on 08-Oct-2026

BERNSTEIN
SOCIETE GENERALE GROUP

US Networking & Communications Equipment

US Networking: Connecting the dots... A primer on networking equipment in the age of AI



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



Mark C. Newman
+1 212 845 7822
mark.newman@bernsteinsg.com



Viola Chen
+1 917 344 8614
viola.chen@bernsteinsg.com



Phoebe Sun
+1 917 344 8481
phoebe.sun@bernsteinsg.com

We initiated on US Networking and Communications Equipment on 9/29 after the close. This note serves as a primer for investors new to the space. For more details on our views, see our initiation report: US Networking: Scaling the bandwidth wall. Initiating with a positive view

Networking allows other devices to communicate and build larger device ecosystems. Networking refers to devices that don't process or store data on their own, but instead help other devices communicate. Networking devices in particular solve 4 main issues: forwarding, performance, reliability, and security.

Networking companies are primarily product companies. The key product categories include switching, enterprise networking, optical transport, and transceivers. Companies in our coverage also supply components, notably optical fiber and optical modules. Networking primarily serves three main product categories: enterprises, service providers, and hyperscalers.

Networking has historically been one of the worst performing subsectors of tech. Networking stocks have historically underperformed, driven by low growth. Low growth in turn was in large part due to high exposure to service providers; the end customers did not grow revenue or capex, pressuring growth of suppliers.

Networking has re-rated on the AI opportunity. As compute performance has outscaled network performance, this has created a "bandwidth wall" in which networking is increasingly the bottleneck to overall AI performance. This has in turn lead to the emergence of new, AI-focused network architectures. These include: Scale-up, which refers to densely connecting accelerators in the same rack so that they behave as a single node. Scale-out, which refers to connecting different racks and different nodes. Scale-across, which refers to direct connections enabling separate data centers to behave a single cluster. Scale-in, which connects end users to accelerated compute.

Networking can seem technical, but not all technology matters. We believe investors should be focusing more on the space, and in particular on technologies that can actually drive the stocks.

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: 08 Oct 2026 08:08 UTC Completion Date: 08 Oct 2026 08:08 UTC

on 08-Oct-2026

BERNSTEIN TICKER TABLE

Ticker	Rating	Cur	7 Oct 2026 Closing Price	Price Target	TTM Rel. Perf.	Cur	Reported EPS 2025A	Reported EPS 2026E	Reported EPS 2027E	Reported P/E (x) 2025A	Reported P/E (x) 2026E	Reported P/E (x) 2027E
ANET (Arista)	O	USD	215.83	250.00	32.4%	USD	2.96	4.11	5.61	72.8	52.5	38.4
CLS (Celestica)	O	USD	371.57	520.00	40.1%	USD	6.06	11.35	23.79	61.3	32.7	15.6
CIEN (Ciena)	O	USD	446.45	440.00	178.5%	USD	2.64	7.15	12.46	169.0	62.5	35.8
CSCO (Cisco)	M	USD	117.39	110.00	54.0%	USD	3.73	4.33	5.11	31.5	27.1	23.0
COHR (Coherent)	O	USD	334.56	350.00	178.4%	USD	3.53	5.58	9.51	94.7	60.0	35.2
GLW (Corning)	M	USD	163.25	140.00	75.7%	USD	2.54	3.22	4.16	64.4	50.7	39.2
LITE (Lumentum)	O	USD	1,111.07	1,220.00	590.8%	USD	2.05	8.37	22.40	543.1	132.7	49.6
HPE (HPE)	M	USD	72.09	62.00	173.2%	USD	1.96	3.84	4.59	36.7	18.8	15.7
SPX			7,801.77									

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

HPE estimate is Adjusted EPS; HPE valuation is Adjusted P/E (x);

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

CSCO (MP, $110): We struggle to get excited about Cisco, given the pivot from hardware to software and the deceleration in the software business.

ANET (OP, $250): We see Arista as a quality compounder that should see further upside to numbers amid share gain.

CLS (OP, $520): We believe the impact of the OpenAI Jalapeno program on Celestica's financials is not fully priced into 2027-28 numbers. Accordingly, we see meaningful upside to numbers, and believe the 10/27 analyst day could be a significant catalyst.

CIEN (OP, $440): While Ciena is likely to be supply constrained in the near term, improving supply availability should unlock revenue upside, and longer-term growth expectations and valuation is undemanding.

LITE (OP, $1,220): We believe that the optical growth story should continue, and Lumentum stands to benefit given high incremental margins.

COHR (OP, $350): We believe that Coherent should also benefit from increasing optical adoption, and see a call option from the potential VCSEL revival.

GLW (MP, $140): We struggle to get comfortable with management's growth frameworks, even baking in significant upside from increasing accelerator content.

HPE (MP, $62): HPE's recent networking wins with Oracle and Vultr and higher networking guidance are encouraging, but sustained improvement will depend on overcoming its historical execution issues and below-market growth. Also, HPE's revenue and profit growth have recently been more driven by its traditional server business despite its strong pivot towards becoming a Networking company. See our debrief of HPE's Networking Day (covered by Mark Newman, US IT Hardware).

on 08-Oct-2026

DETAILS

We would like to thank Sierra Valdez for her contribution to this report.

WHAT IS NETWORKING, REALLY, AND WHY DO WE NEED IT?

Networking products exist to connect devices to other devices.

- **For context: a recap on how IT works.** An individual end-user (like yourself) accesses an end-user device like a laptop, desktop, smartphone, or tablet. However, many of the applications that you use are not actually running on your device, but on a server, which refers to a computer that serves groups of users by interacting with the end-user's devices instead of with the end-user directly. For instance, when you send an email to a coworker, your laptop is probably sending the data to a server, which is running Microsoft Outlook Exchange on the back end, which in turn handles your email. For a more thorough entry-level explanation of how enterprise IT works, see the introduction section of: Technology Primer: An Introduction to the Public / Private Cloud, Hybrid Cloud, Multi-Cloud, Containers, and Kubernetes. Networking is an $89B sector within the broader IT universe (Exhibit 1).
- **Networking allows all of these other devices to connect to each other and build larger device ecosystems.** Most hardware devices exist to perform a function on their own, like running a workload or storing data. By contrast, networking devices serve to help other devices communicate.
- **Networking devices solve 4 main issues: forwarding, performance, reliability, and security.**
  - **Forwarding** deals with getting messages where they need to go. Networking devices package data and find the right location for it so that data is transmitted efficiently.
  - **Performance** means this data needs to be handled quickly and reliably. Networking devices handle huge quantities of data, so this data needs to be transmitted at low latency (low delays) and high scale (lots of data).
  - **Reliability** refers to the stability and consistency of transmission.
  - **Security** is about protecting the data and making sure it only goes to the proper users.
- **The core networking industry can be broadly split into optical transport, data center switching, and enterprise. We also cover optical modules and optical fiber.** Exhibit 2 explores each area in detail.
  - **Switching** refers to products that solve the forwarding problem in a straightforward way by matching data to where it is meant to go. Routers operate similarly. However, while switches follow more simplistic rules than routers[1], switches can handle higher volumes at faster speeds and lower prices compared to routers, and accordingly form the backbone for much of networking architecture. This is especially true for AI back-end switching - the high volume of data sent between AI accelerators puts enormous performance requirements on the network, and has been a key growth area in the industry as a result.
  - **Optical transport** refers to systems that are meant to send data over long distances, typically for service provider (think telco) networks, although many data center operators have their own data center interconnect as well. Over long distances, this often becomes a performance and reliability issue, and so in addition to typical optical fiber systems and switching and routing, optical transport often involves both hardware and software designed for error checking and signal amplification.

[1] A networking switch consists of a number of ports connected to a switch ASIC, and works by maintaining a "switching table" which is essentially a simple lookup table consisting of the address of each device connected to one of its ports. When it receives a data packet, it first checks the sender address, and if it is not a match for the switching table, it updates the entry for that port to match the address of the sender. It then checks the recipient address. If there is a match on the switching table, the switch sends the packet to the corresponding port. If there is no match, it sends the packet to all ports except the sender (in hopes that the recipient is connected somehow). This sounds unsophisticated (and in many ways, it is), and a router performs largely the same function but is able to handle more complex algorithms.

on 08-Oct-2026

- **Enterprise networking** is a bit of a miscellaneous category. It can broadly be defined into enterprise local-area networks (LAN), referring to networking equipment used within a single campus (think the office WiFi router in your office building), and enterprise wide-area networks (WAN), which refers to systems designed to connect different campuses. Because enterprises networks need to manage a variety of different users who may have different levels of trust, enterprise networking often involves sophisticated cybersecurity hardware and software.

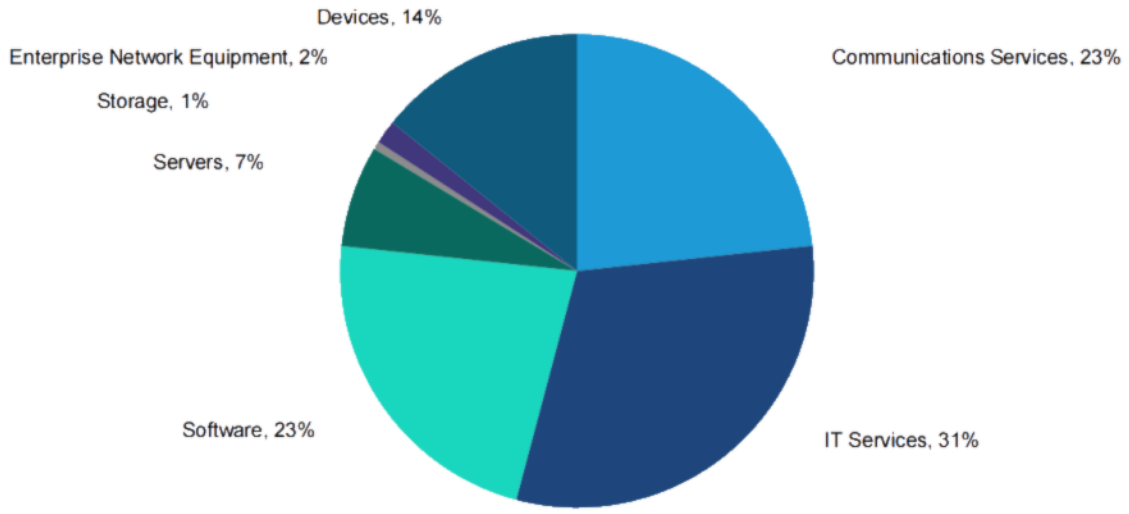
- **Optical fiber** is one of the two physical media that is used in most networking technologies (three, if you count wireless signals), with the other being copper. Optical fiber can transport more data over higher speeds compared to copper, so it is typically used for longer distance communications, although it is increasingly used within the data center as well. For more details on copper connectivity, see the Bernstein multi-industrials team's initiation on Amphenol and TE Connectivity: US Multi Industrials: Can we (inter)connect? Initiating coverage with Outperform on Amphenol and TE Connectivity

- **Transceivers and optical modules** are one of the critical components in the optical ecosystem. A transceiver is a device which sends and receives optical signals (TRANSmits and reCEIVES), or equivalently, converts an electrical signal on copper to a light signal on optical fiber. The optical modules industry refers to companies that produce components for transceivers or other optical systems. In both cases, the highest value component is the laser module, which uses a laser diode that emits a light when powered to covert from electrical to optical signals.

- For more details, can see the Asia Hardware team's *Blackbook*: Artificial Intelligence: Inside the War for AI Data Center Connectivity

This space can seem technical. However, not all tech matters. Networking can seem technical, and filled with acronyms. While some tech is stock-moving, not all technologies matter. While operators talk about the WiFi cycle often, actually looking at Exhibit 3 shows that the officially stated release days of cycles have very little connection with stock moves.

EXHIBIT 1: **IT Spend is a $5.6T dollar industry annually. Networking is $89B of that spend.**



Source: Gartner, Bernstein analysis

on 08-Oct-2026

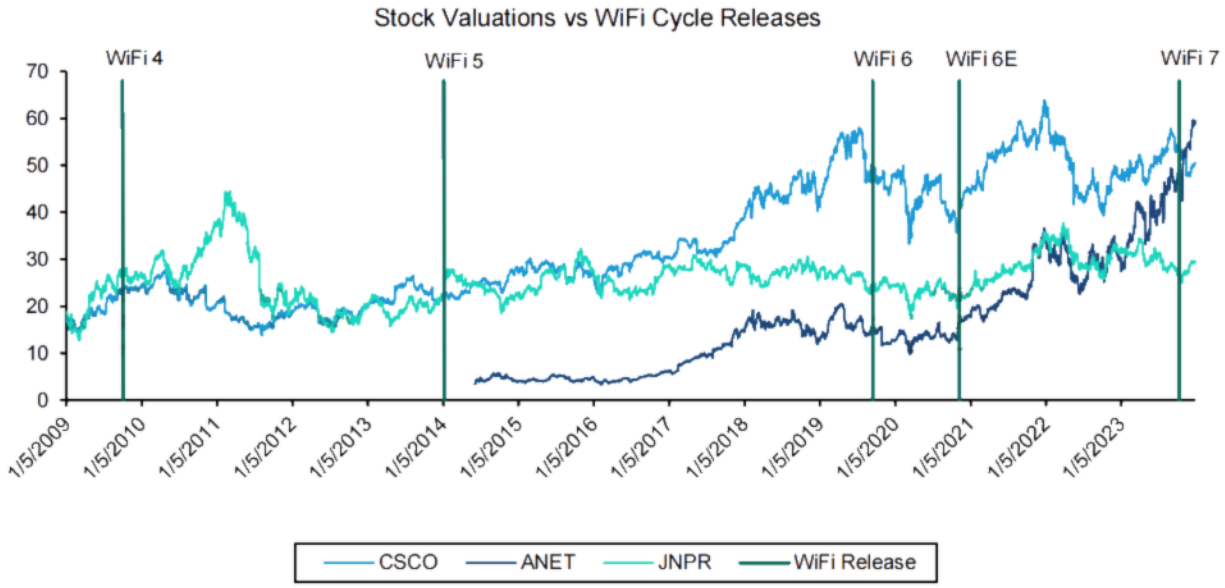
EXHIBIT 2: **Networking and Optical Product Cheat Sheet**

Market	2025 Market Size	What is this product again?
Front End Switching	$18.5B	•Connects external users and systems to internal AI infrastructure •Receives high volumes of data from sender and determines where it is meant to go •Forwards data to the intended recipient to reduce north-south (external-internal) traffic
Back End Switching	$20.1B	•Controls data flow between AI accelerators to manage east-west (internal) data requests •Requires high bandwidth and low latency to quickly connect internal communications •In focus for scale up and scale out
Enterprise LAN	$16.6B	•Local area network that supports an area's IT infrastructure, including access to services like WiFi •Connects many users and devices within a "campus," or a local building •High capacity and low latency is important
Enterprise WAN	$28.5B	•Wide area network that connects an organization's multiple local area networks •Increasingly uses the cloud to connect distant users
Optical Transport	$14.4B	•Allows large volumes of data to be transported in a standardized format •Acts like a digital wrapper by putting everything in uniform "packaging" •Includes built-in error checking and mapping
Optical Fiber	$41.9B	•Replaces copper to support long distance and high bandwidth connections •Acts as the physical medium that carries light signals between AI clusters
Transceiver	$38.9B	•Active device combining a transmitter and a receiver •Interfaces with optical fiber to enable two-way communication •Converts electric data signals into light to be transmitted at extremely high bandwidth

Source: Ciena, Cisco, FS, Dell'Oro, Naddod, Juniper, Gartner, Bernstein analysis

on 08-Oct-2026

EXHIBIT 3: **The industry can seem technical, but not all technologies matter**



Source: Bloomberg, WCA, NYT, IEEE, Bernstein analysis

HOW DO NETWORKING COMPANIES MAKE MONEY?

Networking and optical companies are primarily products companies. This means, for the most part, analyzing company revenues means analyzing the sales of a tangible device. The diversity of products and focuses means each company in our coverage is unique, as highlighted by Exhibit 4.

The three key customer segments that buy networking products are enterprise, service provider, and hyperscalers. Each segment has slightly different architectural needs within the product categories.

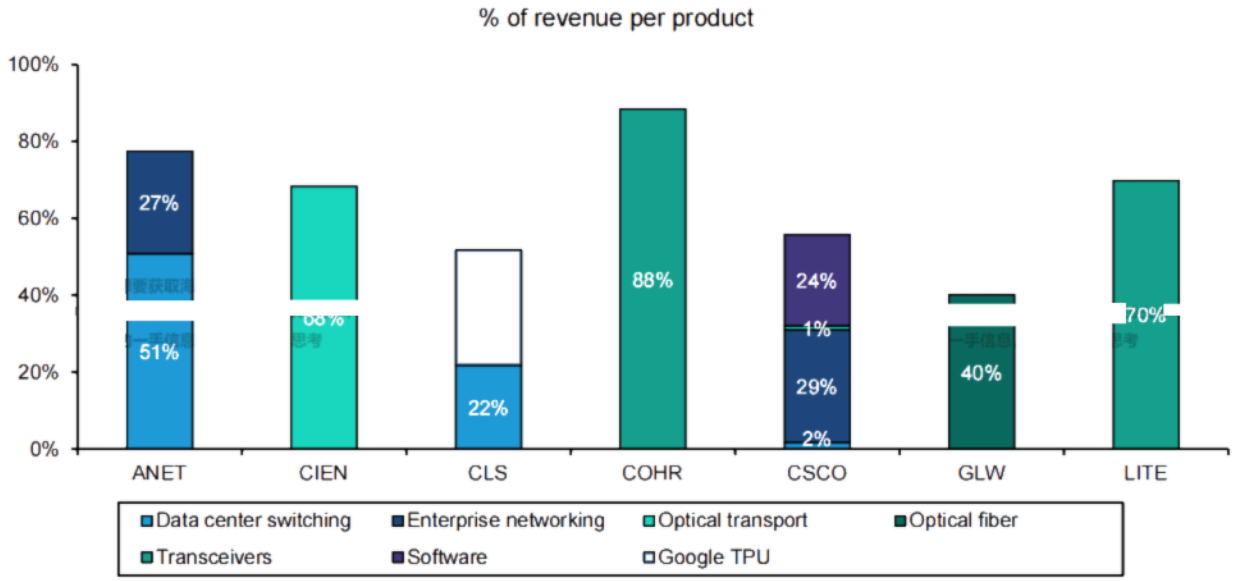
- **Enterprise customers are large companies or organizations that use networking to connect their members.** As explained above, while enterprise is a somewhat miscellaneous category, it basically refers to an organization that needs a closed network for its members to work on. This can look like the WiFi network in an office, or a VPN spanning multiple different campuses.

- **Service providers run large-scale networks that provide connectivity for customers across a wide geographic range.** Examples of service providers are Internet providers and telecom carriers. While service provider used to be more significant, growth has stagnated as the market has become more saturated.

- **Hyperscalers are public cloud service providers with millions (or billions of users).** Hyperscalers maintain data centers with massive storage, compute, and infrastructure capacities. Hyperscalers hold huge quantities of compute capacity, but they are not constantly using all of it. This lets them rent out compute capacity to AI labs for model inferencing and training. The need to meet the ever-growing demands of AI labs means hyperscalers need to improve connectivity within and between data centers, creating a clear room for growth in networking.

on 08-Oct-2026

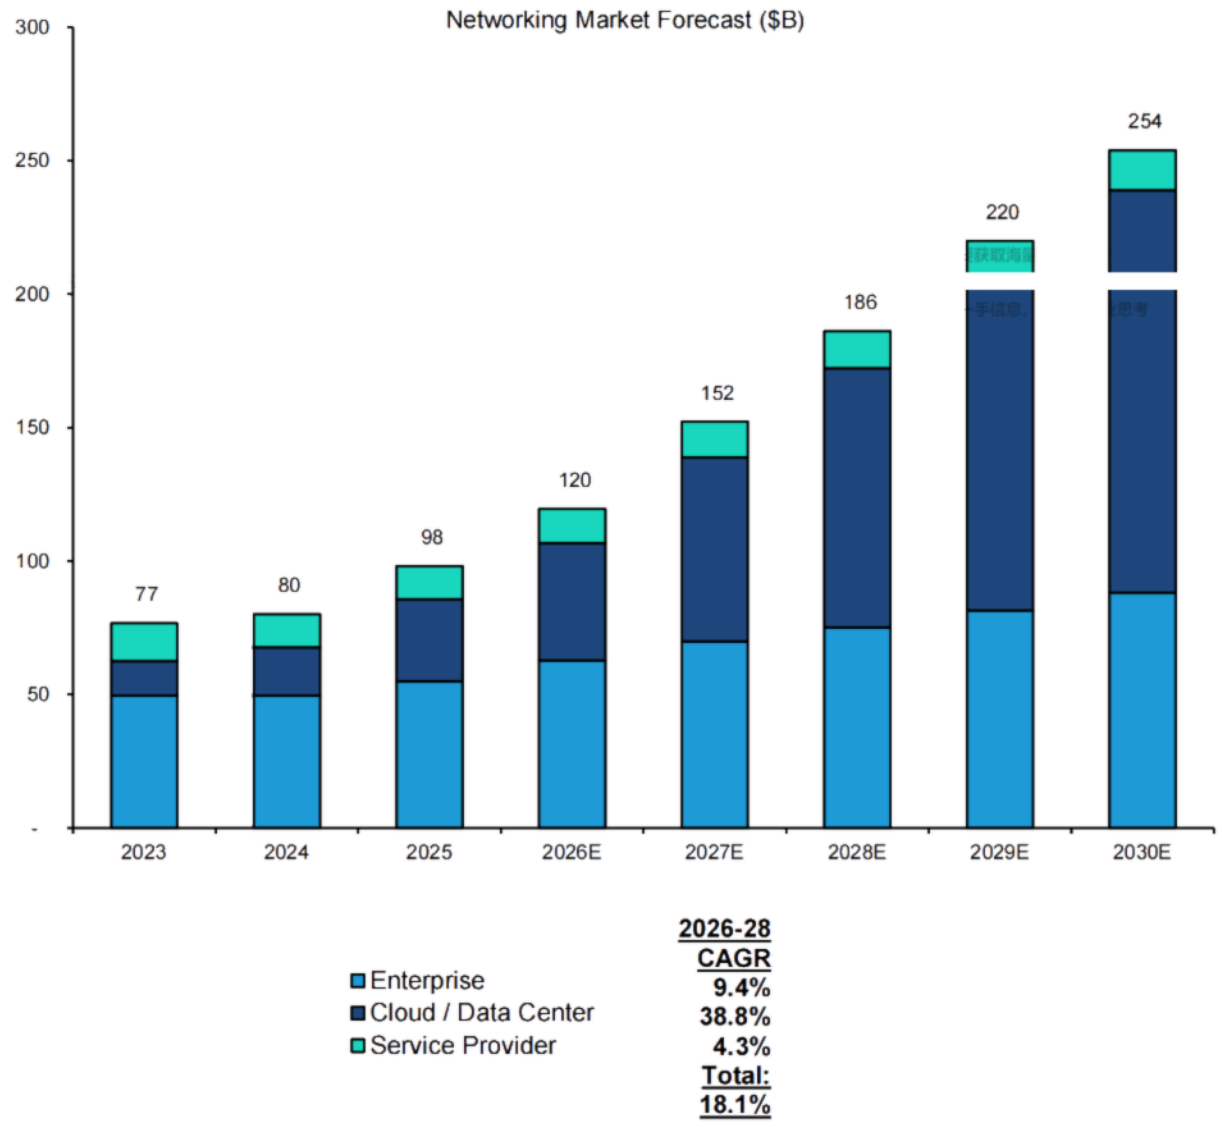
EXHIBIT 4: **Who sells what? Or, companies by market and 2025 revenue mix:**



Source: Company filings, Dell'Oro, Bernstein analysis

on 08-Oct-2026

EXHIBIT 5: **We net out at 18% 2026-28 growth for the industry, including 33% from cloud & AI, 9% from enterprise and 4% from service providers**



Source: Dell'Oro, company filings, Bernstein estimates and analysis

ZERO TO HERO: NETWORKING USED TO BE ONE OF THE WORST PERFORMING TECH SUBSECTORS.

Networking stocks were historically one of the worst performing tech subsectors. Over the past 25 years on average, networking stocks have experienced negative growth (Exhibit 6). This comes from their exposure to the service provider segment and the rise of cloud.

- **The service provider segment has struggled to grow.** Service provider customers create connectivity across wide geographic regions. As teleco and Internet penetration have scaled into almost all regions, very few areas of the world are left without connectivity. This caps the amount of spend for networking to replacement networking equipment, rather than expansion into new regions. Demand for better connectivity is also relatively limited, so service provider exposure has slowed growth.

- **Cloud has generally been a headwind for enterprising networking, with Arista being a notable exception.** As with compute servers and storage, much of the value proposition of working with OEMs like Cisco, Arista, and HPE Juniper,

on 08-Oct-2026

instead of EMS/ODM vendors like Hon Hai and Accton, is that OEMs offer better software and more services support. However, hyperscaler cloud operators generally insource software and services, and accordingly are able to cut OEMs out of the supply chain and work directly with EMS/ODM vendors instead. This in turn means that growing cloud adoption is generally a headwind for OEMs. Arista has proved to be a notable exception, with Arista's EOS software proving to be differentiated even for the hyperscalers, and accordingly Arista has been able to capture much more meaningful share with key cloud customers like Meta and Microsoft than peers. As a result, while cloud adoption has been a headwind for most of the networking OEMs, it has been more neutral to positive for Arista, in turn helping support Arista's track record of share gain.

While these have been challenges to growth, networking stocks have found ways to recover and rerate.

THE HYPE(RSCALER): DATA CENTER NEEDS ARE DRIVING NETWORKING GROWTH

So, is the hype real? We believe it is: the bandwidth wall explains why networking stocks have become structural AI winners. While AI accelerators and chip performance have followed Moore's law, with compute capacity and performance steadily increasing, networking has not scaled at the same rate. Exhibit 7 shows that server racks have increased in FLOPS, bandwidth, and processing speed. Networking devices have to manage these higher performing devices with similarly high speed and processing capacity. Networking is increasingly important for breaking through the bottleneck and delivering higher performance.

Hyperscaler demand has driven a re-rating in networking growth. Networking equipment is critical for the functioning of hyperscaler data centers, and demand for AI is enormous. The shipments of AI accelerators has doubled since 2023, and volumes are projected to continue increasing well into the future, seen in Exhibit 8. Networking is well-positioned to keep up with this demand, as AI front-end and back-end switching estimates project continued growth (Exhibit 9).

- **How does an AI data center network actually work?** Data centers contain many servers. The servers are stacked in racks almost like bookshelves, shown in Exhibit 10. These servers need to communicate with other servers to transfer, compute, and store data. Each rack has a switch on top of it. The switch is connected to cables that connect to servers within the same rack and other racks.

- **There have generally been two ways to expand the data center network: scale up and scale out.** Data center capacity is constrained by the number of switches, since only a limited number of ports fit on each switch. Data centers have to make connections as efficiently as possible to maximize the utility of each port. They scale their connections to do so. Increasing scale for data centers allows for better GPU utilization, redundancy to maximize path efficiency, and improved power management.

  - **Scale up means increasing connections and capacity within the same rack, so servers that are right by each other can communicate.** Scale up is the most low-latency way to increase system capacity, but the amount of connections that can be made within one rack is limited.

  - **Scale out involves connecting more racks to one another.** Scale out helps expand compute without requiring long links between data centers. While scale out effectively increases processing capacity more than scale up, some projects require computing power that exceeds the size of one data center. Individual data centers are constrained by power allocations and real estate, which can make scale across necessary.

    - **Spine and leaf-switching:** One key switch-based network structure, commonly used in AI data centers for scale-out, is the spine and leaf switch. Spine and leaf switching is a multilayer switching network, where every "spine" switch is connected to every "leaf" switch, and then the leaf switches have half their ports connected to spines and the other half connected to either individual servers or (in the case of AI data centers) to scale-up nodes (Exhibit 11).

- **Vendors are increasingly also talking about scale across, which means connecting racks in completely different data centers to each other.** This allows separate data centers to act like one by forming a cluster (Exhibit 12). An example of scale across is Colossus. xAI was unable to get enough power for their project in one state, so they had to divide the Colossus project into two data centers with one on each side of the Tennessee state line.

  - **Scale across is a model training story.** We believe latency matters less for inferencing and for agentic workloads than people think. Inferencing and agentic workloads are part of the process of actually applying a trained model to real requests. While end users like you and I care about latency, the time it takes an AI to respond to a request is much longer

on 08-Oct-2026

than the delay that could be caused by location delays. This means users are trained to tolerate latency. However, for a training run, small delays in communication are amplified over the course of the training process. Actually training the model on data requires AI accelerators like GPUs and TPUs to quickly be able to talk to each other so that the training run takes less time.

- **We believe there is huge potential for upside in scale across.**
  - **Retrofitting existing data centers alone is a massive revenue opportunity.** After analyzing the distances between data centers and the number of potential scale across links, we found that as few as 117 new data center links could result in a nearly $9B dollar market, shown in Exhibit 13 and Exhibit 14. While we do not yet know how many of these links will actually get built, each individual link generates approximately $243M in revenue. This means any new scale across link creates massive financial opportunity. We also think it is likely that many of these links will be built as the sheer size of AI model needs grows.
  - **New data center links could also prove substantial, using Stargate as a case study.** Given Stargate's planned spending of $500B by 2029 on 10 GW of data center capacity, we used the average GW capacity of each data center and the total number of possible connections between sites to size the Stargate targets. We found that new scale across links could be 2.2% of total AI spend on training. Based on AI Capex and training percent, this means new scale across revenue could be ~8B dollars.
  - **Our calculations are supported by industry leaders.** Ciena's CFO, Marc D. Graff, expects scale-across to be "about an $8 to $10 billion dollar market" by 2029. By the end of June, Cisco's quantity of orders from hyperscalers had already exceeded their full-year expectations, at $5.3B by CYQ2 2026 up from the full-year guide of $5B by CYQ4 2026. Nokia doubled their growth guide for Optical and IP Networks, raising to 18-20% up from 10-12%. Since the major hyperscalers have increased their capex, this is unlikely to be the end of the growth story.

- **Continuing with the theme naming, Nvidia has announced the roll-out of scale in, which connects end users to accelerated compute.** Scale-in is often dismissed as a rebranded version of the front-end network, which in turn is not too different from enterprise networks. However, Nvidia is arguing that the emergence of agentic inference will put more demand on the front end, and has branded the new front-end network as scale-in. It remains to be seen how differentiated Nvidia's scale-in is compared to existing front-end solution.

- **Other technologies supporting the ramp up in hyperscaler demand are InfiniBand and Ethernet.** These products both scale compute power. Ethernet is the prevailing network architecture, which is an open-source standard that has dominated networking ever since Xerox PARC researchers published the Ethernet Blue Book in 1980. However, conventional Ethernet based solutions have struggled with resending data that gets lost, and these lost packets are a more acute issue in AI scale-out networks, which have seen spikier traffic. As a result, Nvidia initially developed InfiniBand, which is a combined hardware and software technology that ensures lossless transmission of data through Virtual Lanes. A consortium of industry participants including most incumbent networking players (e.g. Arista, HPE, and Cisco) as well as the hyperscalers (Microsoft, Oracle, and Meta are all steering members) have developed Ultra Ethernet in response, which is an open standard that optimizes ethernet for AI workloads, in turn offering an open, multi-vendor alternative to Nvidia's proprietary Infiniband. For more information, see our colleagues' note U.S. Semiconductors - Infiniwho? Etherwhat? A short primer on AI networking.

- **CPOs, or co-packaged optics, are a new technology to keep on your radar as hyperscaler solutions roll out.** CPOs are advanced transceivers that integrate the optical engine into the switching ASIC. Usually, changing from the optical engine to the switching ASIC causes a delay because the data has to transfer from copper to fiberoptics. This new design allows for lower latency since the components can communicate with each other faster. The CPO also has higher bandwidth. CPOs let transceivers send and receive signals even faster, making switching more efficient.
  - **Since hyperscalers care about speed, the rollout of this product will be important.** However, CPOs have been delayed in rolling out so far, causing speculation as to when they will come onto the market again. CPOs are also expensive; the Asia tech hardware team estimates that CPOs cost 10% more than pluggable solutions, in Future of Tech: Mapping the CPO Value Chain. As demand increases for optical equipment, this remains a technology to watch. For more information, see Co-Packaged Optics (CPO) Primer: Inside the War for Datacenter Connectivity.
  - **Near-packaged optics, or NPOs, have grown as a solution to the CPO delay.** While CPO has been delayed, NPOs have helped fill this demand. While not fully integrating the optical engine with the ASIC, near-packaged optics bring the

on 08-Oct-2026

two very close together, decreasing delays and increasing power efficiency. Though not as efficient as CPO promises to be, NPO is still a step up from other pluggables and are easier to manufacture.

- **AI spend is increasing, and it is flowing to networking.** For each rack of servers, the most expensive components are all networking components, including switches and Ethernet (Exhibit 15, Exhibit 16). Increased AI demand and hyperscaler build outs benefit these stocks greatly.

- **Growth is expected to continue well into 2028.** Our companies have all experienced massive growth in revenue (Exhibit 17), and AI demand is projected to continue bolstering this growth.

THE HYPE IS HERE TO STAY: WE BELIEVE NETWORKING STOCKS ARE CURRENTLY AND WILL CONTINUE TO BE STRUCTURAL AI WINNERS

Networking is continuing to grow, and investors should be sharpening their pencils in this space. The AI opportunity is massive, and networking stocks as a bottleneck are well-positioned to capture a significant portion of revenue from growth. The space shows no signs of slowing down its acceleration in performance either, so we encourage investors to stay tuned as the bonanza carries on. For more information on what stocks we see as particular winners in this space, see our initiation: US Networking: Scaling the bandwidth wall. Initiating with a positive view.

EXHIBIT 6: **U.S. Tech: 2000-2025 Revenue and Market Cap Growth**

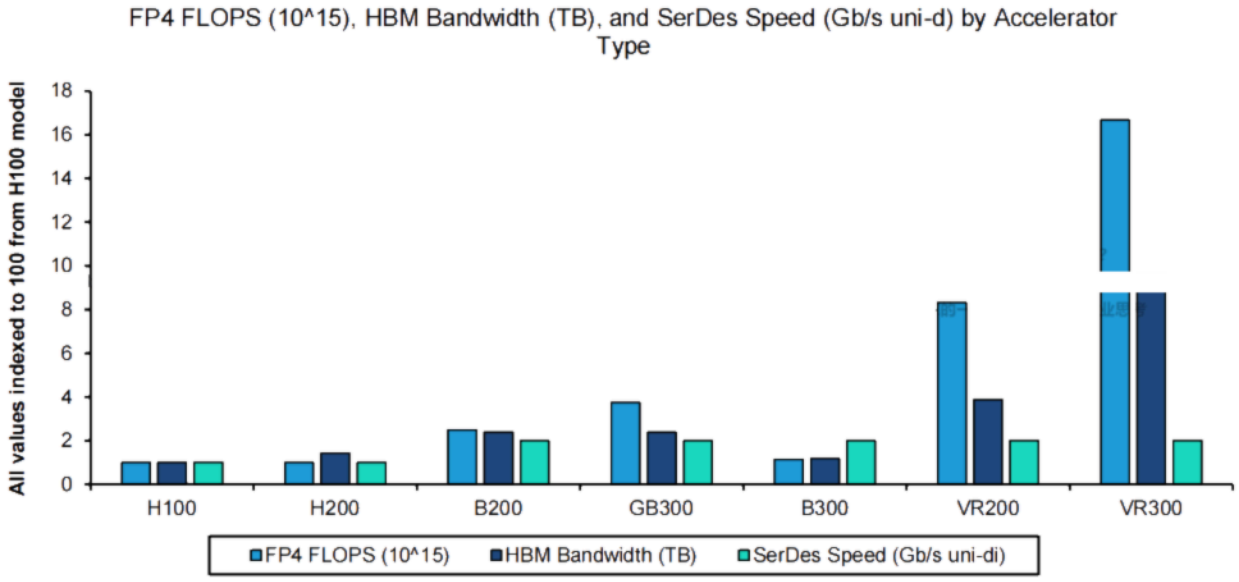
	Hardware ex. Apple		Semis		Software		Tech Services	
	10 years	25 years	10 years	25 years	10 years	25 years	10 years	25 years
Revenue CAGR	1.27%	2%	12.99%	7%	15.26%	10%	5.16%	7.06%
Market Cap ($M)	894,330		7,646,373		7,788,251		875,319	
CAGR	12%	2%	32%	11%	25%	11%	5%	6%

	Networking		Internet			Total	
	10 years	25 years	10 years	25 years		10 years	25 years
Revenue CAGR	-3.75%	-2%	19.97%	16%		8.76%	7%
Market Cap ($M)	472,470		5,893,104			27,043,536	
CAGR	3%	-2%	20%	14%		21%	9%

Source: Bernstein Quantitative Team, Bloomberg (updated as of 09/24/26), Bernstein analysis, cap-weighted

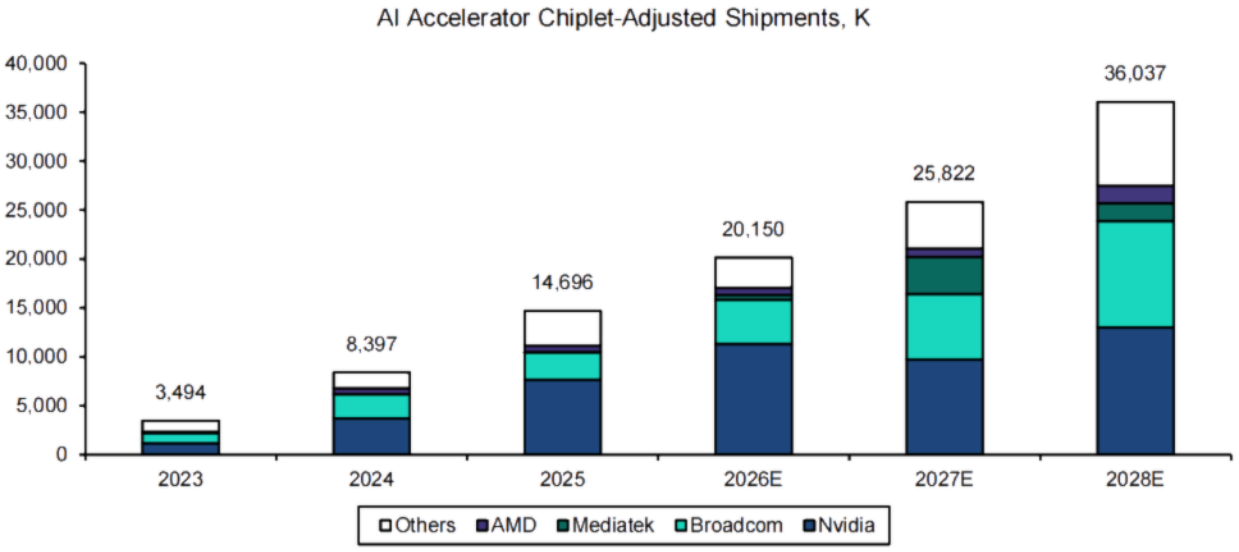
on 08-Oct-2026

EXHIBIT 7: **Accelerated compute has scaled better than networking, creating a "bandwidth wall" in which performance is bandwidth constrained**



Source: Harvard Business School, Silicon Matter Substack, Bernstein estimates and analysis, including from US Semiconductors. Originally published June 8th, 2026.

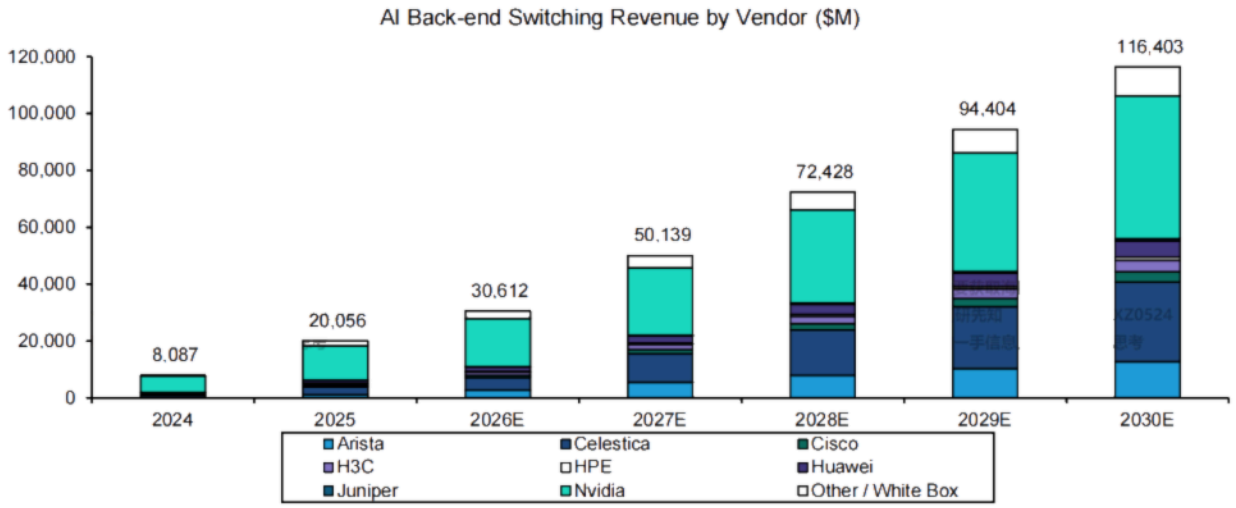
EXHIBIT 8: **We expect accelerator chiplet volumes to double 2025-28E and switch ports to grow proportionally**



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

on 08-Oct-2026

EXHIBIT 9: **AI back-end switching revenue is projected to continue growing**



Source: Dell'Oro, Bernstein estimates and analysis

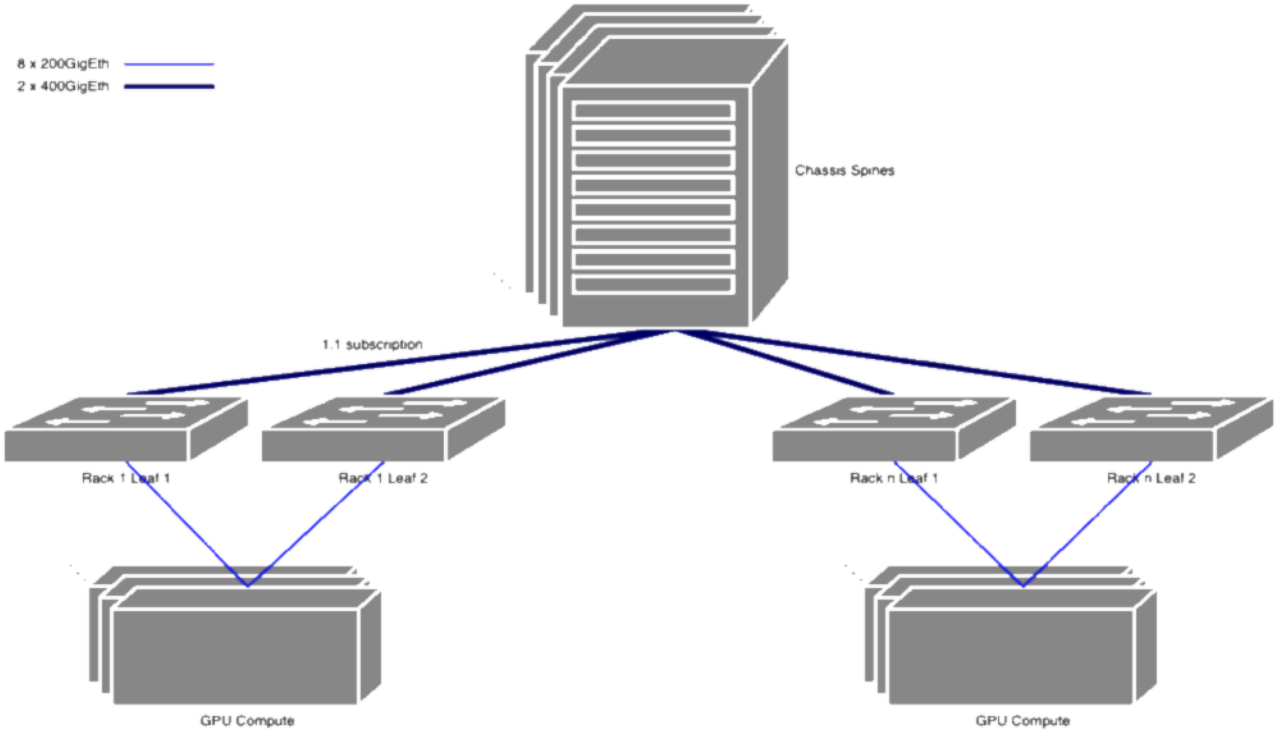
EXHIBIT 10: **Data center server rack**



Source: Alter Playground

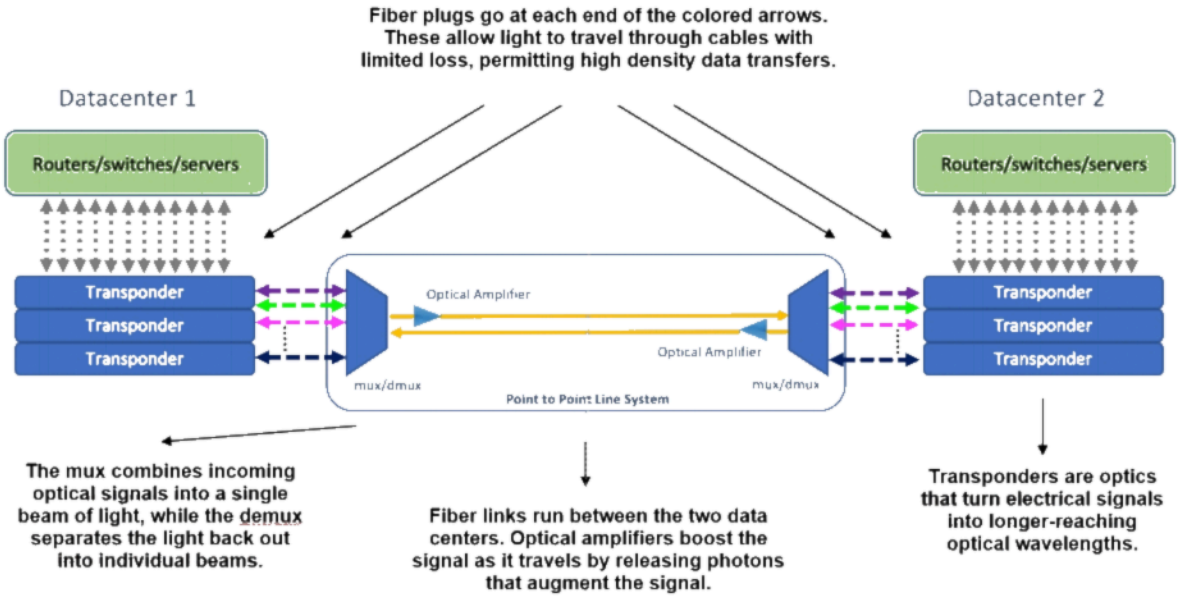
on 08-Oct-2026

EXHIBIT 11: **Spine and Leaf Infrastructure**



Source: IBM Developer, Bernstein analysis

EXHIBIT 12: **Scale-across visualization**



Source: World Wide Technology, Amphenol LTW, Bernstein analysis

on 08-Oct-2026

EXHIBIT 13: **Scale-across calculations based on data center distances and sizes**

Assumptions	
Minimum distance (km)	5
Maximum distance (km)	150
Minimum Size (MW)	100
Depreciation (years)	4
Link Speed (PB/s)	20
ZR+ GB/s	800
Plug Price	3500
Plugs as % of link cost	90%
Gross margin	20%
US + Canada as % of total TAM US is ~60% of data center installed base	80%

Source: Dell'Oro

EXHIBIT 14: **Datacenter Interconnect TAM**

Datacenter Interconnect TAM, $B	
Analysis	
Plugs on each end	25,000
Total Plugs	50,000
Plug value	0.175
Link cost	0.194
Revenue per link	**0.243**
Number of links	117
Total Value	28.4
Annual US + Canada TAM	**7.1**
Global TAM	**8.9**
Sense check	
Data center Capex ($B / GW)	50
Total capex at minimum size	10
Link as % of capex	**2.40%**

Source: Dell'Oro

EXHIBIT 15: **Component Counts by Rack**

Specs	H100 DGX x8	GB200 NVL 72	VR NVL 72
GPU Count	64	72	72
CPU Count	8	36	36
GPU Memory (TB)	5.1	13.4	20.7
CPU Memory (TB)	16	17	54
E1.S NAND Storage (TB)	276	2,248	2,248
ICMS NAND Storage (TB)			1,217

Source: Expert Conversations, Company Disclosures, Omdia, Literature Search, Bernstein estimates and analysis.

Originally published June 8, 2026.

人在草木间

on 08-Oct-2026

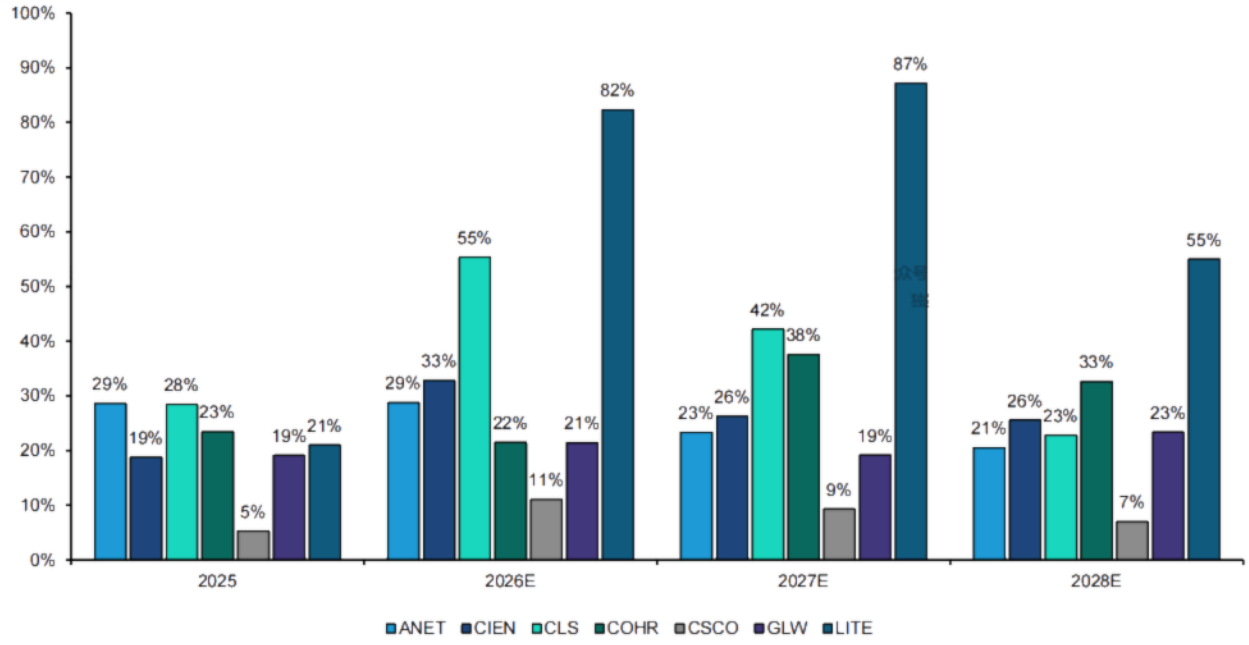
EXHIBIT 16: **Cost of Components by Rack, $K**

	H100 DGX x8	GB200 NVL 72	VR NVL 72
$K / Rack			
CPU ex. DRAM	34	92	180
DRAM	238	252	802
CPU	271	345	982
GPU ex. HBM	1,454	2,112	3,960
HBM	65	185	1,093
GPU	1,519	2,297	5,053
Computing sub-total	1,790	2,642	6,035
NVLink Switch		184	253
Cabling		123	238
Backplanes & Other		102	380
Scale-up	112	287	871
SpectrumX Switch		188	198
Ethernet Fabric		154	198
Scale-out	22	341	396
Networking sub-total	134	628	1,267
Direct-attached Storage	102	832	1,282
			158
Cooling		51	100
Chassis		100	150
Power Delivery		44	100
Others			
Rack Total	2,026	4,296	9,093

Source: Expert conversations, company disclosures, Omdia, Literature Search, Bernstein estimates and analysis

on 08-Oct-2026

EXHIBIT 17: **Revenue Growth by Company, %**



Source: Bloomberg, Bernstein analysis

Glossary

- **Packet:** a unit of data sent across a network
- **AI accelerator:** hardware component that speeds up AI workloads by performing massive computations in parallel
- **Chiplet:** chiplets are mini semi conductor chips that break the functions of a full chip into smaller pieces. This allows many chips to be combined into one system that is more scalable and flexible than a standard chip.
- **Latency:** the time delay between sending data in one place and receiving it in another
- **Bandwidth:** the transmission capacity of a network
- **FLOP:** floating point operations per second, this is a measure of how many calculations a GPU can make
- **Leaf-spine:** two-tier data center design where every server-facing leaf switch connects to every spine switch.
- **Radix:** number of ports on a switch chip or system; higher radix means fewer network tiers.
- **NOS:** network operating system (EOS, IOS XE, NX-OS, open-source SONiC).
- **Merchant silicon:** switch chips bought from a chip vendor rather than designed in-house.
- **White box / blue box:** commodity hardware running customer or third-party software; a blue box is branded hardware sold for open software.
- **Scale-up / scale-out / scale-across:** accelerator links within a rack or pod / across racks in one site / between sites.
- **ARR / RPO:** annualised recurring revenue / contracted revenue not yet recognised.
- **Deferred revenue:** billed or shipped but not yet recognised, often pending customer acceptance.

on 08-Oct-2026

- **PoE:** Power over Ethernet, which powers access points and devices through the switch.
- **SRv6 / MRC:** routing techniques that steer packets across many paths to avoid congestion in AI fabrics.
- **CPO / LPO:** co-packaged optics (optics beside the switch chip) / linear pluggable optics (no signal-processing chip, lower power).