ADLINK



ADLINK COM-HPC-ALT

Breaking Performance Limits in Autonomous Mobility to Accelerate Deployment Readiness at Scale

COM-HPC™

February 2026

TECHNICAL CASE STUDY

Distribuída no Brasil pela LRI | vendas@lri.com.br | 51 2129-3000

Scalable Intelligence for Autonomous Mobility

Building a fully autonomous vehicle requires every system to operate seamlessly in real time. At the core of this capability is the compute platform—a high-performance "brain" that powers autonomous driving software and fuses data from multiple sensors such as LiDAR, radar, and cameras. This compute platform must deliver deterministic performance while meeting strict requirements for thermal management, power efficiency, mechanical integration, and long-term reliability in real-world operating environments.

A leading autonomous mobility provider faced the challenge of scaling its commercial deployments through a partnership with a global ride-hailing and on-demand delivery platform. In this model, the ride-hailing platform provides customer access, marketplace scale, regulatory navigation, and day-to-day fleet operations, including maintenance, cleaning, charging, and rider support. The autonomous mobility provider delivers proprietary self-driving technology, manages vehicle testing, and supplies autonomous hardware for both robotaxi and delivery operations. To succeed at this scale, performance alone was not sufficient, the compute platform within the autonomous vehicles needed to ensure fleet-wide consistency, long-term availability, and rapid integration, keeping pace with aggressive commercialization timelines.

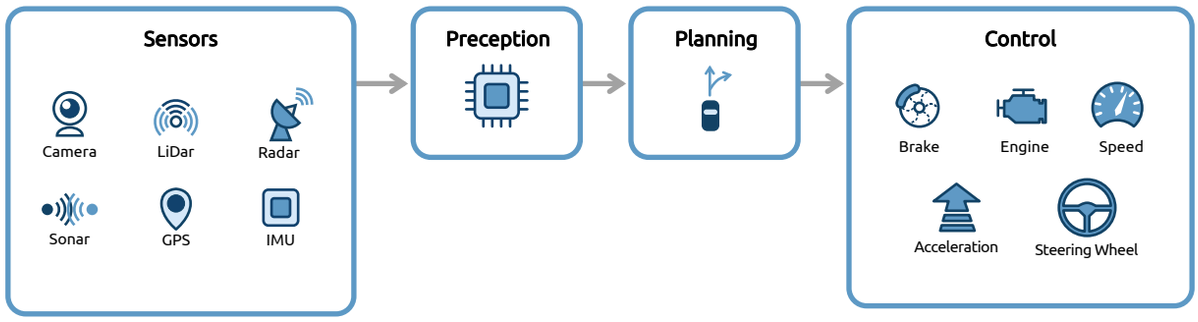


Figure 1: Block diagram of autonomous vehicle

This autonomous mobility provider operates a dual-track model:

- **Autonomous cars / robotaxis** – Retrofitted electric vehicles equipped with a proprietary self-driving stack, designed for fully autonomous ride-hailing and ride-sharing services.
- **Autonomous delivery robots** – Compact, sidewalk-capable robots for last-mile delivery, capable of autonomously navigating pedestrian pathways, crosswalks, and traffic signals.

Both vehicle types share a common AI and hardware architecture, enabling advancements in perception, path planning, and compute efficiency to benefit both passenger mobility and delivery logistics. ADLINK's scalable and reliable compute platforms provide the common foundation that allows these innovations to transition from development to fleet-ready, commercial deployment.



Distribuída no Brasil pela LRI | vendas@lri.com.br | 51 2129-3000

The Autonomous Pivot Point

As the autonomous mobility provider advanced its next-generation autonomous taxi platform, several mission-critical challenges emerged:

- Performance gaps: Existing CPU options could not deliver the required core count and sustained performance, especially under power-throttling conditions.
- Thermal and environmental constraints: The compute platform had to operate reliably across a temperature range of 0°C to 60 °C.

Commercial risk: Failure to meet performance and delivery timelines would have jeopardized a major autonomous taxi partnership.

Rapid development schedule: From initial engagement to mass production readiness, the platform needed to be validated in just four months.

Solving these challenges required more than high-performance hardware—it demanded a partner that could move quickly while delivering industrial-grade reliability. By adopting ADLINK's COM-HPC-ALT platform, powered by Ampere® Altra® processors, the team secured a server-class compute foundation that closed performance gaps, accelerated validation, and helped move autonomous taxi programs from development to real-world deployment.



COM-HPC-ALT Module

Ecosystem Depth as a Competitive Advantage

A key differentiator was ADLINK's deep integration across the CPU ecosystem. ADLINK maintains long-standing, strategic relationships with leading processor vendors, including Intel, AMD, Ampere, NXP, MediaTek, and Qualcomm, enabling early access, roadmap alignment, and rapid platform readiness across architectures. This ecosystem depth allows autonomous mobility providers to adopt the right compute technology at the right time, without the integration risk, lifecycle uncertainty, or scale limitations often encountered with both legacy suppliers and smaller niche vendors.

Strategic Modular Design Advantage for Fleet-Scale Autonomy

Delivering autonomous mobility at scale requires more than raw compute power, it demands a platform that is flexible, upgradeable, and reliable across an entire fleet. ADLINK's COM-HPC modules provide a modular foundation, enabling developers to focus on autonomy, perception, and decision-making while leveraging pre-integrated CPU, memory, I/O, and optional AI acceleration.

This approach accelerates integration, shortens design cycles, and reduces risk, because only vehicle-specific interfaces need to be customized. As compute requirements evolve with AI models, sensor payloads, and fleet expansion, modules can be upgraded without redesigning the overall system, ensuring long-term scalability and operational continuity.



By combining modularity with industrial-grade reliability, ADLINK enables autonomous mobility providers to deploy fleets quickly, adopt new compute technologies seamlessly, and maintain cost-effective, future-ready operations. This ensures that companies of any size or with legacy system constraints can adopt high-performance compute without redesigning their entire platform.

Distribuída no Brasil pela LRI | vendas@lri.com.br | 51 2129-3000

High-Performance Compute — Where Power, Thermals, and Space Work as One

Maximum Compute Density in Automotive Enclosures

Autonomous vehicles demand exceptional compute density. Advanced driving stacks must process data from cameras, lidar, radar, and vehicle systems simultaneously, all within constrained automotive enclosures exposed to wide temperature ranges.

High Core Performance Without Throttling

ADLINK's COM-HPC-ALT module, based on the COM-HPC Type E specification, delivered the performance leap required. Equipped with the Ampere® Altra® M128-26 processor, the platform provided a high core count and strong performance-per-watt and TDP efficiency, enabling the autonomous driving stack to run efficiently without hitting power-throttling limits.

Sustained Performance in Electric Vehicles

By validating and supporting the Ampere® Altra® M128-26 CPU, ADLINK delivered a solution that balanced high core density with efficient power consumption. This ensured sustained performance for autonomous driving workloads without thermal instability, a key requirement inside electric vehicles. Efficient power consumption not only ensured thermal stability, but also helped preserve vehicle range and reduce operating costs at fleet scale.

Server-Class Architecture for Continuous Operation

With support for massive 6x 32 GB/64 GB ECC memory, COM-HPC Server I/O, and a server-grade architecture, the solution met both compute scalability and reliability requirements, critical for continuous operation in autonomous taxis and fleet vehicles. The platform delivers stable 24/7 operation and predictable availability across 2–3-year generation lifecycles, with ongoing support and lifecycle management to ensure fleets remain operational and scalable over the years.

ADLINK's COM-HPC-ALT platform unifies these requirements into a single, optimized solution, delivering scalable performance that fits the physical, electrical, and thermal realities of automotive deployment.

Engineering Collaboration & Scalable Compute Compress Development Timelines

ADLINK's engagement went far beyond hardware delivery. Their engineering team worked closely with the customer to accelerate integration, providing:

- Firmware and software tuning aligned to the autonomous driving workload
- Platform optimization and validation with the Ampere® Altra® M128-26 SoC
- Rapid sample delivery to support customer demonstrations and partner validation

Developing a custom CPU board would have added engineering risk, extended timelines, and increased overall costs. To avoid these constraints, the autonomous mobility provider leveraged ADLINK's COM-HPC-ALT modules to enable faster integration without sacrificing performance and scalability.

Distribuída no Brasil pela LRI | vendas@lri.com.br | 51 2129-3000



Scalable, Cost-Efficient Compute

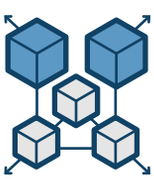
The COM-HPC-ALT architecture significantly reduced system complexity while preserving flexibility for future compute upgrades as autonomous software demands continue to grow.

This hands-on collaboration, paired with a ready-to-deploy, modular platform, helped compress the entire development cycle, enabling the solution to move from inception to production readiness in just four months. At the same time, it provided a cost-effective, scalable foundation for future autonomous mobility innovations.

Measurable Impact:

From Solution Validation to Commercial Success

The results were tangible and immediate:



- **Rapid availability of reliable samples** enabled downstream customer and partner validation.



- **Critical gaps closed,** removing compute bottlenecks in the autonomous driving stack.



- **Accelerated deal closure with swift contract execution,** enabling rapid commercialization of the autonomous taxi solution.

By delivering the right platform at the right time, ADLINK helped transform a critical technical challenge into a business-enabling success.

Distribuída no Brasil pela LRI | vendas@lri.com.br | 51 2129-3000

Blueprint for Autonomous Mobility Innovators — Industrial Reliability You Can Trust

This collaboration highlights how the autonomous mobility provider, across robotaxis, delivery robots, and next-generation mobility platforms, can leverage ADLINK's COM-HPC solutions to stay ahead of tightening timelines, rising performance demands, and cost pressures.

COM-HPC™

Core advantages include:

- High-core-count, power-efficient compute for autonomous driving
- Industrial-grade reliability for automotive environments
- Predictable product longevity for fleet deployment
- Modular COM-HPC architecture for future scalability
- Accelerated time-to-market through hands-on engineering support
- Lower risk via off-the-shelf, cost-efficient server-grade modules

Whether powering robotaxis, autonomous delivery robots, or other edge-AI mobility systems, ADLINK's COM-HPC platforms provide the performance, stability, flexibility, and cost efficiency needed to move from concept to commercialization with confidence.

By combining server-class compute, industrial reliability, and close engineering collaboration, ADLINK turns complex challenges into practical solutions, helping autonomous innovators focus on what matters most: deploying safe, scalable, and intelligent mobility systems at speed.

ADLINK turns challenges into solutions —
accelerate the future of your autonomous mobility with COM-HPC.

ADLINK
LEADING EDGE AI COMPUTING

Tel: +886-3-216-5088
Fax: +886-3-328-5706
Email: ipi.wiki@adlinktech.com
www.adlinktech.com

All products and company names listed are trademarks or trade names of their respective companies Updated February 2026 © ADLINK Technology, Inc. All rights reserved. All specifications, pricing, and product names are subject to change without notice.

LRI Distribuída no Brasil pela LRI | vendas@lri.com.br | 51 2129-3000