semianalysis

June 9, 2026

Powered Down, Lights Off. 800VDC Pushout and CPO Delays

//Multi-Vertical Note

8 minutes

By Konrad Wang, Nigel Chiang, Nick Doyle, Doug OLaughlin and Daniel Nishball

Executive Summary

- **Mass adoption of Nvidia's native single-ended 800VDC design**, with 800VDC volume shipments pushed to 2028+. ±400VDC, a separate HVDC architecture, remains on schedule for hyperscaler's ASIC deployments.
- **Industry chatter suggests that hyperscalers are pushing back on adoption of the single-ended 800VDC architecture driven by Nvidia, given that 800VDC is not a necessity for Rubin which will still take 50VDC.** They believe that taking grid power at 350-450VDC, stepping up to 800VDC and back down to 50VDC to feed the compute tray is inefficient. Instead, we see hyperscalers increasingly pushing for power to be delivered at a higher voltage before being stepped down to the compute tray.
- **±400VDC is still set to proceed in 2H26 as expected, primarily for hyperscalers' in-house ASIC deployments.** We expect ±400VDC sidecar orders to land late this year, with manufacturing ramping in Q1 2027.
- Net-net, this means that **sidecar volumes that were initially going to be driven by Rubin Ultra/Kyber deliveries will be moved into the 2028 window.**
- **2027 CPO expectations look too aggressive – Co-Packaged Optics (CPO) will be delayed versus current Street expectations**. We expect to revise down scale-out CPO shipments in 2026 and 2027 – meanwhile we have always had Scale-up CPO ramping in earnest from 2029 even though the street has been calling for 2028 or even 2027 ramp timelines for Scale-up CPO. Instead – many NPO

projects will be ramping, but this may play more into the hands of transceiver vendors.

- For **Scale-out CPO switches, CPO System-level integration is the gating concern with yield economics remaining a major barrier.** At an optimistic **95%** optical-engine attach yield and **32 COUPEs** per ASIC, system yield is only **~19%**. The Street models 70–100k+ Scale-out CPO switches produced annually by 2027, but these issues are leading to much lower current production levels – putting us off track with respect to those numbers. We will likely revise down Scale-out CPO switch shipments in upcoming Networking model updates.
- Overall, we came away from Computex incrementally more positive on Amphenol, Vertiv, Forgent Power Solutions, Legrand and FormFactor. Conversely, we are incrementally more negative on Lumentum, Himax, Navitas, and Wolfspeed.

Recent Winners become Near-Term Losers and Vice Versa

800VDC and CPO are two themes which have been dominating the semiconductor market narrative this year. We've been positive on both, with dedicated Newsletters on CPO & 800VDC.

With this note, we expect to reset investor expectations and especially sentiment. We think **positioning is stretched** to a degree that matters more than the fundamentals here. The "bottleneck" trade (especially in photonics/optics and power semis) has become the most crowded long in the AI complex, likely funded by shorts against the large-cap platform owners (NVDA, AVGO).

Many of these names sit at or near all-time highs on momentum, with sentiment extremely bullish and risk tolerance maxed. That asymmetry is the opportunity: when the most-extended, most-levered part of a trade gets a negative timing confirmation, the unwind is violent, and the funding shorts squeeze back, **independent of the long-term bull case**, which we still hold. We do not believe this dynamic applies to memory, which is a consistent bottleneck with tighter supply-demand.

Our **biggest Computex 2026 takeaway is that both CPO and 800VDC have a high likelihood of delays** compared to original ramp expectations. CPO system level integration is more complicated than the market appreciates, and we believe shipments will be meaningfully below prior expectations. Industry chatter suggests hyperscalers are pushing back on the single-ended 800VDC architecture driven by Nvidia, given it is not a necessity for Rubin, with volume shipments now pushed to 2028+.

HVDC roadmap in flux – ±400VDC still happening 2H26 as expected, but 800VDC is pushed out

Contrary to our prior expectation of 800VDC being a 2027 story, we now see adoption pushed out to 2028+. Nvidia is still pushing the industry toward 800VDC, and the underlying logic still holds: increased rack power, centralized power delivery, and reduced conversion losses by stepping voltage down closer to the compute tray.

However, penetration will be low in 2H26/2027 because it is not necessary for Vera Rubin. Rubin Ultra and Feynman are more likely to be where 800V becomes necessary, with Rubin Ultra designs only being finalized later this year.

- **Industry chatter suggests that hyperscalers are pushing back on the adoption of 800VDC** given that Nvidia seems to have multiple power architecture options for the Rubin generation, and 800VDC is not a necessity. In their view, taking grid power at 350-450VDC, stepping up to 800VDC and back down to 50VDC to feed the compute tray is highly inefficient. We believe that hyperscalers are increasingly pushing for power to be delivered at higher voltage before being stepped down to the compute tray.
- **Technology adoption pushout, not a cancellation:** Rubin Ultra/Kyber was designed with native 800VDC input to the GPU compute tray, but this has been delayed. We flagged this in our Computex preview, and our Accelerator Team pushed out Kyber expectations. However, 800VDC is a must when the GPU compute tray input requires 800VDC. We remain of the view that at very high tray and package power levels, native HVDC distribution becomes compelling. While not our base case, we note there are benefits to a Rubin Ultra/Oberon design which takes 800VDC input, given that the compute tray will likely exceed 15kW TDP.
- **±400VDC proceeding in 2H26 as expected:** It is important to note that ±400VDC is not the same as 800VDC. The former is a high-voltage direct-current (HVDC) architecture that the hyperscalers are pursuing themselves to 1) centralize power delivery at higher voltage for better efficiency and 2) avoid AC-DC-AC conversion losses with the UPS. We believe that the ±400VDC rack architecture is primarily for hyperscalers' in-house ASIC efforts, and this is still on track. We expect ±400VDC sidecar orders to land late this year, with manufacturing ramping in 1Q27. That said, we are not ruling out hyperscalers using the ±400VDC sidecar to support Nvidia hardware given that the DC-DC power shelves should allow for this.

800VDC Read-Across to Stocks in our Coverage

- **Power rack suppliers:** We see the delay as neutral for power rack suppliers Vertiv (VRT), Delta (2308 TT), and Lite-On (2301 TT), as the sidecar/power rack transition is happening regardless of whether the bus is ±400V or single-ended 800V.
  - Vertiv is particularly well-positioned because the 800VDC delay extends the life of their large UPS business. At Computex, Vertiv demoed a grey space power rack design which locates the PSU and PDUs outside of the IT room to preserve precious white space, that we found compelling.
- **Grey-space electrical equipment suppliers:** We see this as positive for grey space suppliers Forgent Power Solutions (FPS), Legrand (LR FP), Schneider

Electric (SU FP), Hammond Power Solutions (HPS.A CN) and ABB (ABBN SW) that were at risk of losing LV transformer, LV switchgear and busway content. The pushout of 800VDC directly translates to increased upside and a longer growth runway for these names.

- **Board-level VRM / power semiconductors:** Silicon-based passives and power semiconductors win regardless of architecture. Both ±400V and 800V sidecars need superjunction MOSFETs, resistors, inductors, and capacitors for pulse-load transient absorption in power shelves, BBUs, and power racks. Vishay (VSH) supplies superjunction MOSFETs, resistors, and inductors. MLCC content scales linearly with rack TDP: Murata (6981 JP), SEMCO (009150 KS), Yageo (2327 TT), TDK (6762 JP). Infineon (IFX GR) appears to be the best-positioned power semi because they are hedged across Si superjunction, SiC, and GaN. At the board/compute tray level, VRM smart power stages step 48V/12V down to sub-1V at the GPU die, and this is unaffected by 800VDC delay. The 48V-to-sub-1V conversion chain stays the same whether the upstream architecture is AC, ±400V, or 800V.
  - Players: MPS (MPWR), Renesas (6723 JP), Infineon, with TXN and ON qualifying as new entrants.
- **Wide-bandgap pure-plays:** We believe this puts wide-bandgap (WBG) suppliers in an awkward spot. ±400V uses some compound semi (GaN is more efficient at that voltage, though SiC's supply chain is more mature), so the ±400V transition is close to neutral for the WBG vs silicon debate. 800VDC is where WBG content really inflects; a pushout means the incremental WBG content uplift that would justify current multiples on pure-play names like Wolfspeed (WOLF) or Navitas (NVTS) has no meaningful near-term catalyst.

CPO: Costs and Physics are More Prohibitive than Expected

We are positive about copper and pluggable optics relative to CPO. For Scale-up, **we have always seen the huge inflection in volumes coming in 2029** as key projects from AWS, AMD and Feynman ramp in earnest. Optical Engines on Interposers will only be fully ungated and start ramping in that time period, opening the door to true ubiquity of CPO, but only after 2029 and 2030.

In contrast, the Street has misinterpreted COUPE Optical Engine production volumes, first incorrectly assuming it was for a CPO Rubin Ultra Kyber – we debunked this idea well before GTC 2026's announcements made official that this would not happen. We do have volumes of the VRU NVL576 in 2027 and 2028, but this will only use CPO from switch to switch and not to GPUs, and we don't see project being large enough to move the needle and most on the Street model too aggressive of a timeline. Overall, 2029 is a more sensible target for Scale-up CPO volume shipments and street expectations of game-changing jumps in CPO volume in 2028 look optimistic. Our first 2029 estimates from early April took this view, and we see the significant inflections happening in 2029 and beyond, coinciding with Feynman's ramp. In the meantime, many NPO projects will be ramping in decent volume, but this may play more into the hands of transceiver vendors.

Turning to Scale-out, we will likely revise down CPO switch forecasts in upcoming Networking model updates. Scale-out CPO will still lead, but yield issues must be solved first. In addition, we had assumed 85% penetration for CPO switches at Neoclouds. We continue to think that adopting CPO for Neoclouds makes sense, but we have not yet found the adoption rate to be fully universal, though we think adoption will increase over time. The Street currently models 60-100k+ Scale-out CPO switches shipped by 2027, driven by expectations of a sharp scale-out ramp and strong adoption rates at Neoclouds, and we think today's pace of scale-out production is too low to reach this bogey.

The market has been focused on InP based CPO lasers as a structural constraint for the ramp of CPO. While lasers remain an important bottleneck, this framing misses other key deployment gates, specifically COUPE developments. While COUPE development remains on track for Nvidia and other adopters such as Broadcom, AMD, and Ayar Labs, **system-level integration with COUPEs remains challenging**.

- **We expect Spectrum 6 CPO output (SN6810, SN6800) to slip by more than two quarters.** NVIDIA's Spectrum 6 CPO, the first 102.4T switch with second-gen COUPEs, recently showed >3.5 dB of insertion loss in on-board system-level testing, consuming the entire optical channel budget. This is worse than Spectrum 5 CPO ("Agora"), which uses the same COUPE count per ASIC (32), the same connectors, FAU, and similar system design. Neither NVIDIA nor TSMC has identified the source of the problem; efforts have shifted toward fundamental assembly-process redesign.
- **The yield math is punishing.** Industry chatter puts optical engine attach yield at an optimistic 95% today. At 32 COUPEs per Spectrum 6 ASIC, that compounds to 19% system yield (0.95^32). Every COUPE must be perfect post-coupling; there is no rework path on a soldered switch substrate. The industry needs ~99.5% attach yield per engine to make volume economics work, delivering ~85% system yield at 32 engines.
- **NVIDIA's Quantum X3450 (InfiniBand CPO) is in comparatively better shape:** only 3 COUPEs per module, so defective modules can be screened and the best ones selected for assembly. Even if real attach yield is below 95%, which is highly probable, the small module granularity makes the economics manageable.

CPO delays reinforce the case that copper remains the primary interconnect for scale-up networks, with pluggable optics continuing to serve scale-out, supporting sustained demand growth for both categories. We defended AECs earlier this year, and noted last week another positive Credo development.

CPO Read-Across to Stocks in our Coverage

Networking Industry Comparables Table

Name	Ticker	Industry Focus	Market Cap ($B)	CY26 YTD Return	Return Since LTM Peak	NTM P/E	NTM EV/EBITDA
Credo	CRDO US Equity	AEC, Optics	$41	54%	-10%	35.4x	30.1x
Astera Labs	ALAB US Equity	AEC, Switch	$59	108%	-7%	96.8x	102.1x
Marvell	MRVL US Equity	AEC, ACC, ASIC	$253	240%	-11%	56.9x	47.0x
Broadcom	AVGO US Equity	AEC, ACC, ASIC	$1,887	15%	-20%	24.5x	18.7x
Semtech	SMTC US Equity	ACC	$15	115%	-8%	52.0x	41.6x
Macom	MTSI US Equity	ACC	$28	111%	-14%	58.8x	50.9x
Amphenol	APH US Equity	AEC, ACC, Connectors	$177	6%	-14%	27.6x	16.1x
TE Connectivity	TEL US Equity	AEC, ACC	$60	-10%	-19%	16.9x	11.2x
Lumentum	LITE US Equity	Optics	$70	143%	-18%	51.5x	29.7x
Coherent	COHR US Equity	Optics	$79	118%	-9%	49.7x	33.4x
Applied Optoelectronics	AAOI US Equity	Optics	$16	464%	-16%	61.1x	46.1x
Eoptolink	300502 CH Equity	Optics, AEC	$783	82%	-4%	37.2x	25.7x
Innolight	300308 CH Equity	Optics	$1,316	93%	-11%	40.0x	30.0x
Arista Networks	ANET US Equity	Switch, Optics	$197	19%	-13%	39.4x	29.7x
Cisco	CSCO US Equity	Switch, Optics	$489	61%	-5%	26.3x	19.6x
Accton	2345 TT Equity	Switch	$1,389	109%	-8%	29.0x	20.6x

Source: SemiAnalysis, Bloomberg (as of 6/8/2025)

- **Relative to market expectations, we** are **most positive on copper names** Amphenol (APH), Semtech (SMTC), and MACOM (MTSI) – some of these Copper-exposed companies have been largely ignored despite a revenue TAM that we believe will clearly 3x in the next few years.
- We still like Optics companies levered to pluggable transceivers and DSPs rather than CPO, namely (Marvell (MRVL), Innolight (300308.SZ), Eoptolink (300502.SZ), Tower Semiconductor (TSEM), STMicroelectronics (STM), and Astera Labs (ALAB) as we see the pluggable transceiver TAM continuing to grow at a rapid pace. The Networking Model has discussed the growing TAM for optics at projects like Google's TPUv9 and v10. Indeed, many NPO projects will also be ramping in decent volume in the next few years, but this may play more into the hands of transceiver vendors.
- We are **more conservative on names where CPO volume is a material part of the bull case**, including Lumentum (LITE), Coherent (COHR), Himax (HIMX) and Applied Optoelectronics (AAOI), though we still think CPO has a "bright" future once inflection point in 2029-2030 is reached.
- **We continue to be positive on the CPO testing theme.** As we detailed in our CPO test equipment landscape note and subsequent update, CPO test equipment is a picks-and-shovels beneficiary that should see procurement ahead of CPO volume. We remain most constructive on Teradyne (TER), the frontrunner in NVIDIA qualification and ficonTEC partnership; FormFactor (FORM) for PIC wafer probing; Chroma (2360.TT) for die-level and system-level test; and Hon Precision (7769.TT) for handler monopoly in AI/HPC final test. This underscores why testing is so critical: with system-level integration already consuming the entire channel budget at 32 COUPEs, the only lever left is ensuring every optical engine entering assembly is flawless.

Disclaimers

Analyst Certifications and Independence of Research.
Each of the analysts whose names appear in this report hereby certify that all the views expressed in this Report accurately reflect our personal views about any and all of the subject securities or issuers and that no part of our compensation was, is, or will be, directly or indirectly, related to the specific recommendations or views in this Report. SemiAnalysis LLC (the "Company") is an independent equity research provider. The Company is not a member of the FINRA or the SIPC and is not a registered broker dealer or investment adviser. SemiAnalysis has no other regulated or unregulated business activities which conflict with its provision of independent research.