When Driving Intelligence Becomes Abundant

NVIDIA's Alpamayo 2 Super is the clearest industrial-scale test yet of model commoditization in Physical AI. It makes NVIDIA's factory more valuable, and the deployer's Operating Habitat decisive.

Tal Cohen | talcohen.ai | Independent analysis

A new autonomous-vehicle team can now download NVIDIA Alpamayo 2 Super, use it commercially, modify it, and train smaller driving models from it. The model-license fee is zero.

What the team cannot download is a robotaxi.

It still needs a vehicle architecture, live integration with steering and braking, an approved operating territory, emergency procedures, a safety case, an insurer, and an operating record. Alpamayo 2 Super is a teacher-scale development model, not the finished production driver sitting behind the wheel.

That gap is the story.

Open driving research did not begin this week. What changed is market structure. One of Physical AI's most important platform suppliers released a high-capability driving-reasoning foundation under commercially reusable terms, with an explicit pathway through training, simulation, distillation, and in-vehicle compute.

That combination is why I treat Alpamayo as the clearest commoditization event yet in a consequential Physical AI domain. A strategically important capability that typically demanded a heavily funded program can now be reused by outsiders. The direction of scarcity has changed.

Language models showed how quickly intelligence can lose scarcity. In 2023, an internal Google memo argued that open models were closing the gap with proprietary systems. The strongest models still command a premium, but access to capable intelligence increasingly does not.

Most of that transition occurred before model outputs directly moved machines. Physical AI changes the stakes. Digital errors can often be contained, reversed, or retried. A failed maneuver cannot be rolled back. It changes the physical world at road speed. The danger concentrates in the long tail, the rare situations a system meets once in millions of miles, where one small mistake costs lives. Finding those situations and proving a system survives them is the hardest work in the field.

The model is the event. What surrounds it is the strategy.

THE ARGUMENT

Alpamayo makes a major intelligence layer portable. NVIDIA monetizes the reusable factory around it. Durable control migrates to the least portable complement required for deployment: the Operating Habitat.

Six jobs, one useful distinction

Strip away the jargon and a self-driving service has six jobs.

1. See. Turn cameras, radar, and other sensors into a live picture of the road.

2. Decide. Choose what the vehicle should do next, including in rare, messy situations no engineer scripted.

3. Act. Turn the decision into steering, braking, and acceleration on a real machine.

4. Prove. Show, with evidence, that the configured system behaves safely across the situations it is permitted to face, and show it again after every meaningful change.

5. Operate. Run the service every day: maintenance, monitoring, recovery, and remote help when a vehicle gets stuck.

6. Answer. Stand behind the operation when something goes wrong, in front of an insurer, a regulator, a city, or a court.

The first three form the technical driving loop. The last three turn that loop into a service an institution can authorize and defend.

For a decade, building a driver that decides well across the long tail was one of the field's costliest walls, and the one a new entrant could least shortcut. Alpamayo lowers that wall for anyone building a driver. The long tail itself stays where it was. Finding those rare situations, assessing them, and resolving them remains the expensive work, and that work lives in the last three jobs.

Hold on to the six jobs. They describe a warehouse robot, a surgical assistant, and a cargo ship almost as well as a car.

What NVIDIA actually opened

NVIDIA unveiled Alpamayo 2 Super on May 31, 2026. On August 4, it made the weights downloadable for commercial use under OpenMDW-1.1. The distinction matters: May introduced the capability; August changed the economics of who may build on it.

The released interface takes surround-camera video, text, and the vehicle's recent motion. It can output a planned trajectory, a meta-action such as yield or change lanes, and a readable reasoning trace that NVIDIA calls a Chain of Causation. It can also answer questions about a scene, ground answers in image regions, and generate structured reasoning labels for driving clips. This is an end-to-end model in the technical sense. One network carries the work from camera pixels to a driving decision, with no separate perception module handing objects to a separate planner. The reach therefore covers all three layers of the classic stack. Perception, decision, and the trajectory that feeds

the vehicle's controllers now arrive in one downloadable artifact, so perception on its own stops being a place to differentiate.

The trace deserves one caution. A readable explanation helps an engineer inspect a decision, challenge it, and compare it with another system. It does not establish that the decision is correct, or that the written explanation faithfully describes how the model reached it.

A REASONING TRACE IS NOT A SAFETY CASE

Reasoning traces make a model easier to question. Proof of safe behavior comes from testing the configured, deployed system.

For a development team, the model can participate in several workflows:

Role	What it does, in plain terms
Propose	Watch a driving scene and produce a trajectory and intended action.
Label and ground	Add structured descriptions, reasoning labels, and image-region references to recorded footage.
Critique	Review another driving system's decision and help surface weak behavior.
Reference	Provide a common model against which an internal system can be evaluated.
Teach	Produce trajectories and reasoning outputs that help train a smaller model.
Explain	Answer questions about what is happening in a scene and why an action may be appropriate.

Several model-level tasks that once required different foundations can now share one. The surrounding work does not disappear. Scenario generation belongs to tools such as Cosmos-Dreams and simulation; evaluation still requires an evaluation system.

Distillation deserves a plain definition because it is the intended road into the car. A large teacher model produces responses across many driving examples. A smaller model learns to reproduce the useful behavior, is optimized for real-time inference, and can fit the compute and latency envelope of a vehicle. NVIDIA positions Alpamayo 2 Super for cloud-based development and as a teacher. The distilled production model is a separate artifact that still has to be integrated and validated.

NVIDIA reports that, in its testing, Alpamayo 2 Super ranked first on LingoQA among nearly 40 models and first across the autonomous-driving benchmarks NVIDIA evaluated. Those vendor-reported results measure model capability. They do not certify a vehicle, a software release, or a public-road operation.

Now the license. OpenMDW-1.1 grants broad, royalty-free rights to use, modify, and redistribute the released model materials, including commercially and without a field-of-use limit. Redistributions must preserve the license and origin notices. Bringing, maintaining, or voluntarily participating in a non-defensive patent or copyright action alleging that the materials infringe can terminate the grant. The license imposes no restrictions or obligations on model outputs; whether an output is ownable remains a separate question of applicable law. The materials are provided without warranty, and users remain responsible for clearing third-party rights and carrying the deployment risk.

The short version is powerful: a qualified team can build a commercial derivative without paying NVIDIA a model royalty.

The equally important version is longer: compute, storage, adaptation, integration, simulation, validation, certification, operations, and insurance still cost money.

Against the six jobs, Alpamayo collapses seeing and deciding into one artifact and hands a trajectory to the vehicle's controllers. Acting on that trajectory in a specific car, proving the configured system across the long tail, operating the fleet, and answering for the result all stay where they were.

So why spend a fortune building it and charge no model royalty?

The model is the event. The factory is the strategy.

The quick answer is chips. It is directionally right and strategically incomplete.

Technology companies have used this move before: make one important complement abundant so demand grows for the layers you control. Android expanded the market for Google's services. Open-source software expanded the market for cloud infrastructure, managed operation, and support. The IBM PC offers the warning on the other side: opening a layer can enrich somebody else's bottleneck if you fail to control what becomes scarce next.

NVIDIA is trying to control the reusable pathway around open Physical AI. Follow a driving system from cloud to car.

First, it learns. Cosmos 3 Super Reasoner supplies the physical-world foundation beneath Alpamayo. DGX systems and cloud capacity train and adapt models. An open Physical AI Data Factory blueprint organizes high-volume data work, while the compute running it remains commercial.

Then, it rehearses. Omniverse NuRec reconstructs recorded drives. Cosmos-Dreams generates difficult scenarios. AlpaSim evaluates behavior in closed loop, and AlpaGym is an early-stage reinforcement-learning framework whose current public documentation covers Alpamayo 1.5 rather than a finished 2 Super workflow. OVX and RTX systems run simulation workloads. These tools differ in license and maturity, and today they arrive as components rather than as a working toolchain. Closing that gap is itself a business, and it currently sits outside NVIDIA.

Then, it enters the vehicle. DRIVE Hyperion 10 provides a reference compute-and-sensor architecture around DRIVE AGX Thor. NVIDIA Halos OS, built on safety-certified DriveOS, supplies the in-vehicle software and safety foundation. NVIDIA Halos spans design, simulation, validation, runtime guardrails, and safety-case tooling. Customers can bring their own driver, use a partner's, or buy NVIDIA DRIVE AV as a commercial full-stack option.

The economics are therefore not simply “free model, paid software.” Several models, recipes, blueprints, and tools are open or available without a software-license fee. Data can carry separate restrictions. Hardware, cloud capacity, enterprise support, and production platforms remain commercial. Free is a licensing fact, not a total-cost fact.

NVIDIA does not need every customer to buy the complete stack. Each serious workload that lands on CUDA, DGX, OVX, RTX, Thor, Hyperion, DriveOS, or DRIVE AV creates value. NVIDIA is betting that performance, integration, and lower friction will pull builders down its pathway.

This is the factory strategy: open inputs widen the entrance; the factory sells the machinery of participation.

FIGURE 1 | NVIDIA'S PHYSICAL AI MOVE

Independent analysis. Open or free describes access or licensing, not total operating cost; terms vary by component.

NVIDIA's opportunity is the reusable factory. The deployer's durable asset, and unavoidable responsibility, is the Operating Habitat.

The deployment owner is the institution seeking and exercising permission to operate the configured system. It must assemble and maintain the Habitat, even when legal responsibility remains distributed across providers, integrators, automakers, and operators.

THE TWO-OWNER ARCHITECTURE

The factory scales what can travel. The Habitat holds what cannot: local constraints, authority, operational history, and consequences.

Where scarcity moves

Put the same Alpamayo derivative into two cities.

In the first, its operating territory is explicit. Software changes are controlled. Emergency procedures are agreed. Authority is assigned. Maintenance is disciplined. Rollback works. Independent evidence is preserved. An insurer can inspect the record.

In the second, none of those things is settled.

The intelligence is identical. The deployable system is not.

The difference is one persistent system around the machine: the constraints it must obey, the authority that may direct or stop it, the record of what happened, and the institution that carries the consequences.

I call that system the Operating Habitat.

THE OPERATING HABITAT

The persistent organizational, technical, and governance environment that makes a machine deployable, governable, and improvable in one particular place, even when models, machines, or vendors change. Defined in full in The Case for Habitat.

Four pillars

1. Constraints. The mission, operating territory, executable rules, and conditions under which the system may begin, continue, or must stop.

2. Authority. Who may approve, override, stop, restart, hand over, authorize an exception, or escalate.

3. Record. Which configuration acted; what changed and who approved it; what was monitored; where people intervened; how maintenance and recovery were handled; and what the independent evidence shows.

4. Consequences. Who insures the operation, reports to the regulator, carries responsibility after harm, and must change the system afterward.

The Model-Swap Test is a counterfactual, not a claim that Alpamayo can be dropped into a production vehicle. Ask: if the model or vendor changed tomorrow, what would still have to remain intact? The constraints. The authority. The accumulated record. The consequences. Those surviving obligations are the Habitat.

August 4 made that counterfactual economically credible across the industry. The surrounding obligations became easier to see.

Waymo illustrates the result. It may never use Alpamayo. But its defined service areas, published safety framework, remote-assistance operation, first-responder procedures, vehicle and release controls, public crash data, regulatory reporting, and deployed operating history survive any model choice.

NVIDIA can support much of the technical machinery behind the first three pillars and provide evidence used in the fourth. An integrator can assemble components. A validator can test them. A city can impose rules. An insurer can make evidence a condition of coverage. None of them can own the four as one operating system on behalf of every deployment.

That responsibility remains with the deployer.

PRODUCTS CROSS. RESPONSIBILITY DOES NOT.

NVIDIA can mass-produce what the world's deployments share. Each deployer must still build the Habitat where its machines act.

“Going NVIDIA” is not one decision

Between building everything alone and buying NVIDIA's full driving software sit four strategic motions.

1. **Use.** Employ Alpamayo as a teacher, critic, labeler, or reference, and keep your own production driver.
2. **Adapt.** Fine-tune or distill a proprietary derivative around your vehicle, driving policy, data, and territory.
3. **Build on.** Adopt selected pieces of the factory: compute, simulation, runtime, reference architecture, safety tooling, or the commercial DRIVE AV stack.
4. **Operate across.** Run NVIDIA-powered and other vehicles through one service without choosing one universal driver.

The surrounding ecosystem has a fifth posture: plug in. Integrators, sensor makers, data platforms, and independent validators can sell across many models and stacks.

The right motion depends on three questions. What do you already own? Which NVIDIA layer lowers your cost without erasing your differentiation? What must remain yours because responsibility attaches to it?

FIGURE 2 | WHO SHOULD GO NVIDIA, AND HOW FAR?

Independent analysis. Categories describe where a company creates value; a DRIVE, Thor, or Hyperion relationship does not establish use of Alpamayo 2 Super. 'Own' means remain responsible for, not necessarily build alone.

How the map changes each game

Full-stack autonomy developers

Waymo, Tesla, Zoox, Mobileye, Baidu Apollo Go, Pony.ai, WeRide, Aurora, and their peers do not begin from a blank page. They own substantial proprietary drivers, data pipelines, vehicle integrations, and safety programs. Several also hold years of testing or operating history.

Replacing a mature core would impose integration, retraining, and revalidation costs while discarding the assets Alpamayo cannot supply. Their likely motion is **Use**: evaluate the model as a development tool around the production driver, not reset the production core.

Mobileye is the revealing hybrid. It is extending from system supply into a vertically integrated robotaxi business while continuing to sell its own compute, sensing, mapping, and safety architecture. It can use an open NVIDIA model and run a rival platform at the same time. Its scarce complement is the integrated system and the evidence behind it.

AI-driver platforms

Wayve, Waabi, Nuro, PlusAI, and Autobrains wake up to both an instrument and a rival. A strong open reference makes it easier for customers to compare them, and harder to charge merely for access to driving intelligence.

Wayve captures the optionality. It competes with NVIDIA at the driver-model layer, runs a Nissan prototype on Hyperion, and supports Qualcomm hardware as well. It can be a model-layer competitor and a platform-layer partner on the same day.

The defensible value for this group must live beyond access to intelligence: production behavior, transfer across vehicles and geographies, fleet learning, integration, economics, hardware portability, evidence, and lifecycle support. Their motion is usually **Use** plus selective **Build on**.

Automakers

Automakers already committed to parts of the DRIVE family have a shorter path into NVIDIA's factory. NVIDIA has publicly announced DRIVE-family relationships with Mercedes-Benz, Hyundai and Kia, Toyota, GM, Volvo Cars, JLR, Stellantis, BYD, Geely, and Nissan, with different products, scopes, and timelines. None of those announcements establishes use of Alpamayo 2 Super; they show pre-existing integration pathways.

For these companies, **Adapt** is the strategic center: turn a common foundation into brand- and territory-specific behavior while retaining vehicle integration, release authority, customer experience, and product responsibility.

Automakers whose announced driver-assistance paths currently point elsewhere can still **Use** Alpamayo as a reference or hold it as a second path. BMW, Honda, and Volkswagen Group are the visible examples. A credible alternative improves technical comparison and supplier leverage without requiring an immediate architecture reset.

Ambition changes the burden. In assisted-driving programs, where a human still supervises, the automaker must own driver monitoring, human-machine interaction, handback behavior, and road certification. In full autonomy, the system inherits the driving responsibility, and the operating burden expands to territory admission, fallback and remote assistance, incident response, public-authority interfaces, and institutional accountability.

The central question is the same at every level: which layers must remain ours because customers experience them, regulators judge them, or liability attaches to them?

Greenfield and sovereign programs

Tensor, the planned Foxconn, Foxtron, and Kaohsiung program, and NVIDIA's work with the Saudi AI company HUMAIN illustrate why blank-sheet and regional entrants may gain most. A coherent path through training, simulation, compute, sensors, runtime, and safety engineering can remove years of foundation work. Their likely posture is deep **Build on**, usually with an experienced integrator.

Greenfield has no legacy. It also has no history.

These entrants carry the largest Habitat deficit: no accumulated operating record, emergency relationships, institutional authority, or local trust. Someone still has to define the service mission and territory, integrate the vehicle, establish fleet procedures, assemble the safety case, and stand behind the result.

A head start is not a deployment license.

Mobility networks and marketplaces

Uber, Lyft, Bolt, and Grab can aggregate vehicles driven by different autonomy systems. Uber's announced NVIDIA plan is the deepest: an NVIDIA-software-driven fleet intended to span 28 cities by 2028, with automakers supplying vehicles and Uber contributing its network, demand, and operating experience. The details of local fleet operation will differ by market.

Their durable asset is admission and orchestration across many drivers: what evidence permits entry, how incidents escalate, how riders and cities receive answers, and how service continues when a supplier fails.

Alpamayo can become a common yardstick, and NVIDIA-powered vehicles can become another supply rail. The network's strategic motion is **Operate across**.

Suppliers, integrators, and sensors

For Bosch, AUMOVIO, Valeo, Magna, ZF, and the sensor ecosystem, open intelligence becomes more accessible raw material. The opportunity is to turn it into a warranted, vehicle-grade product. The threat is NVIDIA absorbing more of the surrounding integration.

What remains scarce is stubbornly physical: timing, thermal limits, diagnostics, cybersecurity, calibration, redundancy, manufacturing variation, failure detection, and years of support. Alpamayo 2 Super's released interface does not take lidar or radar; Hyperion's broader architecture does. That is why qualified sensor partners such as Aeva, Arbe, Hesai, OMNIVISION, and Sony benefit from the platform without being direct users of the model.

Open reasoning does not clean a dirty lens or prove that degraded radar will fail safely.

Development, simulation, and independent evidence

Foretellix, Applied Intuition, Parallel Domain, Voxel51, and their peers receive a tailwind. One open model can produce many derivatives. Each may run on several vehicle and sensor configurations, inside several territories, and change after every release.

The category is also splitting. Applied Intuition now sells an L2+ driver-assistance stack and a vehicle operating system, and runs its own data-collection fleet, which places it partly on the supplier side of the table it also serves. A strong open baseline may pull companies like it back toward tools, where independence is the product.

Model derivative × vehicle configuration × operating territory × software version = evidence surface.

That is the verification multiplier. Every cell in the multiplication is a configured system someone must prove.

This layer can integrate Alpamayo for labeling, critique, evaluation, and closed-loop testing. The harder product is coverage, the data evidence of how completely a system has been tested against the situations it will meet, which is what turns a pile of passing tests into a safety argument. Its deeper asset is independence. The more one supplier provides the model, compute, simulator, runtime, and reference architecture, the more valuable a referee becomes whose evidence survives a change of vendor.

Cities, regulators, insurers, and road authorities will inspect the same surface. They cannot permit or price “Alpamayo” in the abstract. They need the configured operation: the model and version, vehicle and sensors, territory and rules, operator and procedures, intervention history, and evidence after each meaningful change.

The referee can now hold the same open instrument as the players.

What the release sets in motion

Economics of this kind show direction long before they show timing. Incentives tell you where an industry leans. Adoption cycles, regulation, capital, and the timing of incidents decide when it arrives. Read what follows as pressure, and treat the calendar as its own open question.

The first pressure is multiplication. Making a derivative is now far cheaper than qualifying one, so region-specific, vehicle-specific, and weather-specialized models will probably appear faster than deployments can absorb them. Most will stay experiments. One marker would show the pull is real: before the end of 2027, a company outside NVIDIA's August 4 launch group discloses a supervised public-road pilot using an Alpamayo-derived model. Watch for it, and hold this essay to it.

As those derivatives spread, a benchmark score probably says less each year. It stays useful inside a development team. Outside one, it describes a model that has since been distilled, moved onto different sensors, and confined to a single territory. The likely end state is a score that opens the conversation, with deployment evidence carrying it.

The same spread reaches any company whose advantage was access to driving intelligence. Contracts and integration will slow its arrival, probably by years. The pressure still runs one way, toward a second scarce asset: production behavior, fleet learning, integration depth, economics, trusted validation, or operating history.

Responsibility will likely spread as the derivatives do. The license clears royalties and field-of-use limits, and it leaves every question of statutory, product, and operational responsibility open across the model creator, the fine-tuner, the integrator, the automaker, the marketplace, and the operator. Nobody can schedule the first serious incident involving an open derivative. Whenever it lands, it will make that chain visible in a way no contract has.

NVIDIA will probably have company at the factory level. Qualcomm, Mobileye, Huawei, and Horizon Robotics already contest parts of the compute, software, safety, and ecosystem stack. NVIDIA's advantage lasts while its pathway stays the easiest coherent route, and it erodes if customers resist dependence, if rivals open comparable foundations, or if regional rules favor local stacks. Automakers have run this movie before, with the dashboard.

CarPlay and Android Auto were adopted because drivers demanded them, the largest automakers kept building their own software to reclaim the interface, and GM went as far as dropping CarPlay from its electric vehicles. Expect a similar spread of postures here, over a similar span of years.

Speed remains the honest unknown. A hard incident could stall the whole sequence. A successful sovereign launch could pull it forward. A rival open release could reshuffle the factories entirely. I stake the direction, and I hold the timing loosely.

Abundance should lower costs and raise the technical floor. It also lowers the barrier to releasing a weak or poorly governed derivative. The public benefit depends on evidence, admission controls, and operating discipline keeping pace with the models.

The non-downloadable asset

After an incident, nobody stops at the model card. Investigators ask which version acted, on which vehicle and sensor configuration, inside which territory, after which update, under whose authority, with what interventions, and with what result.

An insurer pricing a robotaxi fleet needs the same record. So does a city renewing a permit, an automaker approving a release, or a board deciding whether to expand service.

The scarce asset is **consequence-linked operational history**: a verified record of what a configured system was permitted to do, what happened, where people intervened, and what the institution changed afterward.

This is why “data is the new moat” is too simple. Alpamayo and the surrounding factory also make labels, teacher outputs, reconstructed scenes, and synthetic scenarios easier to obtain. Consequence-linked history sits beyond their reach because it must be earned under real permission and real consequence.

Synthetic data can accelerate learning. It cannot synthesize a deployment record.

The more reproducible the intelligence becomes, the more valuable the irreproducible deployment history becomes.

The road is a lighthouse for Physical AI

Autonomous driving is not the first field to release open models for machines. It is an industrial-scale preview of what happens when commercially reusable intelligence enters a regulated domain of physical consequence, backed by a platform company selling the surrounding factory. The pattern is already traveling: NVIDIA is applying the same open license family to its GR00T models for humanoid robots.

Consider an autonomous cargo ship leaving Singapore and entering Rotterdam. Its navigation model may cross the border unchanged. Its authority cannot. Which remote operator has command? When does a port instruction override the planned route? Who accepts the handover between jurisdictions? What record must the insurer and port state inspect?

The same split appears on a factory floor when a changed robot model seeks permission to re-enter a live production line, and in an operating room when a surgical robot's proposed move conflicts with the clinician's judgment.

These sound like operational questions because they are. They are also where permission, insurance, responsibility, and adoption live.

The Habitat thesis meets the market

This is the shift the Habitat thesis was built to explain. Since 2024, in *The Case for Habitat* and the research published at talcohen.ai, I have argued that Physical AI is not deployed by intelligence alone. It requires a persistent environment of constraint, authority, record, and consequence. As intelligence becomes more abundant, that environment captures more of the durable value.

Alpamayo does not prove the entire thesis, and NVIDIA did not release it to prove my argument. It makes the economic mechanism visible at industrial scale. NVIDIA can make reusable intelligence and infrastructure more accessible because doing so expands the market for its factory.

The implications are immediate. An operator starts the consequence-linked record now because it cannot be backfilled. An automaker writes down which layers must remain under its control. An investor reprices any company whose only moat is access to intelligence. A regulator or insurer asks for the configured operation and its evidence, not merely the model card.

Alpamayo makes the choice visible. The intelligence can be downloaded this afternoon. The authority to deploy it, the record that earns trust, and the responsibility for what happens next cannot.

A better model, you can buy. A Habitat, you have to build.

Selected official sources

- NVIDIA Alpamayo 2 Super commercial announcement
- NVIDIA Alpamayo 2 Super model card
- OpenMDW-1.1 FAQ
- OpenMDW-1.1 license
- NVIDIA Alpamayo 2 Super launch
- NVIDIA Cosmos 3 announcement
- NVIDIA Physical AI Data Factory blueprint
- NVIDIA PhysicalAI Autonomous Vehicles dataset terms
- NVIDIA Alpamayo recipes
- NVIDIA DRIVE AV
- NVIDIA Halos for autonomous vehicles
- NVIDIA DRIVE Hyperion L4 ecosystem
- NVIDIA DRIVE Hyperion robotaxi-ready ecosystem
- Mobileye vertically integrated robotaxi announcement
- Wayve and Nissan on DRIVE Hyperion
- Wayve and Qualcomm collaboration
- Uber's planned NVIDIA fleet
- Tensor and Lyft

- BMW and Qualcomm automated-driving compute
- Honda and Renesas
- Volkswagen Group, CARIAD, and Bosch
- Volkswagen Group, Valeo, and Mobileye
- Foretellix, Parallel Domain, and Voxel51 NuRec relationships
- Applied Intuition and NVIDIA L2+ collaboration
- NVIDIA Alpamayo 2 Super unveiling at GTC Taipei (May 31, 2026)
- NVIDIA DRIVE Hyperion sensor ecosystem
- Waymo remote assistance
- Waymo independent safety audits

Company examples describe strategic positions and publicly announced relationships. They do not imply confirmed Alpamayo 2 Super deployment. "Own" means remain responsible for, not necessarily build alone.