on 10-Nov-2026

BERNSTEIN
SOCIETE GENERALE GROUP

10 November 2026

US Networking & Communications Equipment

US Networking: How big can the CPO opportunity on Rubin Ultra really be?



Daniel Zhu
+1 917 344 8309
daniel.zhu@bernsteinsg.com

COHR disclosed that a 100 Tbps chip could be a $15k content opportunity. This in turn implies an opportunity of $9.4B in 2027 and $36.1B in 2028 for Coherent on Nvidia Rubin Ultra, with the total Nvidia ecosystem opportunity potentially reaching $47B in 2028. The CPO design Nvidia outlined at OCP points to 16 Pb/s in optical bandwidth for a 8-rack Rubin Ultra pod. Given $15k in content per 100 Tbps, this in turn points to a $5.4M opportunity for optics, up from ~$1M from scale-out pluggables on Vera Rubin NVL72

Corning has disclosed 16 fibers per GPU with copper scale-up, but our math estimates 36 for Vera Rubin. Rubin Ultra could add 240 fibers per GPU, for a total of 276, an 8x increase. We don't know how pricing would compare, but 1.5x content increase per GPU seems plausible, and the opportunity could get larger for fully optical scale-up.

It's hard to get confidence on the supply chain, given limited data. Based on our conversations with industry participants, we believe that InP substrates are the key bottleneck, and despite ~100% growth at AXTI, substrate capacity likely still lags demand growth.

Nvidia has been more aggressive about adoption timelines for scale-up optics in general and CPO in particular than most other key players, but we see other players as making progress on optical adoption as well.

Investors have been incrementally concerned about Chinese competition, with Yuanjie in particular increasingly seen as a credible competitor. While Yuanjie has clearly been making progress, we believe Lumentum and Coherent continue to have a lead. While this is clearly a risk to be monitored, we don't see it as thesis-breaking.

So why is Michael Hurlston selling? That this might just be Hurlston's value creation playbook as a CEO. Investors have worried that Lumentum CEO Hurlston is trying to sell ahead of negative news, especially an InP supply wall or Chinese competition. However, we note that all three companies where Hurlston has served as CEO (Finisar, Synaptics, and now Lumentum) have sold, and wonder if this may be Hurlston's value creation playbook.

See the Disclosure Appendix of this report for required disclosures, analyst certifications and other important information. Alternatively, visit our Global Research Disclosure Website.
First Published: 10 Nov 2026 07:06 UTC Completion Date: 10 Nov 2026 07:06 UTC

on 10-Nov-2026

BERNSTEIN TICKER TABLE

			9 Nov 2026		TTM	Reported EPS				Reported P/E (x)	
			Closing	Price	Rel.						
Ticker	Rating	Cur	Price	Target	Perf.	Cur	2025A	2026E	2027E	2025A	2026E	2027E
LITE (Lumentum)	O	USD	596.16	**700.00**	136.1%	USD	2.05	8.37	23.35	291.4	71.2	25.5
OLD				635.00								
COHR (Coherent)	O	USD	225.19	350.00	33.5%	USD	3.53	5.58	11.12	63.7	40.4	20.3
GLW (Corning)	M	USD	115.95	120.00	23.4%	USD	2.54	3.37	4.40	45.7	34.4	26.4
CIEN (Ciena)	O	USD	269.33	440.00	23.0%	USD	2.64	7.15	12.46	101.9	37.7	21.6
SPX			7,549.94									

PRICE TARGET CHANGE IN BOLD

O - Outperform, M - Market-Perform, U - Underperform, NR - Not Rated, CS - Coverage Suspended

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

We maintain our Outperform ratings on Lumentum, Coherent and Ciena, and our Market-Perform on Corning. We update our price target on Lumentum to $700 to reflect Marvell's latest share price. We maintain our price targets for Coherent, Ciena, and Corning.

on 10-Nov-2026

DETAILS

COHR disclosed that a 100 Tbps chip could be a $15k content opportunity. Based on the CPO design Nvidia outlined at OCP, this would point to a $1.0M scale-up CPO opportunity (and a $2.4M total opportunity) for an 8-rack Rubin Ultra pod. Given CoWoS-based XPU shipments, this in turn implies a $16.2B opportunity for Coherent on Nvidia Rubin Ultra.

- **We believe that Nvidia has 13,824 Tbps in scale-up optics content.** At OCP, Nvidia gave incremental details for Rubin Ultra, which will be an 8-rack scale-up pod with 72 GPUs per rack. The scale-up domain appears to be a two tier network architecture. Each switch plane consisting of an in-rack NVLink 7 CPO Switch tray with 144 ports (2x 200G lanes per port) serving as the leaf. The in-rack switch will have 72 ports connecting to the GPUs in the rack, and 72 connecting to the spine. Assuming the spine is also an 144 port switch, that suggests a 4-spine structure where each spine and leaf is connected by an 18-port bundle instead of a single link (8x18 = 144 on the spine side, 4x18=72 on the leaf side). We assume the GPU-leaf connections are copper, the leaf → spine uplinks are CPO but the spine side is pluggable[1]. 8x72 = 576 ports and 2x200G per port implies 230.4 Tbps in CPO bandwidth, and an equal amount of pluggable per plane. Nvidia also disclosed 3 TB/s of NVLink bandwidth per Rubin GPU, implying 30 planes (3 TB/s is 12 Tbps, or 30 ports of 400G), for a total of 13,824 Tbps (Exhibit 1).

- **Rubin NVL 72 has 32.4 TBps of scale-out; assuming that remains constant, an 8 rack pod would have 2,074 Tbps in scale-out optics.** Assuming scale-out per rack remains constant at 32.4 Terabytes per second (259.2 Terabits per second), optical bandwidth for 8 racks goes from 2,073.6 Tbps (all scale-out) to 15,897.6 Tbps (2,073.6 scale-out; 6,912 scale-up plugs and 6,912 scale-up CPO, which is a 7.7x increase in optical content (Exhibit 1, again).

- **Coherent disclosed a $15k content opportunity for a 100 Tbps chip, which feels conservative to us. But taking the number at face value, and assuming COHR content scales linearly, that implies a ~$2.4M content opportunity for the 8-rack pod.** At Coherent's PhotonLink launch event, it disclosed that a Compute or Switch chip with 100 Tbps I/O and 200 Gbps optical lanes (i.e. 500 lanes) could be a $15k content opportunity, implying $30 per lane. Note that this is significantly lower than the pricing we've heard from Asia contacts (~$0.4 / Gbps or $80 per 200G lane). However, one laser feeding 4 lanes means a 100T switch should need 128 high power lasers, reportedly combined into external laser source modules (ELS) with 8 lasers each, or 16 ELS per switch. The 128 lasers likely cost $30-50 each, or $4-6k, with the remainder coming from packaging, fiber assemblies, optical engines, etc. The Asia pricing is likely based on transceivers which are using a single laser per lane, explaining much of the disconnect. Accordingly, we anchor to Coherent's more conservative pricing for CPO but use Asia pricing for pluggables. We net out at $5.4M in content for an 8-rack pod, up from ~$1M from scale-out pluggables alone (Exhibit 2).

- **We estimate ~1M R300 units in 2027 and ~4M in 2028, implying an opportunity of $9.4B in 2027 and $36.1B in 2028.** The Bernstein Asia Semiconductors team (Mark Li) estimates ~1M R300 units in 2027 and 3.9M in 2028 (Exhibit 3 - industry model linked here). This in turn implies 1.8k Rubin Ultra pods with 576 GPUs each in 2027, and 6.8k pods in 2028. At $5.4M per pod, that in turn implies a $9.4B opportunity in 2027 and $36.1B in 2028. Taking a broader view, if we assume $1M in scale-out optics per 576 GPUs for all VR NVL72 and earlier GPUs, and assume $5.4M per 576 GPU pod for Rubin Ultra and Feynman, our numbers will imply an opportunity growing 91% YoY in 2028, from $24.5B to $46.8B (Exhibit 4).

- **We don't believe that AEC can substitute in this architecture.** While Credo has gained meaningful traction using its AEC technology to extend the reach of copper (instead of switching to optics), we are skeptical that AEC can work for Rubin Ultra or similar architectures. AEC can extend the reach of copper to ~7m, but an AI data center with GPUs arranged in rows would likely have further distances (Exhibit 5). Arranging in a square provides more workable distances, but faceplate capacity becomes a concern, given that the central spine rack must terminate ~5,200 copper cables (Exhibit 6).

Corning has disclosed 16 fibers per GPU with copper scale-up, but our math estimates 36 for Vera Rubin. Rubin Ultra would add 240 fibers per GPU, for a total of 276, an 8x increase. We don't know how pricing would compare, but 1.5x content increase per GPU seems plausible.

- **Corning disclosed 16 fibers per GPU for copper scale-up and optical scale-out, but our math estimates 36 for Vera Rubin**. The potentially too-simple math is that VR200 includes 32.4 TB/s in bidirectional scale-out bandwidth per rack, or

[1] Given only 4 spine switches in the rack, each can be taller than a standard tray, although there may also be some power and cooling equipment in the switch rack.

on 10-Nov-2026

129.6 Tbps one way. That is 648 lanes of 200G. Each lane needs a transmit fiber and a receive fiber per optical hop, and the two-tier fabric has two hops (NIC to leaf, leaf to spine). That works out to 2,592 fibers per rack, or 36 per GPU. Most of the gap is the hop count. A 1.6T NIC needs exactly 16 fibers for a single hop (8 lanes x 2), so Corning appears to be counting only the first tier. Our ~1.8 Tbps per GPU explains the rest.

- **Rubin Ultra would add 240 fibers per GPU, for a total of 276, an 8x increase. We don't know how pricing would compare, but 1.5x content increase per GPU seems plausible.** Rubin Ultra scale-up would add 240 fibers per GPU, for a total of ~276, a ~8x increase. 3 TB/s of NVLink per GPU in each direction is 120 lanes of 200G, or 240 fibers, but across just one optical hop, since the first GPU to NVLink switch stays copper. We don't know how pricing would compare given shorter hop lengths, but we observe that on our fiber counts, 1.5x content implies each incremental scale-up fiber carries ~15% of the dollar value of a scale-out fiber (0.5 x 36 / 120). Given that scale-up runs are shorter, a steep per-fiber discount seems reasonable. Accordingly, we think Corning's 1.5x content increase per GPU seems plausible, or potentially conservative if the pricing for inter-rack scale-up is not that different from scale-out.

- **Fully optical scale-up would add another 240, reaching ~276 fibers per GPU, a ~14x increase in fiber count.** While fully optical scale-up would likely not arrive until the Feynman generation at the earliest, and likely even later, given that copper within rack is probably still viable for the 3.2T switch cycle. Accordingly, this is a little too early to be actionable - but when fully optical scale-up does arrive, it could add another 240 fibers, reaching ~276 per GPU, a ~14x increase compared to today.

It's hard to get confidence on the supply chain, given limited data. Based on our conversations with industry participants, we believe that InP substrates are the key bottleneck, and despite ~100% growth at AXTI, substrate capacity likely still lags demand growth. Given limited data on InP wafers and substrates, it is hard to have conviction on supply numbers. However, based on our conversations with industry participants, we believe that InP substrates are really the key bottleneck. For the substrates, Sumitomo and AXTI (which is US listed but whose capacity sits in China) are the key suppliers, with probably ~40% share each as of 2025, with the remainder being held by smaller metals and mining companies in China and Japan (JX Advanced Metals, Yunnan Germanium, etc.). AXTI has stated that capacity will double by year-end 2026 vs. Q4 2025, and then double again in 2027; given a 9-12 month lag between substrate capacity adds and end customer deliveries, this points to industry supply doubling in 2027 and again in 2028. For its part, Sumitomo announced plans to 3x its FY 24 capacity by FY 28[2], pointing to capacity growth around ~30% CAGR. If we assume the other suppliers grow capacity 33% YoY in 2027 and 2028, in line with Sumitomo, that implies industry blended average capacity grows 60% in 2027 and 67% in 2028 (Exhibit 7), with supply growth for 2028 continuing to lag demand growth (while 2027 supply is likely already booked). Accordingly, we believe supply will remain constrained through 2028, and don't expect a sharp collapse in pricing unless we see broader softening in demand across the AI ecosystem.

Nvidia has been more aggressive about adoption timelines for scale-up optics in general and CPO in particular than most other key players, but we see other players as making progress on optical adoption as well. We note that Nvidia has been more proactive about CPO adoption than other players, especially compared to the hyperscalars, who have generally been more hesitant about scale-up adoption in general and also tended to favor NPO over CPO. That said, we do see increasing adoption in the rest of the ecosystem - AWS is reportedly the other player adopting CPO in 2027 (for Trainium 4), while Chinese media has reported that Google placed a 12M order for NPO modules for TPU scale-up ($12-15B in spend), split 60/40 between Innolight and Eoptolink. Accordingly, we believe the broader industry could see 2028 demand growth above ~70% supply growth, and potentially closer to Nvidia's ~90% growth. However, we believe this is an area where we need to continue to build out our views on the non-Nvidia ecosystems.

Investors have been incrementally concerned about Chinese competition, with Yuanjie in particular increasingly seen as a credible competitor. While Yuanjie has clearly been making progress, we believe Lumentum and Coherent continue to have a lead. While this is clearly a risk to be monitored, we don't see it as thesis-breaking.

- **Investors are increasingly worried about Chinese competition.** While US investors have often seen Lumentum and Coherent as a duopoly, given that they are the only two US-listed pure plays. However, the laser industry is actually quite fragmented, especially on lower power lasers, including AAOI and Broadcom, a range of Asian competitors (e.g. Mitsubushi, Sumitomo) as well as in-house capacity at major customers (Nokia / Infenera, Ciena, etc.). We generally see lower power lasers (esp. 100 mW and below) as being pretty commoditized. Lumentum and Coherent's true technology advantage is on

[2] Sumitomo's fiscal year ends in March

on 10-Nov-2026

the high end of the market, but with Chinese vendors (as well as Sumitomo and Mitsubushi) increasingly able to make higher power EMLs, investors are increasingly worried about the competitive threat.

- **Yuanjie is increasingly seen as a credible competitor, but we believe it still lags Lumentum and Coherent.** Yuanjie Semiconductor Technology (SSE: 688498) is increasingly seen as the leader in China. However, we note that its product specs actually continue to lag Lumentum's. Yuanjie's flagship product is a 200G EML laser, which is likely somewhat comparable to Lumentum's key products in terms of raw power, although industry sources believe Yuanjie still lags on linewidth. Importantly, Yuanjie has begun listing the product in 2025, and is still in the sampling and validation phase with key customers (notably with Nvidia). By contrast, first demoed 200G EMLs in 2022 and is expected to start qualifying 400G EMLs in 2028, which points to Yuanjie still being ~3 years behind even accounting for the potential for later movers to move faster. We also observe that Yuanjie (and other peers) may be benefiting from the Marvell-Lumentum deal, with key Lumentum customers such as Ciena reportedly being increasingly exploring other supply options following the deal.

- **Other major Chinese players including Source Photonics, Accelink, and Everbright.** Beyond Yuanjie, key Chinese players include: 1) Source Photonics, which is a US-domiciled company with most operations in Taiwan, but which was acquired by Dongshan Precision Manufacturing (SHE:002384), a China-listed PCB manufacturer, in September 2025; 2) Accelink 3) Everbright 4) Shijia and others. Notably, Dongshan has reportedly been aggressive in securing access to InP substrate supply from Sumitomo, the key overseas supplier.

- **Investors should clearly monitor this risk, but we don't see it as thesis breaking.** While Chinese competitors are clearly a credible threat in the space, we believe Lumentum and Coherent continue to have a technology advantage at the high end of the market. While they are likely to lose share at the lower end, given less aggressive supply adds, we believe there is sufficient high end demand for them to continue to perform well.

So why is Michael Hurlston selling? Investors have worried that Lumentum CEO Hurlston is trying to sell ahead of negative news, especially an InP supply wall or Chinese competition. We observe that this might just be Hurlston's value creation playbook as a CEO. We note that Lumentum CEO Michael Hurlston's first job as CEO was at Finisar, which he sold to II-VI (which is Coherent today) in 2019. Hurlston then became CEO of Synaptics. While Hurlston left in 2025 to take on the Lumentum job, Synaptics is now being acquired by onsemi, and the proxy stated that management has been exploring potential combinations in recent years without reaching a deal, suggesting that Hurlston already began to prepare Synaptics for a potential sale during his tenure. Given this track record, we wonder if finding a buyer is simply Hurlston's value creation framework as CEO. While Lumentum has become too big for there to be a natural buyer, we note that when Hurlston was named CEO, it was a $6B company, close to the transaction value of Finisar ($3B) and Synaptics ($7B). Preparing Lumentum for a potential sale may well have been part of his mandate, and given that the proxy estimated his change-in-control payout at $283M, he could be well compensated for doing so.

We maintain our Outperform ratings on Lumentum, Coherent and Ciena, and our Market-Perform on Corning. We update our price target on Lumentum to $700 to reflect Marvell's latest share price. We maintain our price targets for Coherent, Ciena, and Corning.

- **On net, we continue to be constructive on the sub-sector.** Optical demand is clearly growing very rapidly, and while the level of visibility on supply is lower, we continue to believe that demand exceeds supply through 2072 and 2028.

- **We maintain our $350 price target and Outperform rating on Coherent.** We increasingly feel that we can underwrite the trajectory of the business rather than worrying that LITE is trying to sell ahead of bad news. We also see potential for Coherent to gain share following the Marvell-Lumentum merger. While investors are clearly incrementally cautious on AI, we also continue to be positive on the theme longer term. That said, we also note that most of the major upside levers ($4B revenue quarter, 45% gross margins, scale-up optics) fall into the end of CY 2027 or CY 2028, putting it into FY 2028, and accordingly we see few catalysts before OFC in March. On net, we maintain our Outperform and our price target of $350. For more on how we've thought about the stock during the quarter, see US Networking: Scaling the bandwidth wall. Initiating with a positive view.

- **We rate Lumentum Outperform. We value the stock at the face value of Marvell's bid (now ~$700), and estimate price implied probability at 70%.** We update the Lumentum deal math: We reflect the latest Lumentum and Marvell share prices of $615 and $157, respectively, implying a face value of the bid at $708, or $675 adjusting for the time value of money. We also update our break prices assuming that the Marvell shareholder vote is the main risk to the deal. We also observe that the investor base is likely increasingly turning over to merger arbs, and arbs would likely want to be conservative on fundamentals when setting a break price for arb math purposes. With that in mind, we assume Lumentum falls to $470 in

on 10-Nov-2026

a break scenario (20x cons. $23.7 in FY 27 EPS). Although we would actually think that Lumentum would be worth more as a standalone company, we think management would likely walk away from a break with limited credibility, and the FY28 EPS guidance in particularly will be seen as suspect. We value Marvell's break price at $150, in line with the current share price. While we don't cover Marvell, and don't have a strong view on the fundamentals, our impression from investor conversations is that Marvell management already has pretty low credibility with investors as is, and so in that context would expect the shares to be flat to up in a break scenario. Based on this, our implied probability of deal close nets out at 65% (Exhibit 8). We continue to value Lumentum at the face value of the bid of 4.5 shares, and update our price target to $700 ($635 previously) to reflect Marvell's share price.

- **We maintain our $440 price target and Outperform rating on Ciena.** Ciena stock has sold off together with the rest of the AI ecosystem around incremental concerns about AI lab health and power and land constraints. However Ciena's AI opportunity is largely from scale-across, which should actually benefit from power and land constraints.

- **We maintain our $120 price target and Market-Perform rating on Corning.** Aggregating our views on the sector, we model Corning reaching $28B in 2028, which is somewhat below what Corning's growth framework implies, and below consensus. Accordingly, we see the growth framework as aggressive, and content upside in optical communications as largely priced in, unless we can get conviction that pricing on scale-up optics will be more in line with scale-out than what Corning's guidance would imply.

on 10-Nov-2026

EXHIBIT 1: **Rubin Ultra Optics Estimate**

Rubin NVL 72	
Lane speed (Gbps)	200
Lanes per port	2
Port speed (Gbps)	**400**
Scale-out (TBps)	32.4
Bits per byte (b/B)	8.0
Scale-out (Tbps)	**259.2**
NVLink bandwidth per GPU, bidirectional (TBps)	3.0
NVLink bandwidth per GPU, bidirectional (Tbps)	24.0
NVLink bandwidth per GPU, unidirectional (Tbps)	**12.0**
Switch planes	**30**
Rubin Ultra	
Copper ports per GPU	2.0
Pluggable ports per GPU	1.0
CPO ports per GPU	1.0
Total Ports per GPU	**4.0**
GPUs per rack	72
Racks per pod	8
GPUs per pod	**576**
Bandwidth (Tbps)	
Pluggable Bandwidth	6,912.0
CPO Bandwidth	6,912.0
Scale-up Optical Bandwidth	**13,824.0**
Copper Bandwidth	13,824.0
Total Scale-up Bandwidth	**27,648.0**
Scale-out bandwidth	2,074
Scale-out optical penetration	100%
Total Optical Bandwidth	**15,897.6**

Source: Company Disclosures, Bernstein estimates and analysis

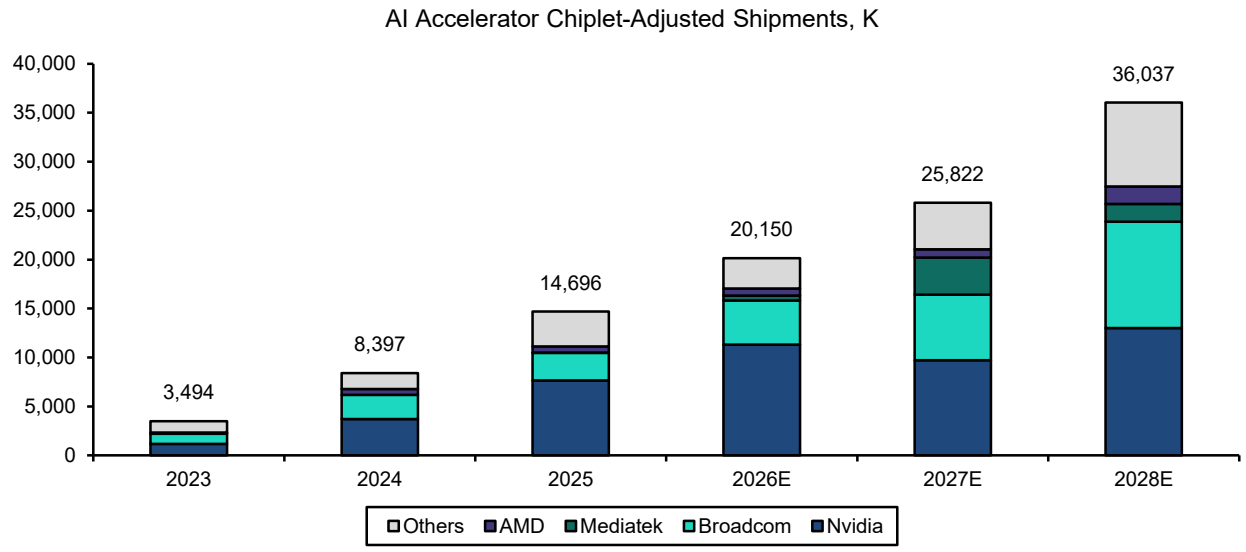
EXHIBIT 2: **Coherent Content Opportunity on Rubin Ultra**

Coherent content	
Coherent PHOTONLINK content per 100 Tbps ($K)	15
PHOTONLINK content ($/Gbps)	0.15
Pluggable content ($/Gbps)	0.48
CPO Content per Rubin Ultra pod ($K)	1,037
Scale-up Pluggable content per Rubin Ultra pod ($K)	3,318
Scale-out Pluggable content per Rubin Ultra pod ($K)	995
Total optical content on Rubin Ultra ($K)	**5,350**

Assumed PHOTONLINK pricing for CPO
Source: Company disclosures, Bernstein estimates and analysis

on 10-Nov-2026

EXHIBIT 3: **We expect accelerator chiplet volumes to double 2025-28E and switch ports to grow proportionally**



Source: Bernstein Asia Semiconductors and Equipment & Global Memory team (Mark Li)

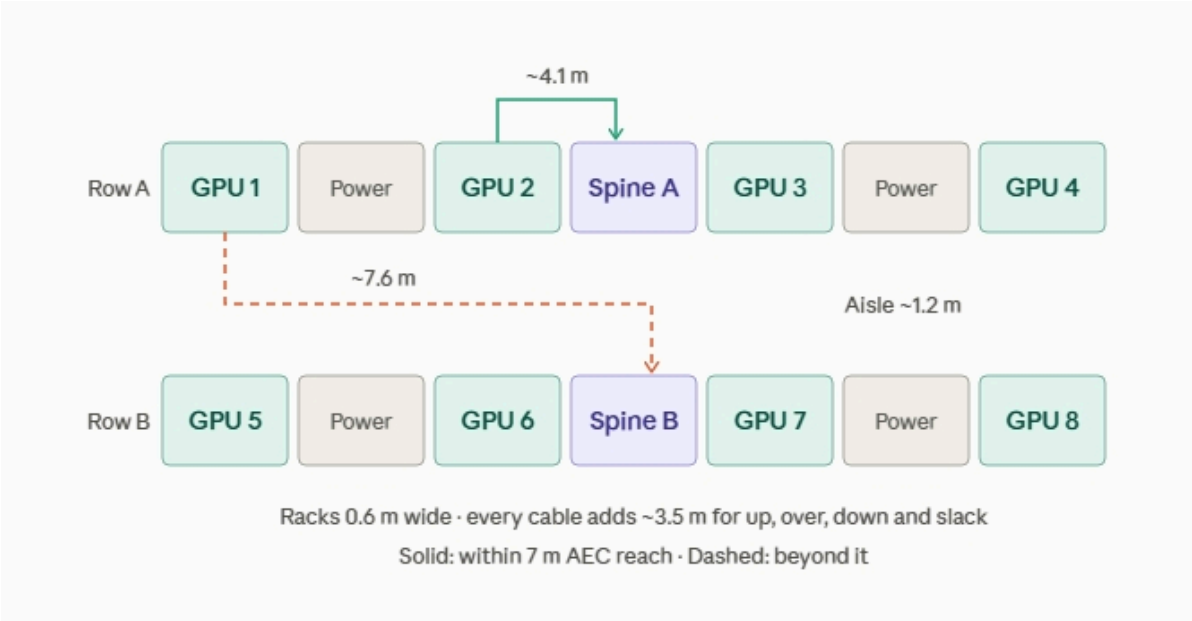
EXHIBIT 4: **Total Optical TAM for Nvidia Ecosystem**

	2026E	2027E	2028E
Units (K)			
R200 and earlier	11,311	8,714	1,802
R300 and F100	0	1,015	4,701
Total	**11,311**	**9,729**	**6,504**
Content per pod ($K)			
R200 and earlier	995	995	995
R300 and F100	5,350	5,350	5,350
Total	**995**	**1,450**	**4,143**
Content per GPU ($K)			
R200 and earlier	1.73	1.73	1.73
R300 and F100	9.29	9.29	9.29
Total	**1.73**	**2.52**	**7.19**
YoY		*46%*	*186%*
Nvidia TAM ($B)			
R200 and earlier	19.5	15.1	3.1
R300 and F100	-	9.4	43.7
Total	**19.5**	**24.5**	**46.8**
YoY		*25%*	*91%*

Source: Bernstein Asia Semiconductor Team (Mark Li), Bernstein estimates and analysis

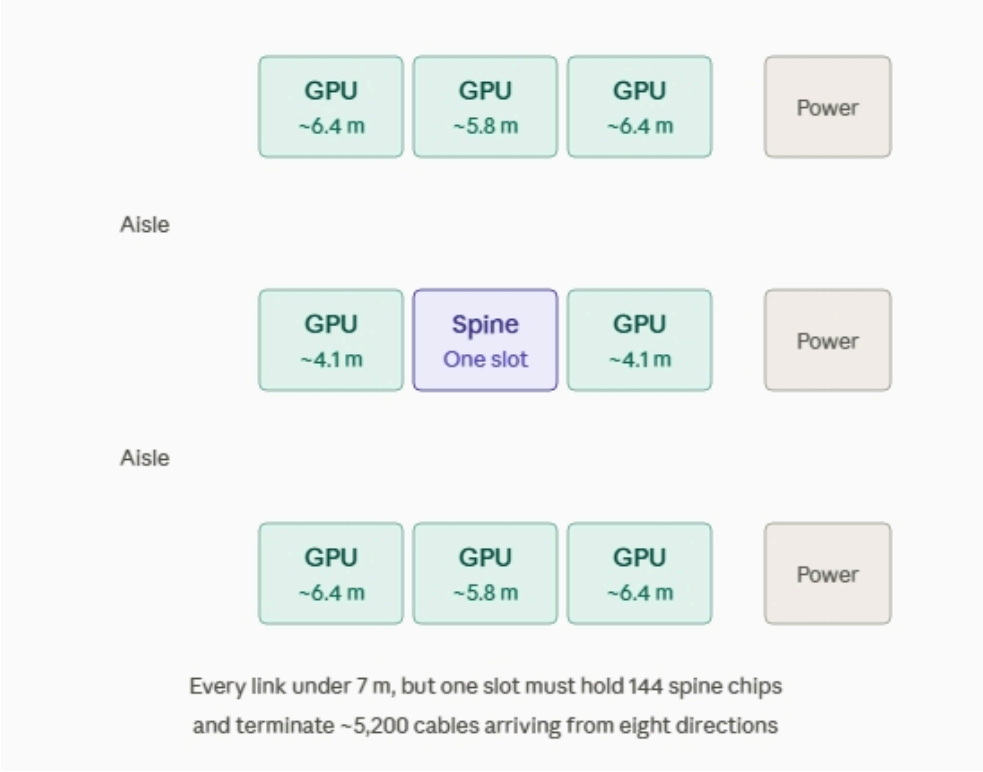
on 10-Nov-2026

EXHIBIT 5: **GPU racks arranged in rows would require distances >7m**



Source: Bernstein estimates and analysis

EXHIBIT 6: **GPU racks arranged in a square would require one slot to hold ~5,200 cables arriving from eight directions**



Source: Bernstein estimates and analysis

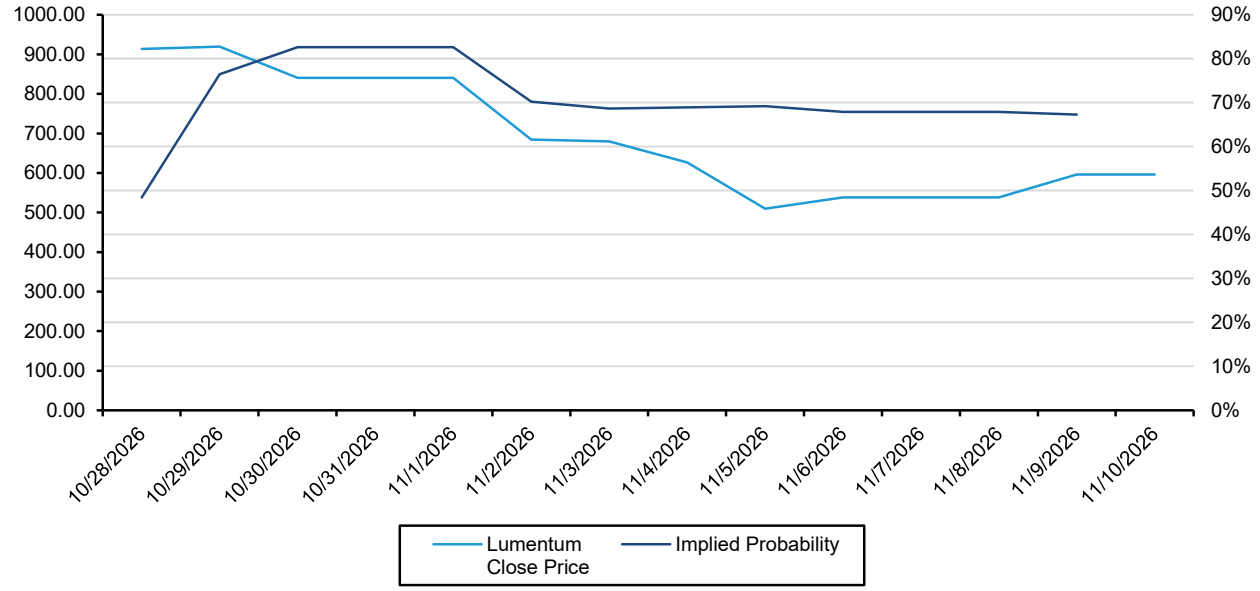
on 10-Nov-2026

EXHIBIT 7: **Illustrative Supply Ramp as % of total 2025 capacity**

	2026E	2027E	2028E
Sumitomo	0.4	0.5	0.7
AXTI	0.4	0.8	1.6
Others	0.2	0.3	0.4
Total	**1.0**	**1.6**	**2.7**
YoY		*60%*	*67%*

Source: Bernstein estimates and analysis

EXHIBIT 8: **LITE Stock Price & Implied Probability of Closing**



Implied probability assumes break price of $822.56 for LITE and $226.49 for MRVL, close date of 1/1/2028, discounted back to current date at a 4.38% risk-free rate,
Source: Bloomberg (updated as of November 10, 2026), Bernstein analysis