 GAAP measures of Ciena's gross margin, operating expense, operating margin, EBITDA,
1298 and net income per share. These measures are not intended to be a substitute for
1299 financial information presented in accordance with GAAP. A reconciliation of non-GAAP
1300 measures used in this presentation to Ciena's GAAP results for the relevant period can
1301 be found in the Appendix to the latest earnings presentation. Additional information can
1302 also be found in our press release filed on September 3, 2026, and in our reports on
1303 Form 10-Q and Form 10-K filed with the Securities and Exchange Commission.

1304 With respect to Ciena's expectations under "Business Outlook", Ciena is not able to
1305 provide a quantitative reconciliation of the adjusted (non-GAAP) gross margin, adjusted
1306 (non-GAAP) operating expense, and adjusted (non-GAAP) operating margin guidance
1307 measures to the corresponding gross margin, operating expense, and operating margin
1308 GAAP measures without unreasonable efforts. Ciena cannot provide meaningful
1309 estimates of the non-recurring charges and credits excluded from these non-GAAP
1310 measures due to the forward-looking nature of these estimates and their inherent
1311 variability and uncertainty. For the same reasons, Ciena is unable to address the
1312 probable significance of the unavailable information.

1313 Note Regarding Market Data: Any market definitions, market share estimates, and
1314 competitive landscape depicted herein are provided for illustrative and discussion
1315 purposes only. They are based on third-party industry analyst information and/or
1316 internal analyses derived from third-party inputs and may not reflect actual market
1317 conditions, competitive dynamics, or an exhaustive list of products, services,
1318 competitors, or technologies that may be relevant for any particular application.

1319 END OF TRANSCRIPT

1320