



Contribution ID: 497

Type: **not specified**

# Toward Mainline Linux Support for UALink AI Scale-UP Interconnect

*Friday 12 December 2025 10:00 (45 minutes)*

## Abstract

UALink is a new open, scale-up interconnect standard for AI that tightly couples CPU, GPU, and NPU memory into a single high-bandwidth, low-latency domain, enabling ultra-fast load/store/atomic operations across hundreds to thousands of accelerators in a pod. Unlike vendor-specific fabrics such as NVIDIA's NVLink, UALink is positioned as a multi-vendor standard for open accelerator interconnects.

This session proposes the bring-up of a dedicated drivers/ualink subsystem and a minimal viable roadmap towards an upstream-first path for new upcoming UALink hardware (Q3 2026).

Key points we'll explore:

- **Vendor-agnostic dma\_buf memory semantics:** Defining tightly coupled memory mappings in a vendor-neutral way via dma\_buf, including pinning, revocation, and peer access across heterogeneous accelerators.
- **UALink daemon for fabric-wide coordination:** Introducing a UALink userspace daemon to coordinate UALink nodes across a rack, enabling cross-node discovery, topology management, and policy enforcement.
- **Security and confidential computing:** Positioning UALink security as a first-class part of confidential computing, protecting multi-tenant and multi-OS boundaries while still allowing high-performance accelerator-to-accelerator data flow.
- **RAS and port migration:** Treating port migration and related scenarios as part of comprehensive RAS handling, and identifying where UALink can reuse kernel components from CXL (enumeration, RAS, error handling) versus where fabric-scale accelerator memory requires new abstractions.

**Goal:** Reach consensus on minimal kernel interfaces to make UALink fabrics a first-class Linux interconnect, complementing CXL while offering an open alternative to NVIDIA's proprietary NVLink and a clear standout from Broadcom SUE in delivering a multi-vendor, fabric-scale accelerator interconnect.

While awaiting new UALink-based devices by Q3 2026, we propose a QEMU-based UALink hardware simulation with a test DRM memory device to create a fully functional proof-of-concept of the proposed UALink core subsystem.

## Tagline

Bringing UALink into Linux: defining dma\_buf semantics and secure CPU/GPU/NPU memory sharing, positioning an open multi-vendor alternative to NVIDIA's proprietary NVLink that complements CXL and clearly differentiates itself from Broadcom SUE.

**Primary author:** SINYUK, Konstantin

**Presenter:** SINYUK, Konstantin

**Session Classification:** Birds of a Feather (BoF)

**Track Classification:** Birds of a Feather (BoF)