ULTRA ACCELERATOR LINK™

NEWS RELEASE

UALink Consortium Releases the Ultra Accelerator Link 200G 1.0 Specification

News Highlights

- Ultra Accelerator Link™ (UALink™) is the open scale-up interconnect for next generation AI workloads.
- The open industry standard is defined by the UALink Consortium – representing more than 85 member companies.
- The UALink 200G 1.0 Specification is now available to the public.

Beaverton, Ore. – April 8, 2025 – The UALink Consortium today announced the ratification of the UALink 200G 1.0 Specification, which defines a low-latency, high-bandwidth interconnect for communication between accelerators and switches in AI computing pods. The UALink 1.0 Specification enables 200G per lane scale-up connection for up to 1,024 accelerators within an AI computing pod, delivering the open standard interconnect for next-generation AI cluster performance.

"As the demand for AI compute grows, we are delighted to deliver an essential, open industry standard technology that enables next-generation AI/ML applications to the market," said Kurtis Bowman, UALink Consortium Board Chair. "UALink is the only memory semantic solution for scale-up AI optimized for lower power, latency and cost while increasing effective bandwidth. The groundbreaking performance made possible with the UALink 200G 1.0 Specification will revolutionize how Cloud Service Providers, System OEMs, and IP/Silicon Providers approach AI workloads."

UALink creates a switch ecosystem for accelerators – supporting critical performance for emerging AI and HPC workloads. It enables accelerator-to-accelerator communication across system nodes using read, write, and atomic transactions and defines a set of protocols and interfaces enabling the creation of multi-node systems for AI applications.

Key Benefits of UALink

- High Performance
  - Low-latency, high-bandwidth interconnect for hundreds of accelerators in a pod.
  - Provides a simple load/store protocol with the same raw speed as Ethernet with the latency of PCIe® switches.
  - Designed for deterministic performance achieving 93% effective peak bandwidth.
- Low Power
  - Enables a highly efficient switch design that reduces power and complexity.
- Cost Efficient
  - Uses significantly smaller die area for link stack, lowering power and acquisition costs, resulting in decreased Total Cost of Ownership (TCO).
  - Increased bandwidth efficiency further enables lower TCO.
- Open and Standardized
  - Multiple vendors are developing UALink accelerators and switches.
  - Harnesses member company innovation to drive leading-edge features into the specification and interoperable products to the market.

"AI is advancing at an unprecedented pace, ushering in a new era of AI reasoning with new scaling laws. As the demand for compute surges and speed requirements continue to grow exponentially, scale-up

interconnect solutions must evolve to keep pace with these rapidly changing AI workload requirements," said Sameh Boujelbene, VP at Dell'Oro Group. "We are thrilled to see the release of the UALink 1.0 Specification, which rises to this challenge by enabling 200G per lane scale-up connections for up to 1,024 accelerators within the same AI computing pod. This milestone marks a significant step forward in addressing the demand of next-generation AI infrastructure."

"With the release of the UALink 200G 1.0 Specification, the UALink Consortium's member companies are actively building an open ecosystem for scale-up accelerator connectivity," said Peter Onufryk, UALink Consortium President. "We are excited to witness the variety of solutions that will soon be entering the market and enabling future AI applications."

The UALink 200G 1.0 Specification is available for public download at https://ualinkconsortium.org/specification/

Additional Resources:

- White Paper: Introducing the UALink 200G 1.0 Specification
- UALink Consortium Member Company Statements of Support

About Ultra Accelerator Link Consortium

The Ultra Accelerator Link (UALink) Consortium, incorporated in October 2024, is the open industry standard group dedicated to developing the UALink specifications, a high-speed, scale-up accelerator interconnect technology that advances next-generation AI & HPC cluster performance. The consortium is led by a board made up of stalwarts of the industry; Alibaba, AMD, Apple, Astera Labs, AWS, Cisco, Google, HPE, Intel, Meta, Microsoft, and Synopsys. The Consortium develops technical specifications that facilitate breakthrough performance for emerging AI usage models while supporting an open ecosystem for data center accelerators. For more information on the UALink Consortium, please visit www.UALinkConsortium.org

Contacts
Nolan Morgan
+1 971-271-2657
press@members.ualinkconsortium.org

###