



# Full Autonomous Mobility — Beyond the Single Stack

SDV · Robotics · AAM // Reactive · Predictive · Cooperative Perception Commons

*System 1, System 2, and the cooperative path across cars, robots (industrial + domestic), and aircraft.*

SDV → AIDV

ROBOTICS · industrial · lights-out · domestic

AAM → AIDA

► BRIEF MAP · 4-PART FRAMEWORK

**I THE THESIS**  
 SD → AID era shift · L4-L5 gap · communications substrate · fig.01 hybrid architecture (System 1 · 2 · commons).  
 pp. 2-3 · the\_thesis

**II THE DEBATE**  
 Three architectural schools at L4+ · three verticals on one topology (SDV · Robotics · AAM).  
 p. 4 · the\_debate

**III THE CLOSER**  
 Topology answer · six paths forward · Houston + lights-out manufacturing · validation commons · Vol 03 teaser.  
 p. 5 · the\_closer

**IV ENGAGE · THE PRACTICE**  
 Ronna-X scopes · indicative bands · single CTA. Omitted in screen-safe / candidate PDFs.  
 p. 6 · the\_practice · commercial only

# Full Autonomous Mobility — Beyond the Single Stack

BRIEF MAP · PART I OF IV · THE\_THESIS

System 1, System 2, and the cooperative path across cars, robots (industrial + domestic), and aircraft.

\$ hervé-oudin

circulation permitted with attribution




by Ronna-X

PART I THE THESIS


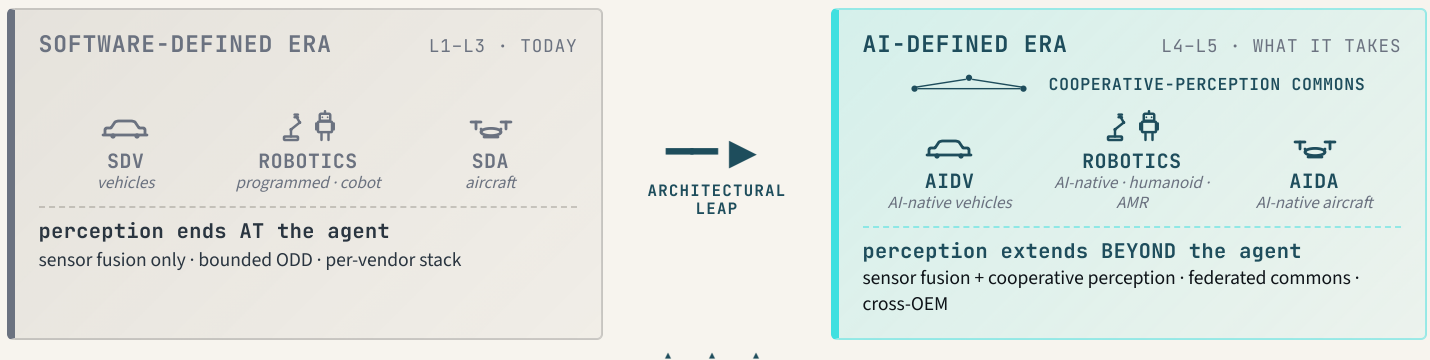
pp. 2-3 · the\_thesis

**Autonomous mobility is on a two-regime path.** Levels 1–3 are the **Software-Defined era**—sensor fusion inside a bounded ODD, shipping at scale across **SDV**, **SDA**, and programmed / cobot robotics. Levels 4–5 are the **AI-Defined era**—sensor fusion **plus cooperative perception**, emergent phenomena, and validation no single OEM can field alone. The leap is not "a better model." It is **Physical AI**, a hybrid **reactive + predictive** topology, and a cooperative-perception commons.

## where\_we\_are.l1\_l3 //Part I · Software-Defined era — sensor fusion alone, bounded ODD



**ODD** = Operational Design Domain — the bounded envelope (geography · weather · infrastructure · sensor stack) within which an autonomy system is designed, validated, and permitted to operate.



The boundary of perception ends at the agent. On-platform **sensor fusion** (camera · radar · LiDAR · IMU · GNSS) scales beautifully inside a bounded ODD: more data and compute compound, the long tail is manageable, and validation is per-scenario regression testing. The same logic governs **SDV** on supervised lanes, **SDA** under VFR in charted corridors, and programmed / cobot robotics in structured cells, lines, and homes. **This regime is shipping.**

L1–L3 ships inside bounded ODDs — mapped metro envelopes, factory cells, charted air corridors — with perception ending at the agent. SAE L4 in geofenced ODDs is a real, revenue-bearing class today; in this author's view it remains **bounded ODD automation**, not open-ODD generalization, until perception extends through a cooperative commons. The cross-lane shift is the same: **SDV** → **AIDV**, programmed robotics → AI-native, **SDA** → **AIDA** — one architectural leap, three modes.

The architectural leap above is what L4–L5 demands: perception that extends **beyond** the agent, carried by a cooperative-perception commons and the communications substrate beneath it. Page 3 diagrams **fig.01** — System 1 reactive, System 2 predictive, and the shared commons that couples SDV, robotics, and AAM on one topology.

L4-L5 forces cooperative perception beyond the agent. Fig.01 makes the mechanism explicit — System 1 reactive, System 2 predictive, and a cooperative-perception commons shared by every lane.

## what\_it\_takes.hybrid\_architecture //Part I · fig.01 · reactive · predictive · cooperative perception commons

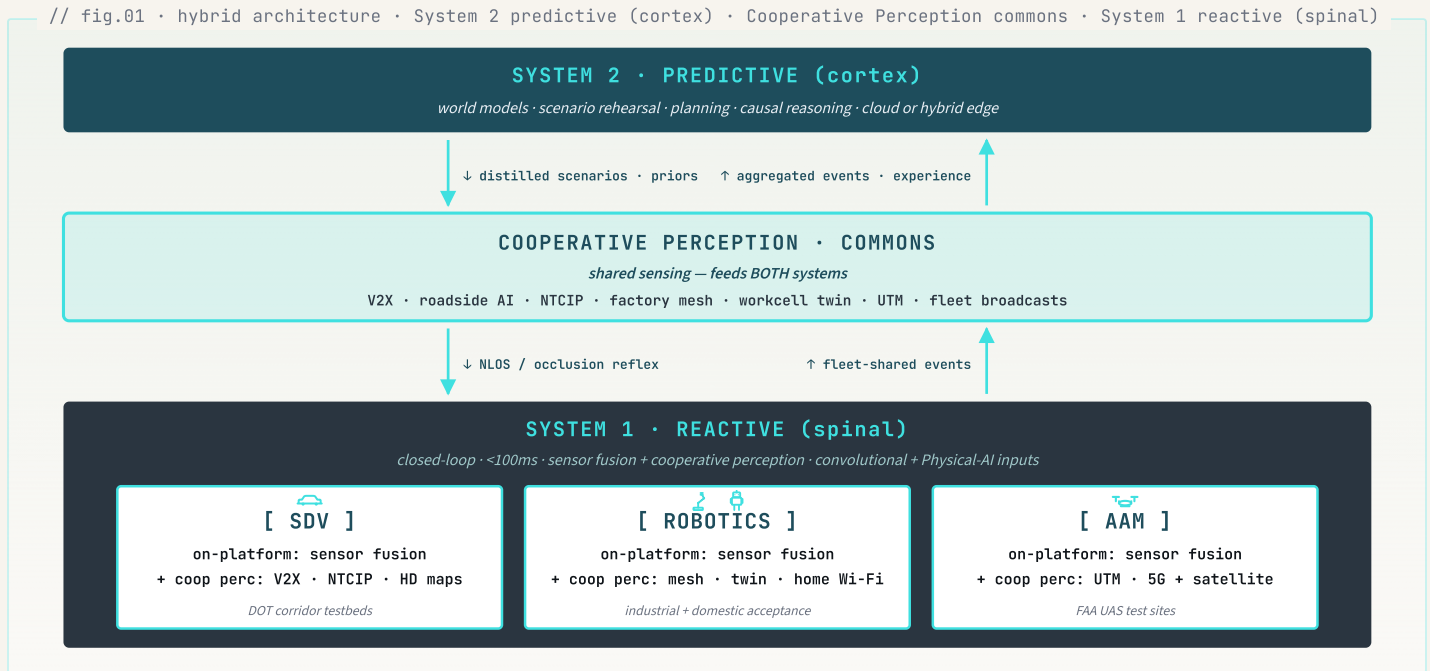


Fig.01 — the thesis. Cooperative perception is not an add-on; it is the middle layer that feeds both reactive execution and predictive planning. V2X, factory mesh, workcell digital twins, and UTM are not separate product categories — they are the commons interfaces that let System 1 and System 2 stay coupled while ODDs expand. The L4+ unlock is the coupling of all three tiers — across SDV, industrial and domestic robotics, and AAM.

## the\_l4\_l5\_gap //Part I · AI-Defined era — perception extends beyond the agent

Beyond L3, the long tail becomes emergent and causal: occluded pedestrians, intent, physics, novel scenes, counterfactuals. The architectural mechanism is the expansion of the perception boundary: sensor fusion alone is the L1-L3 substrate; sensor fusion plus cooperative perception is the L4-L5 substrate. For open-ODD autonomy, the gap is no longer primarily more parameters or more fleet data alone — inside geofenced L4 envelopes those still compound — but what no single agent can perceive alone. This is where SDV → AIDV, SDA → AIDA, and programmed robotics → AI-native all converge on the same need: Physical AI that predicts how the world evolves, fed by a shared sensing commons.

Worked example — unsignalized intersection, cross-traffic occluded by a box truck: onboard-only fusion sees the ego lane and brakes conservatively on uncertainty; commons-augmented cooperative perception adds a roadside camera track or C-V2X alert from the conflicting approach, surfacing the pedestrian or cyclist in the shared commons before System 1 commits. Same geometry — different perception boundary.

Physical AI — multi-modal, real-time, safety-critical AI that reasons about and acts in the physical world; spans autonomous driving, mobile robotics (industrial + domestic), advanced air mobility, and smart infrastructure.

> PART II PREVIEW · THREE SCHOOLS AT L4+  
01 Scale the recipe — L4 in geofenced ODDs today, not open-ODD autonomy · 02 Predictive world models — V-JEPA 2, GAIA-1, Cosmos-class causal reasoning · 03 Modular + cooperative — maps, twins, V2X/UTM inside ODDs. The industry debate is which school wins; the architecture answer is that open-ODD autonomy likely requires all three, coupled through the commons.

## three\_schools.l4\_plus\_debate //Part II · which architecture wins at L4+?

01 SCALE THE SAME RECIPE

TODAY

End-to-end nets, more parameters, more fleet data. SAE L4 geofenced-ODD automation is a real class today — mapped metro envelopes, campus and airport loops, warehouse yards, fixed shuttle corridors (US metro deployments at fleet scale · 2026) — revenue and miles inside the envelope, but in this author's view bounded ODD automation, not open-ODD / geographically generalized autonomy (industry classifies these as L4; open-world operation remains unproven). Bets the next leap is a quantitative continuation of the current curve.

02 PREDICTIVE WORLD MODELS

EMERGING

Self-supervised, generative, scenario-aware (Meta V-JEPA 2 · Wayve GAIA-1 · NVIDIA Cosmos, adjacent research). Designed for the causal reasoning L4+ demands. Earlier in the curve—and the most likely architectural unlock for the open-ODD, geographically-general L5 regime.

03 MODULAR + COOPERATIVE

SHIPPING · ODD-BOUNDED

Explicit perception / prediction / planning paired with HD maps, factory twins, aeronautical charts, V2X, UTM. Operates today inside ODDs. Scales slowly geographically but produces real revenue and real safety cases.

The honest answer at L4+ is probably not pick-one. Scale-the-recipe, predictive world models, and modular cooperative stacks each ship real value inside bounded ODDs — but geographically open autonomy demands coupling all three through the commons fig.01 diagrams: reactive execution on the agent, predictive rehearsal in cloud or hybrid edge, and shared sensing that feeds both.

## three\_verticals.one\_architecture //Part II · same topology · different infrastructure per lane

Three mobility modes · one architecture. Infrastructure differs by lane; the System 1 / System 2 / commons split does not.

[ SDV ]

Software-Defined Vehicles

maturity · V2X + corridor testbeds: shipping pilots

ODD: geofence · weather envelope · sensor stack Infrastructure: V2X (C-V2X; DSRC spectrum sunseting) · HD maps · NTCIP traffic controllers Validation: DOT corridor testbeds · Mcity · 5GAA · FHWA Standards: SAE J3016.

Canonical field examples: Mcity Test Facility (Ann Arbor) · 5GAA corridor demos · THEA Tampa CV Pilot · Texas TRUST V2X corridor (Greater Houston + College Station · TTI-led).

[ ROBOTICS ]

Industrial (line · cell · lights-out factory) + Domestic · Humanoid · AMR · Cobot

maturity · factory mesh + cell twin: production ODDs; home commons: early

ODD: cell · line · warehouse aisle · lights-out / dark factory (industrial) · home · care environment (domestic) Infrastructure: factory mesh · private 5G · digital twin of the workcell · edge-AI inference stack (on-prem, IP-sovereign, sub-100ms) (industrial) · home Wi-Fi · consumer edge AI (domestic) Validation: ISO 10218 + ISO/TS 15066 (industrial) · emerging humanoid + consumer-product safety regimes (domestic) Players: Aptronik · Figure · 1X · Boston Dynamics · Tesla Optimus.

Canonical field examples: Aptronik–Mercedes-Benz · Figure–BMW · 1X · Boston Dynamics–Hyundai · Tesla Optimus — plus lights-out factory adopters below (Foxconn · Siemens Amberg · BMW iFACTORY · Giga Texas).

[ AAM ]

Advanced Air Mobility · eVTOL

maturity · UTM + test sites: federated research; operator ops early

ODD: vertiport · corridor · weather + visibility envelope Infrastructure: UTM (Unmanned Traffic Management) · 5G + satellite redundancy · vertiport ground systems Validation: FAA UAS test sites · NASA UTM partner programs Standards: ASTM F38 · FAA Part 135 · JARUS.

Canonical field examples: FAA UAS Test Site program (7 federally designated sites) — NUAIR (NY) · Nevada · Alaska / Pan-Pacific (UAF) · Northern Plains (ND) · Virginia Tech (MAAP) · TAMU-CC Autonomy Research Institute (formerly Lone Star UAS Center, TX) · Choctaw Nation (OK). Plus Houston Spaceport — FAA-licensed commercial spaceport at Ellington — as vertiport / UAM corridor anchor.

ODD · Operational Design Domain — bounded operating envelope per lane (defined Part I · where\_we\_are.l1\_13) · Maturity tags on each lane label — infrastructure readiness varies; open-ODD generalization is unproven in every lane.

Three modes, one topology: infrastructure differs by lane, but the System 1 / System 2 / commons split does not. Part III maps what to build next — federated testbeds, interface contracts, and the validation commons that make open-ODD Physical AI economically shippable.

> THE TOPOLOGY ANSWER · REACTIVE + PREDICTIVE + A SHARED COMMONS  
SYSTEM 1 + SYSTEM 2 + COOPERATIVE PERCEPTION

**System 1 — reactive (spinal):** *closed-loop, <100ms, sensor fusion + cooperative perception integrated as one nervous system.*

**System 2 — predictive (cortex):** *world models, scenario rehearsal, planning — cloud or hybrid edge. Cooperative perception —*

**commons:** *shared sensing that feeds both tiers. The L4+ unlock is the coupling of all three—across cars, robots, and aircraft.*

Borrowing Kahneman's System 1/2 framing as inspiration, not a literal mapping — the engineering reality is reactive vs. predictive, with a third commons layer that has no analog in pure cognitive theory.

## six\_paths\_forward //Part III · industry actions · SDV · Robotics · AAM · Manufacturing

1. Fund **corridor-scale federated testbeds** — cross-state, cross-OEM, cross-mode (road + factory + sky).
2. Couple **cloud foundation / world models** with **distilled specialized edge models** — scale *and* specialize, not one or the other.
3. Standardize **cooperative-perception interface contracts**: V2X (ground), mesh + digital twin (factory), UTM (air).
4. Deploy **infrastructure-side AI on existing assets first** (CCTV, PLCs, roadside sensors, vertiport ground systems) — rip-and-replace last.
5. Adopt **non-intrusive validation attestation** across vehicle, robot, aircraft, and infrastructure domains — replay, not redeploy.
6. Share **SoS reference models** so the ecosystem converges on "good enough" Physical AI past L4 — across every mode.

## regional\_spotlight.houston //Part III · one metro · three lanes · factory corridor

> **REGIONAL SPOTLIGHT · HOUSTON** One US metro overlaps **all three lanes**: the FAA-licensed Houston Spaceport (vertiport + UAM anchor), TAMU-CC Autonomy Research Institute (FAA UAS test site), Texas TRUST V2X corridor (TTI-led), plus active logistics-robotics pilots across the Gulf manufacturing corridor. Natural canvas for a federated, cross-mode validation commons — same topology as the three vehicle lanes, anchored in a region that already has the freight, the energy, the airspace, the academic depth, and the regulatory permissive surface to host it at scale.

> **DEEP-DIVE · LIGHTS-OUT MANUFACTURING** Robotics' industrial side is shifting on three layers together — *agent*: programmed cobot → AI-native autonomous cell · *system*: **Industry 4.0** → **5.0** (data-driven → human + AI co-operating) · *network*: IIoT → **IIAIoT** (telemetry rails → intelligent agents on the rails). Convergent form: **lights-out / dark-factory** operations running on **sovereign on-prem edge-AI inference** stacks — IP cannot leave the plant and latency budgets cannot tolerate the cloud. Adopters at scale: Foxconn Shenzhen · Siemens Amberg · BMW iFACTORY · Tesla Giga Texas industrial AI lines.

*Different infrastructure. Same topology. The commons the sector still has to build is shared across all three lanes — and across the factory floor.*

## system\_of\_systems.validation\_commons //Part III · why validation is infrastructure, not a vendor moat

Cooperative perception is the architectural commons — the shared sensing and intent layer that turns long-tail edge cases from a per-vendor problem into a shared dataset. No single OEM, robotics company, or eVTOL operator can build the testbed surface that L4+ economics demands alone. The pattern is structurally **federated**: cross-state, cross-OEM, cross-mode (road + factory + sky), governed by interface contracts — not by vendor stacks.

**Evidence the pattern already exists:** Mcity and 5GAA corridor programs have spent a decade turning cooperative-driving interfaces into repeatable test infrastructure; NASA's UTM partner ecosystem does the same for low-altitude airspace — federated governance, interface contracts, shared scenario libraries — before any single OEM or operator owns the full stack.

**Physical AI** is what ImageNet played for vision — a shared substrate that lets every player build on top instead of starting from scratch. The validation commons is the same idea, applied to cyber-physical systems. The ones who win past L4 won't be the ones with the biggest model. They'll be the ones embedded in the deepest federated validation commons.

*The next era belongs to corridors, factories, and vertiports that treat **cooperative validation** as infrastructure — not just to platforms with the loudest model.*

// next → Vol 03 · **Power Sovereignty as Physical AI** — same topology applied to the smart grid (SDG → AIDG) and the AI-DC build-out: where the load lands, where the compute sits, where the validation commons gets governed. In drafting · Q3 2026.

# From Brief to Build – The Ronna-X practice.

The same architecture-first lens applied above — now deployed inside your team.

**Ronna-X** is a Houston-anchored research-and-build practice for full-autonomous-mobility programs. We help SDV, robotics, AAM, and infrastructure-side teams turn architectural complexity into shippable strategy: **federated testbed design · cooperative-perception roadmaps · partner ecosystem maps · GTM playbooks · pilot architectures**. Independent of any vendor stack. Operator-side voice. Engagements are sprint-based with the option to extend into fractional advisory.

## engagement\_formats // four standard scopes · price bands indicative · all scoped per engagement

▸ 01 · ENTRY 1 day · on-site or remote

### EXECUTIVE BRIEFING

A one-day intensive for an exec or program team that wants the cross-vertical architecture made *actionable* in their context. Structured Q&A, scenario walk-through, ecosystem orientation.

DELIVERABLES

- 1-day workshop (½ day intro + ½ day deep-dive)
- Tailored 8–12 page readout PDF
- 2 follow-up office-hours sessions (30 min each)

---

USD 6K – 12K FLAT · TRAVEL SEPARATE

▸ 02 · DIAGNOSTIC 2–3 weeks

### DIAGNOSTIC SPRINT

A focused architecture-and-GTM review for a specific program (SDV pilot, V2X corridor, vertiport operator stack, humanoid deployment). Where are the moats? Where is the topology debt? Who should you partner with?

DELIVERABLES

- Architecture & topology gap analysis (~20 pages)
- Partner ecosystem map + named shortlist
- GTM readiness scorecard with 90-day actions
- Final readout to leadership (60–90 min)

---

USD 20K – 35K FIXED-FEE · 2 WORK-STREAMS

▸ 03 · BUILD 6–8 weeks

### BUILD SPRINT

Full GTM-to-pilot scope — the end-to-end output a program team can take to commit. Strategy, playbook, partner contracts framework, pilot architecture, validation plan, KPI framework. Where Ronna-X most directly bridges thesis to deployment.

DELIVERABLES

- Complete GTM playbook + ecosystem map + value-prop matrix
- Pilot architecture + cooperative-perception interface spec
- Named partner contact list + warm-intro execution
- KPI framework + 12-month operating cadence
- Weekly 90-min working sessions throughout

---

USD 60K – 120K MILESTONE-BASED · 3 WORK-STREAMS

▸ 04 · EMBEDDED monthly retainer · 3-month minimum

### FRACTIONAL ADVISORY

Embedded at ~1 day/week with the product, strategy, or program-management team. On-call for exec working sessions, deal reviews, partner conversations, board-prep memos. Best for programs already in motion needing senior architectural judgment continuously.

DELIVERABLES

- ~1 day/week embedded + on-call slack/email
- Monthly board-ready memo on program posture
- Co-author of strategy docs, partner pitches, RFP responses
- Quarterly off-site facilitation

---

USD 10K – 18K / month 3 MO. MIN

▸ WHERE CLIENTS HIRE RONNA-X

- SDV · V2X corridor strategy
- Cooperative-perception roadmaps
- Federated testbed design
- AAM · vertiport ecosystem
- UTM partner architecture
- Industrial humanoid pilot framing
- 5G + edge-AI program management
- Houston / Gulf market entry

▸ HOW A CONVERSATION STARTS

<p>01 <b>DISCOVER</b></p> <p>30-min intro call. We map your program against the System 1 / 2 / commons frame and identify whether there's a real engagement fit.</p> <p>no charge · no slides</p>	<p>→</p>	<p>02 <b>SCOPE</b></p> <p>Within 48h: a written 1-page Statement of Work — outcome, timeline, deliverables, milestones, price. You decide. No follow-up pressure.</p> <p>48h SLA</p>
		<p>→</p>
		<p>03 <b>SHIP</b></p> <p>Sprint or retainer kicks off. Weekly checkpoints. Working sessions with your team. Concrete artifacts at every milestone.</p> <p>milestone-based</p>

```
$ ./initiate-engagement.sh
```

**If this brief mapped to a real program on your side – let's talk for 30 minutes.**

---

EMAIL [herve@ronna-x.com](mailto:herve@ronna-x.com) LINKEDIN [linkedin.com/in/herveoudin](https://www.linkedin.com/in/herveoudin) WEB [ronna-x.com](https://ronna-x.com)

// 48h response · NDA-friendly · no sales sequence · no pitch deck · 15+ years operator-side · Test & Measurement · Digital Twins · Cyber-Physical · 5G · V2X · Mcity ITS PoC · 5GAA · smart-grid GTM · Houston metro