



Statements of Support for the UALink™ 200G 1.0 Specification

AMD

"Ultra Accelerator Link represents the industry working together to develop an open, high-speed, scale-up accelerator interconnect technology that advances next-generation AI cluster performance. As a founding member, AMD is excited to see UALink fostering broad support across the ecosystem, enabling rapid innovation and diverse solutions. UALink utilizes 224Gbps SerDes and delivers a power-efficient design and cost-effectiveness to deliver 224Gbps data rate per lane, making it an ideal choice for demanding workloads in AI and HPC, and we expect to adopt the standard in our future products."

Forrest Norrod, Executive Vice President, General Manager, Data Center Solutions, AMD.

AnaPass, Inc.

As the demand for high-speed, low-latency networks in scale-up AI systems continues to grow, the release of the UALink 200G 1.0 Specification marks a significant milestone for the industry. By establishing an open standard, UALink enables more efficient communication between AI accelerators and fosters greater collaboration among industry participants. AnaPass brings deep expertise in high-speed interconnect technology and a strong track record in standardization efforts with global partners. Our solutions are designed for seamless compatibility with leading GPU companies. We look forward to contributing to the ongoing advancements of the UALink Consortium.

Sukkyun "Albert" Hong, Director of R&D Department, AnaPass

Astera Labs

"The launch of the UALink™ 200G 1.0 Specification is a game-changer for AI infrastructure, delivering an open scale-up memory semantic AI fabric to connect hundreds of GPUs and accelerators. Hyperscalers and data center operators will be able to scale AI workloads efficiently, while gaining performance and flexibility to support next-generation AI models. Astera Labs proudly supports this new standard and is uniquely positioned to lead in this space with our deep expertise on silicon-based connectivity solutions for scalable AI infrastructure."

Chris Petersen, UALink Consortium Board Director and Fellow, Technology & Ecosystems, Astera Labs

Auradine

"At Auradine, we believe that the UALink Consortium and the 200G 1.0 Specification are pivotal in supporting the rapid growth of AI. With the 1.0 specification fostering open standards for scale-up networking, UALink will enable transformative improvements in AI data center performance for training and inference use cases. We look forward to working with ecosystem partners and customers to leverage UALink in driving world-class solutions that deliver breakthrough value for AI infrastructure."

Rajiv Khemani, Co-Founder and CEO, Auradine.



Cisco

"The UALink Consortium is making rapid progress in addressing key challenges with building optimized AI supercomputers. Cisco is committed to accelerating open, scalable technology and solutions for our customers, and we look forward to continued collaboration with UALink members as we advance this vision."

Rakesh Chopra, SVP and Cisco Fellow, Common Hardware Group, Cisco

Clustered Systems

"UALink opens the door to huge gains in GPU compute efficiency. Flat networks with 1,600 Tbps bi-directional bandwidth connecting over 1,000 GPU ports assures efficient processing of multi-trillion parameter LLMs without sacrificing accuracy or performance. We are proud to be a part of the visionary UALink team developing products for this key technology."

Robert Lipp, CEO, Clustered Systems

Enfabrica

"The rapid evolution of AI computing workloads necessitates an open interconnect standard between accelerators - one that is rooted in the native communication semantics of GPUs and GPU-like processors. UALink provides this essential standard, providing AI infrastructure builders choices along the performance-efficiency curve for flexibly scaling up AI compute racks and pods through disaggregated switching elements. Enfabrica wholeheartedly supports the UALink Consortium's mission, and as a contributing member, will help shape the product landscape for open, ultra-scalable AI infrastructure."

Shrijeet Mukherjee, Chief Development Officer & Co-Founder, Enfabrica

Hewlett Packard Enterprise

"Hewlett Packard Enterprise celebrates the latest milestone release of the UALink 200G 1.0 specification which advances open standards for high performance networking that we anticipate will significantly progress AI system designs in the near future. These specifications will improve common understanding, leading to more accessible supercomputing and large AI clusters focused on model training, tuning, and inferencing."

Trish Damkroger, Senior Vice President and General Manager, HPC & AI Infrastructure Solutions, Hewlett Packard Enterprise

Intel

"Intel is proud to support and contribute to UALink 1.0, an open, standards-based scale-up interconnect as we look to our next-generation AI products. We recognize the critical role of memory semantics in enabling low-latency all-reduce for LLM inference applications."

Mark Davis, Fellow and Chief AI Product and System Architect, Intel

Marvell

"Marvell is pleased to have contributed to this new 200G UALink specification. The specification will serve to ensure predictable, lossless, low- latency communication required for next-



generation AI compute architectures. Marvell has long been a champion of open industry standards and is proud to continue that legacy with the UALink Consortium.”

Nick Kucharewski, Senior Vice President and General Manager, Network Switching Business Unit and Cloud Platform Business Unit, Marvell.

Synopsys

“To meet the demands of ever-growing AI workloads, hyperscale data centers need to maximize performance and bandwidth of every connection while maintaining ecosystem interoperability. As a leading provider for best-in-class interface IP and board members to the UALink consortium, we are committed to enabling our customers with open, standardized solutions like UALink to address the industry's need for scalable and high-performance interconnects.”

Neeraj Paliwal, Senior Vice President of IP Product Management, Synopsys

UnifabriX

“UnifabriX is excited to welcome the release of the UALink 200G 1.0 specification, a pivotal and long-awaited milestone that marks the beginning of a journey that will reshape the way AI compute is built and used. UALink technology provides the fundamental backbone for high-performance communication between GPUs, purpose-built AI accelerators, and memory, enabling a massive scale-up of compute power to tackle the increasingly demanding tasks assigned to AI workloads. UnifabriX is committed to supporting this growing ecosystem and contributing to its success.”

Ronen Hyatt, CEO, Chief Architect and Founder, UnifabriX