ARCHITECTURE

The fabric that locked the rack is now a part, and its margin is falling

5 November 2026 · Strategic Intelligence Division · Architecture

Astera Labs' Scorpio fabric switches became its largest product family in the quarter to 30 September 2026, at a 72.3% GAAP gross margin, down from 76.2% a year earlier and guided to about 72%. Arista announced an open rack fabric the same week with seven named partners. Astera expects OEM and ODM tray designs from Q1 2027, so the fabric clause is written before the alternative is orderable.

IN SHORT

- The merchant fabric keeps 72.3 cents of gross profit on every dollar of Astera's revenue and 63.6 cents on Arista's, and Astera's own margin has fallen from 76.2% in a year.
- AMD's 10-Q for the quarter ended 26 September 2026 states that it does "not manufacture or sell the completed Helios rack-scale solutions", so one accelerator vendor has already surrendered the rack layer.
- Astera expects Orion rack designs through its OEM and ODM partners beginning in Q1 2027, which ends on 31 March 2027, so a commitment signed before that date fixes a fabric whose replacement is a re-qualification.
- NVIDIA's own filing states that the tenant will deploy "NVIDIA's full-stack DSX AI factory platform, subject to limited exceptions", and the exceptions are not enumerated anywhere in it.
- Astera sells UALink and NVLink Fusion out of one platform and Marvell runs a UALink switch team while partnering NVIDIA, so two open camps buy the buyer optionality rather than a price war.

The scale-up fabric stopped being the lock on the rack and became a part whose margin is thinning, and the fall is already in the filing. Astera Labs' release of 3 November 2026 reported a GAAP gross margin of 72.3% for the quarter to 30 September 2026, against 76.2% a year earlier, and put Orion rack reference and system designs on track for availability through OEM and ODM partners beginning in Q1 2027. Arista published the other side of the same trade the same day: an open rack fabric built with seven named partners. A capacity commitment signed before 31 March 2027 fixes a fabric whose replacement is a re-qualification of the pod, which makes the fabric clause a contract decision with a date rather than a preference to hold.

Who sells which layer of an AI rack now, and the sentence in the filing that says so

Layer of the rack	Sold now as a merchant part by	The filing's own words
Scale-up fabric	Astera Labs, Scorpio X-Series; Arista, Etherlink SU-144; Marvell, XConn	"Scorpio fabric switches became the company's largest product family"
Signal conditioning	Astera Labs, Aries 7	"allowing customers to select different signal conditioners without redesigning the board"
Rack assembly	AMD's customers and their manufacturing partners	"we do not manufacture or sell the completed Helios rack-scale solutions"
Completed rack and stack	NVIDIA, with its own accelerators	"full-stack DSX AI factory platform, subject to limited exceptions"

Source: Astera Labs Form 8-K, Exhibit 99.1, 3 November 2026; Arista Networks Form 8-K, Exhibit 99.1, 3 November 2026; Marvell Technology Form 10-Q, 28 August 2026; Advanced Micro Devices Form 10-Q, quarter ended 26 September 2026; NVIDIA Form 8-K, 17 August 2026.

AMD has filed that it does not sell the completed rack, and the layer it hands over is the one that decides which accelerators can share a pod. Its 10-Q for the quarter ended 26 September 2026 puts Helios among the system-level architectures "which we develop and license for implementation by customers and their manufacturing partners", and then states that Helios racks "incorporate components that we do not design, manufacture or supply, including networking". That networking layer is the scale-up fabric. Two merchants announced a mixable version of it on 3 November 2026, and the third, Marvell, was already in the layer, having completed its XConn acquisition on 10 February 2026. This reaches a reader who commits AI capacity years before delivery, or who single-sources a fabric whose replacement is a re-qualification of the whole pod.

72.3%
Astera GAAP gross margin, quarter to 30 September 2026

76.2%
The same margin a year earlier

144
Accelerators per rack in a single hop, Etherlink SU-144

Q1 2027
Astera's stated availability for Orion rack designs

Source: Astera Labs Form 8-K, Exhibit 99.1, 3 November 2026; Arista Networks Form 8-K, Exhibit 99.1, 3 November 2026.

The lock was the fabric, and from Q1 2027 it is a part number

Astera's release of 3 November 2026 is the document to read closely, because its own description of Orion is a specification for a rack whose compute is not its own: "a compute-agnostic rack-scale reference architecture that allows customers to choose their own GPU and CPU", built on "its high-radix 320-lane PCIe switch fabric is designed to scale from 8 to 160 accelerators per pod". At the Open Compute Project summit it showed "a powered Orion rack featuring interchangeable, pre-validated compute trays with different CPU and accelerator configurations", and Aries 7 carries the same idea down to the signal conditioner, "allowing customers to select different signal conditioners without redesigning the board". The sentence with a date in it is this one: "on track for availability through Astera Labs' OEM and ODM partners beginning in Q1 2027".

Arista announced the other half on the same day, and its wording is careful. The company "launched open standards-based rack-scale architectures for scale-up and scale-out AI networks" "In partnership with AMD, Arm, Broadcom, d-Matrix, Meta, Microsoft, and Qualcomm", built on "the Etherlink SU-144 architecture, supporting up to 144 accelerators per rack in a single hop and scaling up to 1,024 XPUs via cross-rack topologies". Seven firms are named as partners in a launch. The release names no price, no order and no customer.

What the incumbent's own filing asks the tenant to accept

"OpenAI, as the tenant, will utilize capacity at the Portsmouth Site supporting approximately 4.25 gigawatts of IT load to deploy NVIDIA's full-stack DSX AI factory platform, subject to limited exceptions." Source: NVIDIA Form 8-K, Item 1.01, 17 August 2026. The filing does not enumerate the exceptions.

The toll is 72.3 cents on the dollar, and it has already come down from 76.2

The merchant fabric is not cheap. It is thinning. In the release of 3 November 2026 Astera reported a GAAP gross margin of 72.3% in the quarter to 30 September 2026, against 76.2% in the same quarter of 2025 and 73.3% in the quarter to 30 June 2026, with approximately 72% guided for the quarter to 31 December 2026. That is 397 basis points gone in a year on the layer that once justified selling a rack rather than a part. The layer now has three merchants in it: Marvell, which completed the XConn acquisition on 10 February 2026, and Astera and Arista, which both filed on 3 November 2026. Arista, the other half of the same trade, earned a

GAAP gross margin of 63.6% on $3.66 billion of quarterly revenue, which grew 58.5% year over year.

Put it as a unit rather than a total: 72.3 cents of gross profit on every dollar of Astera's revenue, 63.6 cents on every dollar of Arista's, and the reader supplies the fabric budget. What the buyer should not assume is that the open pod is the cheaper one. Neither company has published a price for an Orion tray or an SU-144 fabric, so no comparison exists to make.

The merchant has taken the build-ahead risk with it. Astera held $200.9 million of inventory at 30 September 2026 against $59.0 million at 31 December 2025, which is 113 days of its $162.1 million quarterly cost of revenue. That working capital is Astera's to carry for as long as the buyer has not ordered, and it is the opposite of what a single-vendor rack did, where the trays arrived on the operator's purchase order years before the pod lit.

Both camps are staffed by the same firms, so the buyer gets optionality and not a price war

Marvell is the clearest case and it is dated. Marvell's Form 10-Q for the quarter ended 1 August 2026, filed 28 August 2026, records a total purchase consideration of $3.5 billion for Celestial AI on 2 February 2026, "a provider of a Photonic Fabric technology platform purpose-built for next-generation scale-up interconnect", and eight days later $469.0 million for XConn, "a provider of advanced PCIe and CXL switching silicon", an acquisition the same note says "expands the Company's switching portfolio" and augments a UALink scale-up switch team. In the first quarter of fiscal 2027 the same filing records that Marvell and NVIDIA "announced a strategic partnership to connect our custom XPUs and compatible scale-up networking with NVIDIA's AI infrastructure ecosystem". A firm building the open switch team took a partnership with the closed one, in the same quarter, and filed both.

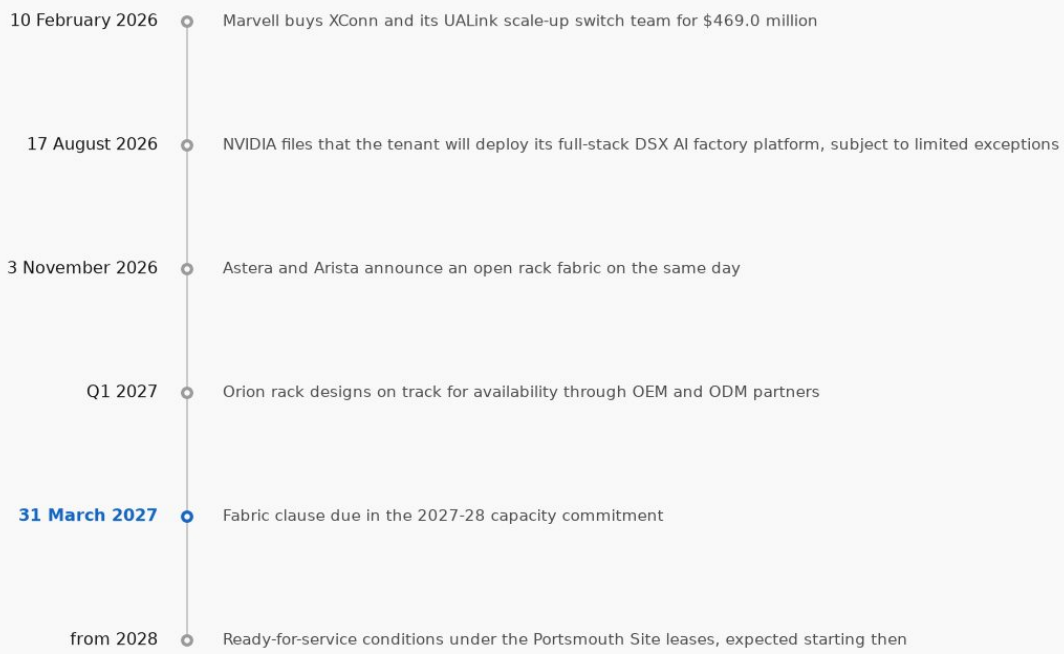
Astera sells both camps out of one platform that "integrates CXL, Ethernet, NVLink Fusion, PCIe, and UALink semiconductor-based technologies", a merchant part that speaks the incumbent's own protocol. Credo's 10-Q of 2 September 2026 puts its connectivity products in the same room, "optimized for optical and electrical Ethernet, PCIe and emerging UALink, ESUN and SUE applications". Two camps and two protocols do not produce a price war when the same firms sell into both.

The open camp carries a single-source risk of its own, in Arista's words rather than anyone else's. That release's cautionary language names "our reliance upon a predominant merchant silicon vendor", and Broadcom is one of the seven firms named in the launch.

Two open camps, and what a buyer gives up in each

Camp	What is on offer	Named in it	What the buyer gives up
Ethernet, ESUN	One hop to 144 accelerators per rack, and 1,024 XPUs across racks	Arista, with AMD, Arm, Broadcom, d-Matrix, Meta, Microsoft and Qualcomm	A specification whose launch partners include Meta and Microsoft, and which names no customer
PCIe, UALink	A 320-lane switch fabric, 8 to 160 accelerators per pod	Astera, with UALink and NVLink Fusion in one platform; Marvell's UALink switch team	A merchant part whose gross margin has already fallen 397 basis points in a year

Source: Arista Networks Form 8-K, Exhibit 99.1, 3 November 2026; Astera Labs Form 8-K, Exhibit 99.1, 3 November 2026; Marvell Technology Form 10-Q, 28 August 2026.



The window, from Marvell's switch-team acquisition to the quarter Astera says the neutral tray becomes orderable, and the date a commitment signed now outlives. Sources: Astera Labs and Arista Networks Form 8-K, Exhibit 99.1, 3 November 2026; Marvell Technology Form 10-Q, 28 August 2026; NVIDIA Form 8-K, 17 August 2026.

The clause is due by 31 March 2027, and one date serves four seats

The window closes on the last day of the quarter in which the alternative becomes orderable. Astera's sentence is "beginning in Q1 2027", Q1 2027 ends on 31 March 2027, and a capacity commitment signed before that date fixes a fabric whose replacement is a re-qualification of the pod. The re-population right and the tray swap price have to be written into the contract, because neither seller publishes either, and the buyer who waits to see the price list is buying a specification already written.

Four seats, one date, and the artefact each one owes

Seat	The artefact	By when	Why this one
CIO or chief AI officer	A fabric clause in the 2027-28 capacity commitment that names the scale-up fabric, names the second source, and buys a re-population right at refresh	31 March 2027	Astera's own availability date is the last day of that quarter, so a commitment signed earlier fixes the fabric before the alternative can be ordered
Network architect	The conformance document for whichever open fabric is shortlisted, ESUN or UALink, in place of the launch announcement	31 March 2027	Arista names Meta and Microsoft as launch partners and names no customer, so joining the standard is the only lever an outside operator holds

Seat	The artefact	By when	Why this one
General counsel	The written list behind "subject to limited exceptions", obtained before any commitment on a full-stack platform is signed	31 March 2027	NVIDIA's Form 8-K states the deployment subject to that phrase and does not enumerate it, so the exception is only as wide as what is written down
CFO	The fabric carried as its own line in the capacity model, at a merchant gross margin between 63.6% and 72.3%, rather than inside a rack price	31 March 2027	That is where the record puts the gross profit, and a bundled rack price hides whether the buyer is paying Astera's 72.3% or Arista's 63.6%

Source: Dates from Astera Labs Form 8-K, Exhibit 99.1, 3 November 2026, which states availability "beginning in Q1 2027"; margins from the same release and from Arista Networks Form 8-K, Exhibit 99.1, 3 November 2026.

The board decides by 31 March 2027, and six questions test the decision

The decision that reaches the board before 31 March 2027 is one line: approve the 2027-28 capacity commitment only with the scale-up fabric named, a second source named and a re-population right bought at refresh. After that date the alternative is orderable and the clause is no longer there to negotiate, so the board either writes the condition now or accepts the fabric that arrives with the pod.

- Which scale-up fabric is named in each capacity commitment we have signed, and which seat chose it?
- If we had to replace the accelerator trays in a pod at refresh, what would have to be re-qualified, and who pays for it?
- What does our rack price include, and what share of it is the fabric?
- Which open fabric are our architects inside, and who represents us where it is written?
- What is the written list behind the exceptions clause in the platform we have committed to?
- If the neutral tray is orderable after 31 March 2027, what in the 2027-28 plan assumes we do not switch, and who signed that assumption?

What to watch, and the two dates that settle it

31 March 2027 ends the quarter in which Astera says Orion rack designs reach the market through its OEM and ODM partners. If no OEM or ODM has published a price for an Orion rack or an Etherlink SU-144 fabric by that date, the neutral pod was not orderable in the quarter named, and the specification window described here had not opened.

Then the margin, which is the falsifiable half. Astera reported 76.2%, 73.3% and 72.3% for the quarters to 30 September 2025, 30 June 2026 and 30 September 2026, and guided the quarter to 31 December 2026 to approximately 72%. A GAAP gross margin at or above 76.2% in any quarter it reports for a period ending in 2027 says the fabric is a toll that holds. Below 65% says the toll is eroding faster than this reading claims. Both arrive in a filing, on a date, and neither needs a second opinion.

The six filings this piece rests on

- Astera Labs, Form 8-K filed 3 November 2026, accession 0001736297-26-000044, Exhibit 99.1, third quarter fiscal 2026 results.

- Arista Networks, Form 8-K filed 3 November 2026, accession 0001596532-26-000252, Exhibit 99.1, third quarter 2026 results.
- Advanced Micro Devices, Form 10-Q for the quarter ended 26 September 2026, signed 3 November 2026, accession 0000002488-26-000191, risk factors.
- Marvell Technology, Form 10-Q for the quarter ended 1 August 2026, filed 28 August 2026, accession 0001835632-26-000025, Note 4, Business Combinations.
- Credo Technology Group, Form 10-Q for the quarter ended 1 August 2026, filed 2 September 2026, accession 0001628280-26-060111.
- NVIDIA, Form 8-K filed 17 August 2026, accession 0001045810-26-000069, Items 1.01, 2.03 and 7.01.