



Ultra Ethernet (UEC) vs RoCEv2: What Changes for AI Fabrics

Plan and build private AI compute with GPU Smith. We connect workload requirements with GPU hardware, networking, deployment and operating decisions so your infrastructure fits the work it must perform.

Published by GPU Smith · 2026-10-29 · ultra ethernet vs rocev2 · ultra ethernet consortium 1.0 specification · ultra ethernet transport uet · rocev2 vs ultra ethernet packet spraying · rocev2 pfc ecn tuning ai cluster · techniques for tuning rocev2 in 1.6t environments

Original article: <https://gpusmith.com/articles/en/ultra-ethernet-uec-1-0-vs-rocev2-ai-fabrics>



In this article

1. The short answer
 2. Why RoCEv2 strains at AI scale
 3. What UEC is, and which version is current
 4. Layered architecture
 5. Delivery modes: RUD, ROD, RUDI, UUD
 6. Packet spraying and multipathing
 7. Congestion control: NSCC and RCCC
 8. Link Level Retry and Credit-Based Flow Control
 9. Side-by-side: RoCEv2 vs Ultra Ethernet
 10. Who claims UEC support (as of 28 October 2026)
 11. Procurement questions that separate claims from capability
 12. Decision guide
 13. Sources
-

The short answer

RoCEv2 carries InfiniBand's transport over UDP/IP and assumes a lossless, in-order network, which in practice means Priority Flow Control (PFC) on every hop. Ultra Ethernet (UE) keeps standard Ethernet physical layers but replaces the transport with the Ultra Ethernet Transport (UET). UET is designed to run on best-effort (lossy) networks. It sprays packets across all paths, accepts out-of-order arrival, retransmits selectively, and runs congestion control at the endpoints. PFC becomes optional rather than load-bearing (Hoefler et al., [arXiv 2508.08906](#), 2025).

As of 28 October 2026, the specification is no longer at 1.0. The Ultra Ethernet Consortium (UEC) lists 1.0.1, 1.0.2 and 1.0.3 as follow-ups, and **1.1 (12 October 2026) is the current published version (UEC [specification history](#))**. One switch ASIC claims UEC compliance and is shipping (Broadcom Tomahawk Ultra). One 800G NIC claims full UEC feature compliance but was announced as sampling (Broadcom Thor Ultra). AMD sells a "UEC-ready" 400G NIC. We found no independent certification program. For most 2026 builds the defensible plan is still a well-tuned RoCEv2 fabric whose NICs and switches have a documented path to UET.

For a product-level comparison of fabrics, see [InfiniBand vs Spectrum-X vs RoCE](#).

Why RoCEv2 strains at AI scale

The UE specification's authors summarize the problem this way: RoCE adopted InfiniBand's transport "largely unchanged". That transport requires lossless delivery with strict packet ordering, and on Ethernet the guarantee comes mainly from PFC. PFC in turn needs a separate traffic class with substantial headroom buffer and suffers from congestion spreading and head-of-line blocking. In-order delivery also limits path choice ([arXiv 2508.08906](#), §1).

Microsoft's account of running RoCEv2 in production shows these mechanics in practice ([Guo et al., SIGCOMM 2016](#)):

- **Path pinning.** The RoCEv2 destination UDP port is always 4791. The source port is chosen per queue pair (QP), and switches hash the five-tuple. All traffic on one QP therefore follows one path.
- **Headroom.** When an ingress queue crosses XOFF, the switch sends a pause frame upstream, and packets keep arriving until the pause takes effect. Each lossless priority needs reserved headroom sized by the MTU, the PFC reaction time and "most importantly, the propagation delay". With 9 MB and 12 MB shallow-buffer switches and links up to 300 m, Microsoft could afford only two lossless classes out of eight.
- **Safety failures.** The team hit a PFC-induced deadlock ("yes, it happened!"), an RDMA transport livelock and NIC PFC pause-frame storms. They built PFC storm watchdogs on both NICs and switches to contain them.

End-to-end congestion control was layered on top to reduce how often PFC fires. DCQCN, the ECN-based scheme Microsoft and Mellanox built for RoCEv2, was motivated by PFC's head-of-line blocking and unfairness ([Zhu et al., SIGCOMM 2015](#)).

Worked example 1: PFC headroom grows with link speed

This is an illustration, not a vendor formula. Assumptions: 100 m of fiber at about 5 ns/m, an assumed 1 μ s combined pause generation and reaction time, 9,216-byte jumbo frames, one lossless class, and 64 ports.

At 800 Gb/s:

1. One-way propagation: $100 \text{ m} \times 5 \text{ ns/m} = 500 \text{ ns}$. The pause has to travel upstream while data keeps flowing downstream, so use the round trip: 1 μ s.
2. Bytes in flight over 1 μ s: $800 \times 10^9 \text{ b/s} \times 1 \times 10^{-6} \text{ s} = 800,000 \text{ bits} = 100,000 \text{ bytes}$.
3. Reaction time (assumed 1 μ s): another 100,000 bytes.
4. MTU allowance (one frame in transmission at each end): $2 \times 9,216 = 18,432 \text{ bytes}$.
5. Per port: $100,000 + 100,000 + 18,432 = 218,432 \text{ bytes}$. Across 64 ports: about 14.0 MB for a single lossless class.

At 1.6 Tb/s, the two time-based terms double: $200,000 + 200,000 + 18,432 = 418,432 \text{ bytes per port}$, or about 26.8 MB across 64 ports.

Doubling the link rate roughly doubles the buffer that PFC has to reserve, and a second lossless class doubles it again. Distance multiplies everything. This is the practical core of "tuning RoCEv2 at 1.6T": keep lossless classes to a minimum, keep cable runs short where PFC is enabled, confirm the switch vendor's own headroom calculation for your exact optics and lengths, and lean on ECN-based congestion control so PFC stays a backstop. Our [RoCE acceptance testing](#) runbook covers how to verify that on delivered hardware.

Worked example 2: why per-flow hashing leaves links idle

Assume 8 large flows leaving a leaf over 8 equal uplinks, with each flow hashed to a uniformly random uplink. Real ECMP hashing is not perfectly random, so treat this as an idealization.

- Probability that no two flows collide: $8! / 8^8 = 40,320 / 16,777,216 \approx 0.24\%$.
- Expected number of idle uplinks: $8 \times (7/8)^8 \approx 8 \times 0.344 \approx 2.75$.

In the typical case almost three of the eight uplinks carry nothing while others carry two or more flows at half rate each. Synchronized collectives finish at the pace of the slowest flow. Spreading each flow over several QPs helps, because different QPs can take different paths ([Guo et al.](#)). Per-packet spraying removes the problem at the source, because every flow's load is divided across all eight links. For how collective traffic patterns load the fabric, see [NCCL algorithm selection](#).

What UEC is, and which version is current

UEC announced Specification 1.0 on 11 June 2025. It described the spec as covering "all layers of the networking stack—including NICs, switches, optics, and cables" and said it "scales to millions of endpoints" ([Linux Foundation, 2025-06-11](#)). The 1.0 specification ran to 562 pages; 1.1 runs to 764. Its authors (Broadcom, HPE, OpenAI, Microsoft, AMD and others) published a readable overview that says it expects "several rounds of errata" ([arXiv 2508.08906](#)). The errata arrived as follows:

Version	Date	UEC description
1.0	2025-06-11	Initial public release
1.0.1	2025-09-05	Editorial clarifications and correction to the RCCC source algorithm
1.0.2	2026-01-28 (2026-01-21 in the 1.1 release notes)	Correction to CMS congestion control algorithms and other editorial clarifications
1.0.3	2026-07-16	Editorial fixes, 200 Gb/s-per-lane signaling, a new Negotiation CP type, CBFC/LLR race fixes
1.1	2026-10-12	Major release with new functionality; current published version

Source: [UEC specification history](#) and the [1.1 release notes](#), both checked 2026-10-28.

Two of the 1.0.x point releases changed congestion-control algorithms, and 1.0.3 fixed link-layer races, so the changes go beyond wording. A NIC that claims "UEC 1.0" support should therefore state which point release its firmware implements.

Per UEC's 1.1 release notes, 1.1 adds:

- a Unified Local profile for scale-up networks, including a reliable ordered delivery mode for local networks (RODL)
- a Unified Forwarding Header (UFH) and Ethernet idle reduction
- In-Network Collectives (INC)
- Congestion Signaling (CSIG), with CMS and LLDP enhancements to use it
- Programmable Congestion Management (PCM)
- a UE verbs interface mapping
- a management chapter (Redfish and YANG models)

UEC says 1.1 stays backward compatible; only the Unified Local profile adds mandatory requirements. Compliance is self-attestation: UEC publishes profile matrices and checklists "for the purpose of Ultra Ethernet [UE] Compliance self-attestation" ([UEC compliance](#)). [Network World \(2026-10-14\)](#) quotes UEC's technical

chair as expecting some features to reach existing hardware through firmware.

The protocol description below follows the 1.0 overview, the base that 1.0.x and 1.1 build on. We spot-checked ports, profiles, delivery modes, SACK and TSS details against the [1.1 PDF](#).

Layered architecture

Layer	RoCEv2	Ultra Ethernet (1.0)
Physical	Standard Ethernet	Standard Ethernet, essentially unchanged. First products at 100G or 200G per lane (§3)
Link	PFC (and ETS) for lossless classes	Optional Link Level Retry (LLR) and Credit-Based Flow Control (CBFC), negotiated via LLDP (§3.5)
Network	IP	IP
Transport	InfiniBand transport in UDP, port 4791	UET in UDP port 4793 (IANA), or directly over IP with a 4-byte entropy header instead of the 8-byte UDP header (§2.1, §3.2.2)
Software API	Verbs	libfabric; 1.1 adds a verbs mapping (§2.3 of the 1.1 spec)

UET has four sublayers (§3):

- **SES (Semantics):** message semantics. Connects to libfabric and commits RMA operations at the target.
- **PDS (Packet Delivery):** reliability and delivery modes.
- **CMS (Congestion Management):** byte-level congestion windows.
- **TSS (Transport Security):** optional encryption and authentication, using AES-GCM with a 256-bit key by default.

1.0 defined three profiles that set how much semantics a device implements (§2.2); 1.1 adds a fourth, Unified Local, for scale-up links (§1.1.2 of the 1.1 spec):

- **AI Base:** the simplest. Tag matching is left to libfabric or the collective library.
- **AI Full:** adds exact tag matching and deferrable sends, and mandates RMA read.
- **HPC:** adds wildcard matching for MPI and OpenSHMEM.

Pin down profile support first in any NIC datasheet.

The change that matters most for scale is how connection state works. RoCEv2 reliability is tied to QPs that are set up before traffic flows. UE explicitly works "without queue-pairs acting as connections". Instead it uses ephemeral Packet Delivery Contexts (PDCs), which are created from state carried in the first data packets "without incurring additional latency at first packet arrival" (§3.1.1, §3.2). This is how UE claims to support millions of endpoints without per-peer QP state on every NIC.

Delivery modes: RUD, ROD, RUDI, UUD

UET's PDS offers four modes ([§3.2.1](#)):

Mode	Meaning	Use
RUD	Reliable Unordered Delivery	Default bulk mode. The most efficient reliable mode because it permits packet spraying
ROD	Reliable Ordered Delivery	Strict ordering, for example HPC wildcard matching. Uses a single path per flowlet, so the paper calls it "much less efficient"
RUDI	Reliable Unordered Delivery for Idempotent operations	No receiver state needed. HPC profile only
UUD	Unreliable Unordered Delivery	Datagram-style traffic, software protocols, management

1.1 adds RODL, an ordered mode for the Unified Local profile.

RUD lets the NIC accept packets in any order and place payloads directly into memory, because each packet carries enough addressing for the receiver to do so. Classic RoCEv2 reliable connections expect in-order arrival, so reordering looks like loss to them.

Loss recovery changes as well. UE acknowledgements carry a cumulative ACK and a 64-bit SACK bitmap, so the sender retransmits only what is missing. The paper treats timeouts as a poor loss detector and adds optional faster signals: packet trimming, out-of-order counting and entropy-based schemes ([§3.2.4](#), [§3.2.5](#)).

Packet spraying and multipathing

Standard ECMP hashes a flow onto one path. UE redefines a header field as an Entropy Value (EV), which is the UDP source port when UDP/IP is used. The sender changes the EV packet by packet, so ordinary ECMP switches spread one message over many paths. Hash collisions still happen, but they are short-lived and switch buffers absorb them. The spec also describes path-aware alternatives such as REPS ([§2.1](#), [§3.3.5](#)).

Because spraying reuses ECMP, switches need no new forwarding model, only path diversity (uplinks, planes) and sensible ECN marking. If oversubscription is already part of your design, revisit [fabric sizing and oversubscription](#). Spraying makes a non-blocking Clos usable, but it does not create bandwidth that oversubscription removed.

Congestion control: NSCC and RCCC

UE 1.0 defines two algorithms that can run separately or together ([§3.3](#)):

- **NSCC (Network Signal-based Congestion Control):** sender-based. Adjusts a window using ECN marks combined with RTT measurements, plus a "Quick Adapt" response to loss.
- **RCCC (Receiver Credit-based Congestion Control):** the receiver hands out credits. It is good at last-hop incast but cannot see in-network congestion or outcast at the source.

The paper says the simple RCCC algorithm "is expected to perform well in isolation" for workloads without outcast, but "UE recommends enabling NSCC in oversubscribed networks as well", and recommends NSCC whenever RCCC is enabled (§3.3.2, §3.3.3). Congestion management "was designed and tested in simulation" for best-effort networks (§3.3). Field data on large mixed-vendor deployments is not yet public.

Packet trimming is an optional switch feature. When a packet would be dropped, the switch forwards only its header, giving the receiver an immediate signal to request a retransmit. It detects congestion drops, not corruption (§3.2.4).

The contrast with RoCEv2: there, loss is rare only because PFC prevents it. UE assumes loss can happen, recovers cheaply, and spreads load so it happens less often.

Link Level Retry and Credit-Based Flow Control

Both link features are optional and negotiated hop by hop via LLDP (§3.5):

- **LLR** retransmits locally (go-back-N) when a link corrupts frames, so a bit error on one hop does not become an end-to-end retransmit and tail-latency spike.
- **CBFC** is credit-based flow control that "can be used in conjunction with or in place of" PFC and generally needs less headroom buffer (§3.5.2).

UE was designed primarily for lossy networks, but "it also supports lossless networks, such as Slingshot" (§3.1.4). UEC's own summary says UET "performs well over lossless networks where Priority Flow Control is used" (UEC, 2024). A UET migration therefore does not force PFC to be ripped out on day one.

Public interoperability evidence for these link features exists but is narrow. Keysight reported a demonstration of UEC LLR and CBFC at 800GE line rate between its tester and a Broadcom Tomahawk Ultra switch at OFC in March 2026 (Keysight, 2026-03-16). That is tester-to-switch, not a multi-vendor NIC-to-switch fabric.

Side-by-side: RoCEv2 vs Ultra Ethernet

Dimension	RoCEv2	Ultra Ethernet (UET, 1.0 base)
Transport origin	InfiniBand transport over UDP/IP	New transport (UET) designed for Ethernet
Network assumption	Lossless (PFC) for good performance	Best effort by default. Lossless supported
Ordering	In-order on a reliable connection	RUD default (unordered). ROD available
Multipath	Per-QP ECMP hash. More QPs add entropy	Per-packet EV spraying over standard ECMP
Loss recovery	Go-back-N style on the connection	SACK bitmap, selective retransmit, optional trimming
Congestion control	ECN-based (commonly DCQCN), vendor-specific variants	NSCC (ECN + RTT) and/or RCCC (receiver credits). 1.1 adds Programmable Congestion Management
Link layer	PFC, PFC watchdogs	Optional LLR, optional CBFC (with or instead of PFC)

Dimension	RoCEv2	Ultra Ethernet (UET, 1.0 base)
Connection state	Pre-established QPs	Ephemeral PDCs, no QP construct
API	Verbs	libfabric. Verbs mapping added in 1.1
Security	No transport-level encryption defined; added separately by vendors	Optional TSS (AES-GCM)
Maturity (Oct 2026)	Production at large scale for years	Spec at 1.1. First silicon claims since 2025. Interop evidence limited

Who claims UEC support (as of 28 October 2026)

The table separates what each vendor says from what we could confirm. "Ready" and "compliant" are vendor terms. We found no third-party certification list; UEC compliance is vendor self-attestation.

Product	Type	Vendor wording	Status per vendor	Source
Broadcom Thor Ultra	800G NIC, PCIe Gen6 x16, 100G/200G SerDes	"fully feature compliant with UEC specification". Lists packet-level multipathing, out-of-order placement, selective retransmit, sender- and receiver-based CC, packet trimming and CSIG with Tomahawk 5/6	"now sampling" (2025-10-14). No later GA statement found	Broadcom, 2025-10-14
Broadcom Tomahawk Ultra	51.2 Tb/s switch ASIC	"also compliant with the UEC standard". Implements LLR and CBFC	Shipping (2025-07-15)	Broadcom, 2025-07-15
AMD Pensando Pollara 400	400G NIC, PCIe Gen5 x16	"industry's first Ultra Ethernet Consortium (UEC) ready" AI NIC. RoCEv2 compatible. P4-programmable	Available for purchase (2025-04-09), which predates the 1.0 spec	AMD brief, AMD blog
AMD Pensando Vulcano 800	800G NIC	"UEC-ready RDMA" alongside MRC and MetaRoCE (product brief, October 2026). July 2026 blog leads with MRC	No UEC 1.x compliance claim in AMD's page or brief	AMD Pensando page, Vulcano brief, AMD blog, 2026-07-23
Arista Etherlink switches	Switch family	"support the emerging Ultra Ethernet Consortium (UEC) standards"	Claim dates from June 2024, before the spec existed	Arista, 2024-06-05

Notes on reading this table:

- Thor Ultra's release says it works with "any UEC compliant switch". That makes the switch side of a UEC fabric a separate claim to verify per platform and NOS release.
- AMD's Pollara brief reports an average 25% All-Reduce improvement for its UEC-ready RDMA over RoCEv2 (4 QPs, 512 MB to 16 GB messages). This is a vendor measurement on a vendor-chosen setup ([AMD brief](#)).

- AMD's 2026 emphasis on MRC, described as "co-developed for OpenAI", shows that some large buyers are standardizing on multipath transports adjacent to UET rather than waiting for UEC conformance (AMD, 2026-07-23).
- We found no NVIDIA or Cisco statement claiming UEC 1.x compliance for a specific NIC or switch as of this date. That is a search result, not proof of absence, so ask.

Procurement questions that separate claims from capability

Ask every NIC vendor:

1. **Which spec point release** (1.0.x or 1.1) does the shipping firmware implement? Is a written conformance statement available per profile (AI Base, AI Full, HPC) and per delivery mode (RUD, ROD, UUD)?
2. **Which congestion-control algorithms** are implemented (NSCC, RCCC, both), at which revision? Remember that 1.0.1 and 1.0.2 corrected RCCC and CMS algorithms. Can CC be updated in firmware, as 1.1's programmable CC implies?
3. **Is it shipping, sampling or roadmap**, in which form factor (PCIe CEM, OCP 3.0), and at which port speed and SerDes rate?
4. **What runs in the same firmware**: RoCEv2 with what CC? Can RoCEv2 and UET coexist on one port during migration?
5. **Software stack**: libfabric provider maturity, and how your collective library (NCCL, RCCL or vendor fork) reaches UET. Verbs-only stacks need the 1.1 mapping or a shim.

Ask every switch vendor:

1. Which UE switch features are supported: packet trimming, CSIG, LLR, CBFC? On which ASIC and which NOS release?
2. ECN marking behavior suitable for NSCC, and telemetry to see spraying imbalance.
3. Interop evidence with **your** NIC, not a tester. Ask for a lab report naming firmware versions.

For MoE-heavy training, where all-to-all dominates, test spraying and incast behavior specifically.

Decision guide

Situation in 2026	Recommendation
Cluster must be in production within ~6 months	Build RoCEv2 with minimal lossless classes, ECN-based CC, PFC watchdogs and multi-QP collectives. Choose NICs and switches whose vendors publish a UET firmware path
1k+ GPU fabric procured for 2027 delivery	Make UET a scored requirement, not a gate. Require written conformance (point release, profile, CC), plus a mixed-vendor NIC-switch interop test before final acceptance
Single-vendor NIC+switch stack already chosen	Its proprietary RoCE extensions may deliver similar spraying and CC today. UEC mainly buys optionality for the next refresh

Situation in 2026	Recommendation
Mixed-vendor or multi-generation fleet	Plan for coexistence. UET supports lossless networks, so keep PFC on during migration and treat CBFC as a later change

Test gates before accepting either path:

1. **Spraying works:** per-uplink utilization stays balanced under all-reduce and all-to-all. Compare against the worked example's collision baseline.
2. **Loss is cheap:** inject loss or corruption on one link. Job completion time should degrade gracefully (selective retransmit, LLR) instead of collapsing.
3. **Incast is contained:** many-to-one tests at the target scale, with ECN and CC counters recorded.
4. **PFC is a backstop:** on RoCEv2, PFC pause counters stay low in steady state and watchdogs are configured and tested.
5. **Versions are pinned:** NIC firmware, NOS and libfabric or collective versions are recorded in the acceptance report.

GPU Smith helps teams write and run these fabric acceptance tests independently of the vendors involved. [Contact us.](#)

Sources

1. Hoefler, T. et al. "Ultra Ethernet's Design Principles and Architectural Innovations." arXiv 2508.08906 v1, 2025-08-12. <https://arxiv.org/abs/2508.08906> (HTML: <https://arxiv.org/html/2508.08906>)
2. Ultra Ethernet Consortium. "Ultra Ethernet Specification v1.1" and "v1.1 Release Notes", 2026-10-12. <https://ultraethernet.org/wp-content/uploads/sites/20/2026/10/UE-Specification-1.1.pdf> and <https://ultraethernet.org/wp-content/uploads/sites/20/2026/10/20261026-3711-UE-Specification-1.1-release-notes.pdf>
3. Ultra Ethernet Consortium. "Specification History" and "Compliance", checked 2026-10-28. <https://ultraethernet.org/specification-history/> and <https://ultraethernet.org/compliance/>
4. Linux Foundation / UEC. "UEC Launches Spec 1.0 Transforming Ethernet for AI and HPC at Scale," 2025-06-11. <https://www.linuxfoundation.org/press/uec-launches-spec-1.0>
5. Ultra Ethernet Consortium. "Ultra Ethernet Specification Update," updated 2024-08-29. <https://ultraethernet.org/ultra-ethernet-specification-update/>
6. Kerner, S. M. "Ultra Ethernet 1.1 extends the specification to scale-up networks and adds programmable congestion control." Network World, 2026-10-14. <https://www.networkworld.com/article/4235340/ultra-ethernet-1-1-extends-the-specification-to-scale-up-networks-and-adds-programmable-congestion-control.html>
7. Guo, C. et al. "RDMA over Commodity Ethernet at Scale." ACM SIGCOMM 2016, DOI 10.1145/2934872.2934908. <https://courses.grainger.illinois.edu/cs598hpn/fa2019/papers/microsoftrdma.pdf>

8. Zhu, Y. et al. "Congestion Control for Large-Scale RDMA Deployments." ACM SIGCOMM 2015, Microsoft Research. <https://www.microsoft.com/en-us/research/publication/congestion-control-for-large-scale-rdma-deployments/>
9. Broadcom. "Broadcom Introduces Industry's First 800G AI Ethernet NIC" (Thor Ultra), 2025-10-14. <https://investors.broadcom.com/news-releases/news-release-details/broadcom-introduces-industrys-first-800g-ai-ethernet-nic>
10. Broadcom. "Broadcom Ships Tomahawk Ultra: Reimagining the Ethernet Switch for HPC and AI Scale-up," 2025-07-15. <https://investors.broadcom.com/node/63341/pdf>
11. Keysight. "Keysight Advances AI Networking with Ultra Ethernet LLR and CBFC Interoperability Demonstration at OFC 2026," 2026-03-16. https://www.keysight.com/us/en/about/newsroom/news-releases/2026/0316_pr26-051-keysight-advances-ai-networking-with-ultra-ethernet-llr-and-cbfc-interoperability-demonstration-at-ofc-2026.html
12. AMD. "AMD Pensando Pollara 400G Adapter Card: Product Brief." <https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/product-briefs/pollara-product-brief.pdf>
13. AMD. "Powering Next Gen AI Infrastructure – A New Era of Scale-out with AMD Pensando AI NICs," 2025-04-09. <https://www.amd.com/en/blogs/2025/powering-next-gen-ai-infrastructure--a-new-era-of.html>
14. AMD. "AMD Pensando AI NICs" product page, Vulcano 800 product brief (PDF dated 2026-10-08) and "AMD Pensando Vulcano 800 AI NIC: Built to Scale out and Across," 2026-07-23. <https://www.amd.com/en/products/network-interface-cards/pensando.html>, <https://www.amd.com/content/dam/amd/en/documents/pensando-business-docs/product-brief/amd-vulcano-ai-nic-networking-product-brief.pdf> and <https://www.amd.com/en/blogs/2026/amd-pensando-vulcano-800-ai-nic-scale-out-and-across.html>
15. Arista Networks. "Arista Unveils Etherlink AI Networking Platforms," 2024-06-05. <https://www.arista.com/en/company/news/press-release/19493-pr-20240605>

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://arxiv.org/html/2508.08906>
2. <https://ultraethernet.org/specification-history>
3. <https://courses.grainger.illinois.edu/cs598hpn/fa2019/papers/microsoftrdma.pdf>
4. <https://www.microsoft.com/en-us/research/publication/congestion-control-for-large-scale-rdma-deployments>
5. <https://www.linuxfoundation.org/press/uec-launches-spec-1.0>
6. <https://arxiv.org/abs/2508.08906>
7. <https://ultraethernet.org/wp-content/uploads/sites/20/2026/10/20261026-3711-UE-Specification-1.1-release-notes.pdf>
8. <https://ultraethernet.org/compliance>
9. <https://www.networkworld.com/article/4235340/ultra-ethernet-1-1-extends-the-specification-to-scale-up-networks-and-adds-programmable-congestion-control.html>
10. <https://ultraethernet.org/wp-content/uploads/sites/20/2026/10/UE-Specification-1.1.pdf>
11. <https://ultraethernet.org/ultra-ethernet-specification-update>

12. https://www.keysight.com/us/en/about/newsroom/news-releases/2026/0316_pr26-051-keysight-advances-ai-networking-with-ultra-ethernet-llr-and-cbfc-interoperability-demonstration-at-ofc-2026.html
13. <https://investors.broadcom.com/news-releases/news-release-details/broadcom-introduces-industrys-first-800g-ai-ethernet-nic>
14. <https://investors.broadcom.com/node/63341/pdf>
15. <https://www.amd.com/content/dam/amd/en/documents/pensando-technical-docs/product-briefs/pollara-product-brief.pdf>
16. <https://www.amd.com/en/blogs/2025/powering-next-gen-ai-infrastructure--a-new-era-of.html>
17. <https://www.amd.com/en/products/network-interface-cards/pensando.html>
18. <https://www.amd.com/content/dam/amd/en/documents/pensando-business-docs/product-brief/amd-vulcano-ai-nic-networking-product-brief.pdf>
19. <https://www.amd.com/en/blogs/2026/amd-pensando-vulcano-800-ai-nic-scale-out-and-across.html>
20. <https://www.arista.com/en/company/news/press-release/19493-pr-20240605>

THE PUBLISHER

About GPU Smith

GPU Smith provides independent engineering and research for private AI infrastructure. We help technical buyers and operators reason about GPU workload sizing, hardware selection, procurement, deployment and inference operations, from the compute system to the networking, power and cooling requirements around it.

Start with the workload

Useful infrastructure decisions begin with the models, data, throughput, latency and operating constraints a team actually has. GPU Smith helps connect those requirements with system design and deployment choices. The goal is a privately controlled AI environment whose capacity and operational demands are understood before a purchasing decision.

Hardware and supplier research

Our public reference library covers GPUs, complete systems and networking components. The server-vendor directory and manufacturing research help teams investigate suppliers, compare options and follow sources behind technical claims. We distinguish manufacturer specifications from measured benchmarks and advertised capabilities from tested configurations. Reference pages are research resources, not a live inventory listing or a binding equipment quote.

Deployment and operations

GPU Smith's engineering scope includes infrastructure integration, networking, power, cooling and the practical operation of inference workloads. Our published articles explain the tradeoffs behind deployment, procurement and ongoing operations so teams can ask better questions and document decisions.

Work with GPU Smith

Explore the [hardware library](#), [GPU references](#), [networking references](#) and [vendor research](#). [Contact GPU Smith](#) with your workloads and deployment constraints to discuss a project and confirm current scope, pricing and availability.

Inclusion of a manufacturer or product in our research does not imply a partnership, certification or endorsement.

Website: <https://gpusmith.com>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.